

GERBRAND

So what I want to do today is finish up Monte Carlo simulation and go into a little more advanced things in Monte Carlo simulation and then also go into free energy integrations and other coarse-graining methods so when direct simulation won't work. If you remember, I introduced Monte Carlo simulation as a way of importance sampling and contrasted it against simple sampling. What I want to show you today, there's actually quite a few things in between that, in some cases, can be practical.

CEDER:

Remember the idea of simple sampling is that you sample randomly and then you weigh with a probability that's correct relatively, where the states have their correct relative probability to one another. Importance sampling is sampling with the correct probability and then just averaging the quantity you sample, but you can actually sample with things in between. So here you're sampling with a probability proportional to that exponential of that Hamiltonian.

Here you're sampling with no Hamiltonian and therefore, you have to correct the probability later. You can sample with any Hamiltonian H_0 , which is not the true Hamiltonian of your system, and then correct for it in the probabilities. So if you sample with H_0 , then the states you sample have to be corrected with the relative probability of that state, the relative probability that you would get with the correct Hamiltonian versus what you get with the Hamiltonian which you decided to sample. OK?

So this is the Hamiltonian which we decided to pick, the states that go in your sample, and then this is the probability correction. Sorry. And so whenever you sample with what's not the proper Hamiltonian of your system, people call it non-Boltzmann sampling. Why would you want to do non-Boltzmann sampling? Well, it can be relevant if the particular quantity you're after in your system is essentially not determined by the relevant thermodynamics states.

Remember, I introduced this by say, what if you're looking at, say, average energy or average volume or average magnetization, then what you really need to know is what's the energy of the states that the system spent most of its time in? But what if a system spends only a small amount of its time in certain states, but those have the relevant property that you want to sample? For example, let's say that this is a phase space, this [INAUDIBLE].

So with importance sampling, you'd be drawn towards these states. OK? I don't know, maybe the ones that live here are optically active or something like that and the other ones aren't and you want to somehow get a lot of information of the optical activity of the material, you may want to build a Hamiltonian that drives you towards these states. OK? And as long as you then correct for the proper probability, you'll get a proper ensemble.

It's one that's just much more efficiently biased towards the regional phase space where you want to get. So that's non-Boltzmann sampling. Another obvious way is you can use it to sample phase space more efficiently. If you had a phase base with a lot of local minima, you may want to define a new Hamiltonian that looks like this. And this is especially relevant if your Monte Carlo has some form of dynamics in it.

So in the blue Hamiltonian, the true Hamiltonian, it may be very hard to get out of this minimum into the next one. Whereas if you raise the potential well, OK, in the red Hamiltonian it's going to be much easier to get out of it. OK? So essentially your flattening your phase space with the new Hamiltonian, so it's going to be much easier to get out of local minima and then you can correct for that probability.

There's a lot of algorithms these days built on this idea, not just in Monte Carlo, but there are molecular dynamic schemes, there are all kinds of schemes that are built on this idea of lifting up the potential wells and then correcting the relative probability or the relative vibrational frequency or the relative time you spend in each potential well. OK. You can also do non-Metropolis Monte Carlo. This gets even a little more odd.

Remember that we defined an acceptable Metropolis as one where the a priori probabilities were equal. So basically in the Markov chain, the rate at which I pick the i state from the j one as a potential next step in the Markov chain is symmetric. It's the same as the rate at which I pick the j state from the i state. j to i and i to j is the same. You don't necessarily have to do that. You can actually make these rates-- so these are the picking rates, the rates at which I try one state from the other, non-symmetric, and even more, you can make them dependent on the Hamiltonian.

And why would you want to do that? I'll show you an example in a second. It's a way of forcing systems downhill in energy space much faster. So as long as you're correct in your detailed balance argument, you'll be OK. So the important thing is the detailed balance criteria because the detailed balance criteria is the one that ultimately ensures that you sample phase space or that you weigh states and phase space with the correct probability.

And so as long as you correct here-- so if you changed the W_0 , you can still get satisfied detail balanced by essentially changing what the acceptance rates are. $P_{I \rightarrow J}$ and $P_{J \rightarrow I}$. You can write out what they are. Essentially the ratio of $P_{I \rightarrow J}$ to $P_{J \rightarrow I}$ is the factor we had before, but now you correct by the relative weight, which we did the picking of the states. OK? So in the end it's a fairly trivial correction.

I haven't seen too many examples of this, but one is this one. It's force-bias Monte Carlo, which almost looks like a hybrid between Monte Carlo and molecular dynamics. You know, I've given you several times this example of a liquid. How would you sample a liquid?

In Monte Carlo, you just could pick random displacements of an atom within a given range. One way to actually drive the system to low energy faster is not take a random displacement of the atoms, but take one that's biased along the force on that atom. If you move an atom down with force, you're obviously going to lower the energy mathematically because the force is the gradient of the energy. So you could take a displacement factor that has a random component and that has some component that depends on the force.

Actually if you only do this, that's actually how a lot of static relaxation schemes work. For example, in density functional theory you calculate the forces on atoms and then you relax them assuming a certain spring constant typically. But because you still have a random element, this is probably somewhat better described as mixed dynamics Monte Carlo methods. And if you would like to read more about it, I think this was one of the first papers that kind of introduced them.

So what I want to do is a couple more case studies and then show you some of the complications that can arise when you do difficult systems with Monte Carlo and see how we solve them, and that'll lead us into free energy integration. So this is sort of a classic paper on-- it's from 1985-- using Monte Carlo to study surface segregation in copper nickel. It uses the embedded atom method as a potential so, one of the reasons I picked it because it integrates a lot of the things you've seen in class.

So the idea was to take random copper nickel alloys and see which element segregates to the surface and then see how the segregation pattern away from the surface is, because there's been a lot of discussion about that in the literature, especially around that time. So what they do is they set up, just like you would probably, supercells. So I think they either had 24 or 24 or 48 layers, and then with a certain width.

So here's your supercell. Again, you can define your Hamiltonian multiple ways. You could just seed the system with a particular concentration and then you have to decide which kind of Monte Carlo moves to allow. A copper nickel form [INAUDIBLE] solid solutions. So what kind of moves would you do to equilibrate the system if you had to set this up yourself?

So again, essentially you would seed the system, probably with some concentration of copper and nickel except that there's a lot more atoms. So how would you do a Monte Carlo on this? Remember that you're going to have an embedded atom Hamiltonian essentially, so it's one that you can fairly rapidly evaluate.

So in the end, it's a very simple energy function. It's a bunch of pairwise sums and then you just stick the result in a function and you're done. It's not like quantum mechanics. You have, essentially, a Hamiltonian you can abuse, so it's important to realize. So how would you do Monte Carlo on this? This is the one thing-- Monte Carlo is about choosing your perturbations, the rest is automatic after that essentially. It's just doing it a lot of times.

So how would you equilibrate the system? It's too early in the morning, you're all still equilibrated. So first of all, you need to segregate some species to the surface assuming there will be segregation. So how do you do that? Again, you could do diffusive hops. Diffusive like hops, so do canonical like interchanges.

You could interchange them nearby or you could interchange them far away. I got to stick to my colors here. So you could interchange positions of copper and nickel, but again, the slightly faster way to do it is to do it grand canonically, to define actually a Hamiltonian that's the energy minus μ_{copper} , n_{copper} minus μ_{nickel} , n_{nickel} . OK?

And since you're actually going to be interchanging the identity of atoms, n_{copper} and n_{nickel} are not independent variables. You keep total number of atoms constant. So you can actually write this as $E - \mu_{\text{copper}} n_{\text{copper}} - \mu_{\text{nickel}} n_{\text{nickel}}$, since any change of a nickel is compensated by a change of copper. So we have essentially only one controlling chemical potential, which is the difference of the two species chemical potentials. OK. Yes, sir?

AUDIENCE: If a [INAUDIBLE] in a specific ratio, nickel or copper, [INAUDIBLE] very long to find the correct value of [INAUDIBLE]?

GERBRAND Yes.

CEDER:

AUDIENCE: [INAUDIBLE]

GERBRAND Yes. Yes, so you're controlling the external chemical potential, so you don't control the composition. True. What
CEDER: people sometimes do is that-- what's of interest to you is the bulk chemical potential, that's the source for the surface. You first do bulk simulations and then you roughly know what chemical potential corresponds to what compositions, because it's actually a little tricky.

While energetically often, the center of the system does not influence the surface much so you think you have convergence. You still deplete it as a reservoir as you segregate one species to the surface, so that's sometimes why you actually need more bulk in the system. But yes, what you say is true. If you're only interested in one composition it can be better to stick there with that. OK.

But let's say you do these exchange moves. Do you do it on a lattice? Do you do it in sort of a continuous space? I mean, how would you do it? This is an FCC solid. This is an FCC solid solution. So you could if you want to do it on a lattice, but would you? So you know what I mean with a lattice? You have pre-defined positions where the atoms can sit.

Somebody has to be able to say something. I'm just going to stand here. And it's not good because you're being taped for OCW, so they're going to think these MIT students have nothing to say. Somebody in Western Africa sees this they're going to go like, why did they let these people in at MIT? So what else would you add?

First of all, if you do it on a fixed lattice, think of what you end up with thermodynamically? What have you sampled? If you're on a fixed lattice, you've sampled the configurational entropy, essentially the possible distributions of copper and nickel on these fixed lattice sites. But if one atom is bigger than the other, for example, then they don't sit on a fixed lattice, then atoms will relax away from lattice positions.

You will also not sample vibrational excitation because you never go away from your lattice position. So you could add small displacements to the perturbations. So remember, you do the sort of big into the chemical interchange as you interchange copper to nickel, but then to that you could add small displacements of the atoms so you not sample deviations from their equilibrium positions.

And if you're done them, then you'll have both configurational and vibrational entropy. So remember that by the possible perturbations you pick for your system, you are essentially defining the phase space you sample. So you're defining this phase space you integrate over and that's essentially telling you what forms of entropy you allow in your system.

Because, you know, what forms of entropy is the same as what forms of disorder or excitations I allow in your system. OK? So the cool thing about it is that you can look at the effects of both, and I'll show that in a bit. You could actually say, first I'm going to do it on a rigid lattice and I'll see the result, and then I'll also include displacement away from the rigid lattice and you can look at the difference, and that's essentially telling you what the effect of vibrational excitations are on your thermodynamics.

OK, so here's the result. You know this was 1985? You'd be amazed how easy this is to do now if you wanted to do something like this yourself. With a Hamiltonian-like embedded atom it really goes very fast and you could equilibrate systems like this in no time. But here's the result. So this is a function of the bulk concentration of copper.

This is the amount of copper in the first and second and third layer. So the green curve is the first layer. So clearly, copper segregates to the surface in this material. So as you can see, even for 20% copper in the bulk you have an almost perfect copper layer on the surface. I think this was actually a 111 surface.

Copper in the second layer, interestingly, is depleted. So it's actually below the average bulk concentration. And then copper in the third layer essentially starts tracking the bulk concentration. So why do you think the copper in the second layer is depleted? Because first of all, why does copper go to the surface?

It has a lower surface energy, so this is not rocket science. Most of the time the thing with the lower surface energy goes to the surface. OK? But why is depleted in the second layer? Because I sort of think, well, this is the lower surface energy. Even in the second layer it's not perfectly bonded, so maybe it still wants to enhance its concentration there but actually depletes it.

This is not atypical, it's very common surface concentration. It's actually because copper and nickel in this system have an ordering interaction. So it's the fact that the first layer is almost pure copper, but the second layer, for chemical reasons, really wants to be nickel. OK?

And this is a pure surface effect. Copper and nickel in the bulk are actually at low temperature phase separating, but in the surface because of strain effects they form an ordering interaction. And so the second layer wants to be nickel because the first layer is so much copper. And so you often see this in compound forming systems that you have damped oscillations in the concentration away from the surface.

OK. So I want to show you one example of something that gets much more complicated, and you'll see how you start running into problems with Monte Carlo. So I've shown you a little bit how you would detect phase transitions in Monte Carlo. You'd look at things like the heat capacity especially, particularly powerful for detecting second order transitions.

For first order transitions, you look at discontinuities in things like the energy, things like the relation between μ and see chemical potential and concentration, and you'll pick that up. The problem you'll face is often that Monte Carlo systems, just like real systems, often face significant hysteresis. So here's an example from quite a few years ago of a fairly complicated lattice model Hamiltonian. This, again, was on an FCC lattice.

This is the palladium vanadium system, OK? Phase diagram. This is palladium on this side, this side is vanadium, and this is a lattice model with a bunch of interactions in the Hamiltonian. A nearest neighbor pair, second neighbor pair, third neighbor pair, and so on. Four body interactions, four body here, three body here.

It's a fairly complicated Hamiltonian that tends to give you a lot of local minima and how would you get this phase diagram? Well, what would you do? Again, to equilibrate faster you wouldn't keep the concentration fixed, you would scan the chemical potential and evolve in concentration from one end to the other. OK? Let me actually do that on the next slide.

So you have some chemical potential that gives you the difference in palladium and vanadium amounts or the concentration. OK. And so you would plot something like-- so you would have a driving chemical potential and you would plot the average spin, we'd say is a concentration of vanadium. And let's say you scan at this temperature, so what would you see? If you start here, you would go through a solid solution regime.

So you would see your concentration go up when you hit the two phase region. OK? The chemical potential is constant there, so that means at that chemical potential you have a discontinuity in the concentration, so this would go straight up. Then you would go in a single phase region. OK? So in a single phase region, the chemical potential changes rapidly with concentration. So there, you actually tend to be kind of flat.

It's not quite flat, but it tends to just be sloped a lot less. Then you would, again, go in to a two phase region. So you would, again, form a step. You would go into a single phase. So we basically see this kind of behavior, and where you have these discontinuities in concentration, you would know that there's a first order transition.

You could also plot the energy, and the energy is discontinuous as well. The internal energy system is discontinuous at a first order transition. You don't necessarily see things in the heat capacity. Now what happens in reality? If you think about it, if you come from this end so you have a solid solution, you go through two phase region, and you have to form this compound, which is called nickel [INAUDIBLE].

So often in Monte Carlo, just like in a real system, you have nucleation problems. If you have to form a phase or an arrangement that's extremely different from the host from which it forms, it will often just not nucleate, and well, what happens is that you will overshoot. So you will actually push the disordered phase way too far out of equilibrium and then at some point, you will be so far out of equilibrium that you literally shoot into the ordered phase.

And then what happens when you come back? You have hysteresis the other way. So you'll overshoot the order phase sometimes and disorder, although disordering is a little easier to do. So just like in real systems at strong first order conditions you'll see a lot of hysteresis, and it's often very hard to get rid of. I think some of you may have noticed that if you did the molecular dynamics lab.

Seeing first order transitions in a molecular dynamics lab, you also are often plagued with hysteresis. I mean, if you study melting, for example, it's very asymmetrical hysteresis. If you heat up a solid, it's pretty easy to melt it. If you cool down a liquid, it's about impossible to nucleate the solid unless you force it. So there you have enormous amounts of hysteresis. So the point of this is that to study strong first order transitions, sometimes the best way to find their transition temperatures or concentrations is not by direct simulation.

Essentially, what you need to do is go through a thermodynamic route. Try to extract free energies and find where they cross, because first order transitions are defined by where free energies of the two phases intersect, and that's by far the most accurate way of detecting phase transitions. There's a real problem here is that, how do you get free energy?

The fundamental difference between free energy and energy is that energy is an average quantity and free energy is not. See, energy is the average of something that is defined in the microscopic states. For every microscopic state I go through, I can define an energy, and then the internal energy, the thermodynamic energy of the system is just the average of that quantity.

For free energy or entropy-- which are essentially the same because if I know the entropy, I know the free energy-- it is not the average of a quantity that's defined in the microscopic state. You know? If you put the atoms in a fixed position somewhere with some velocity, that's a microscopic state. You cannot define the entropy of that. The entropy is a property of the ensemble as a whole, not of a given macroscopic state.

So entropy and free energy cannot be obtained as averages. They're actually integrals. You can see that the entropy is a sum over all phase space of this quantity, $P \log P$. You can write the free energy same way as an integrated quantity over all of phase space of this quantity.

If you rewrite it, you sort of see it better. You are averaging something. If you look at this, you are averaging something because you have a probability here, but what's the quantity you're averaging? It's essentially the free energy itself. So it's a quantity that's flat in phase space. OK? So that makes it extremely difficult to sample.

Actually this is not a mathematical problem, this is a physical problem. It's just the same as what happens in nature. You can measure energy, you can measure volume, you cannot measure free energy. There are no free energy meters because it's an extensive quantity that that is determined by the whole ensemble. You only really get free energies ever indirectly. You can get free energy difference, but you get free energy by integrating lower order quantities.

OK, so you can write it as an average but it's a little misleading. You can actually write it as the average of this thing here. If you can calculate the exponential of beta the Hamiltonian and average that over phase space-- and notice I didn't make an error there. I did not drop a minus sign. It's the exponential of the positive beta times the Hamiltonian. If you can average that quantity, you can show that you can actually have the free energy.

And the proof is given here, it's sort of an almost trivial proof. But you see that's kind of problematic. The Hamiltonian is an extensive quantity, so it scales with the size of the system. So you're taking the exponential of something that's extensive, so that gets very big. So first of all, you can only do this for finite systems and for systems that are really, really small. Because otherwise, essentially, that quantity, the exponential of beta the Hamiltonian, the difference between two states becomes excessively large as your system size gets bigger because that quantity in the exponential is extensive. OK?

You see, if I'm calculating the Hamiltonian difference, say, between two phases, the Hamiltonian value goes in there as the extensive quantity. It's not the normalized one. It's not the one, say, per unit cell or something like that. So that energy difference is infinite in the extensive limit, OK? So even if two phases only are one joule a part per molecule, per mol set, one joule per mol. In the extensive limit, they're still an infinite energy apart. OK?

So you can't practically actually sample that, so how do people get free energies? Well, there are hundreds of papers on free energy integration and the reason that there's hundreds of papers is that it's such a difficult thing and everybody claims to have the magic potion to do it. The first thing to realize is that you almost never need free energy, you always need free energy differences between two things, and that's a powerful statement because that's a lot easier to do as you'll see.

The second thing is that you probably shouldn't believe half of what you read in papers. I used to track that field and everybody said, oh, I have a great method to get free energies out of a single simulation, because that's sort of the problem. You'll see in a second that the way we get free energy is that we have to do Monte Carlo at a lot of different points to get the free energy at one point.

There's tons of papers that write that you can do it with one simulation. Well, either implicitly they do a lot of simulations within that one and just call it one, or you can do it in limited cases if you know a lot about the form of your phase space. So if you have extremely simple Hamiltonians-- if I give you a nearest neighbor Ising model, the magnetic model which is the nearest neighbor interaction, essentially the amount of excitations out of the low energy states there is very finite.

Like I said, you could flip one isolate, spin, you get a times j . So you could almost numerically start writing out what the free energy becomes. So in very simple models, if you know the form of the phase space, you can actually get towards free energy models. In general, it's pretty much an unsolved problem. And the way we practically get it is with three types of methods.

One is free energy integration, and I put lambda integration under there. I'll show in a second what it is. And the second one is overlapping distribution methods, which is slightly less important, especially for solids. So it's really only 2 previous others, which I used to cover and I don't even do anymore now. OK.

Let me first show you overlapping distribution methods, which I'm less familiar with because I never use it, but the idea of it is quite simple. So if you want to know the free energy difference between two states, remember that the free energy is $KT \log$ the partition function q . So the delta, the difference is the log of the ratio of the partition function, OK? You are with us?

So you can write out what that is. The partition function is the sum over all the states of the exponential of minus beta the Hamiltonian. If I multiply this by 1 -- and I'm going to multiply this by 1 , then I'm going to write 1 as the exponential of beta H_1 . Sorry. Got to be consistent here.

H_1 nu times exponential minus beta H_1 . OK, so that's one. So then I can collect the terms here, and what I get is that I get I still sum over all the states. I get the exponential, the Hamiltonian difference to n_1 weighted by the probability of this state in the ensemble of Hamiltonian one. This is essentially the probability of that state is the exponential minus beta H over the partition function taken with Hamiltonian 1 .

So what have I written here? The exponential of the probability of that state weighted in Hamiltonian 1 of this quantity. So essentially, what I'm averaging is the exponential of minus beta the Hamiltonian difference between state two and one, but I average it in the ensemble of one and that gives me the free energy difference. The reason that's called overlapping distribution methods is that-- let me show you that in a second-- you're trying to say something about state two, but you're sampling in the ensemble of one.

So the only way this is ever going to work is if one and 2 are not too far apart so that the states that are relevant for two are also sampled to some extent when you're in one, and that's why it's called overlapping distribution method. So if I sample in one, essentially let's look at the energy. I sample energy states around the average with some spread, spread could be the heat capacity.

If I'm in two I do the same thing around the average energy of two. And essentially what I'm doing is I'm integrating things for ensemble two just by the way I walk through ensemble one. And so it's a relatively elegant way of doing it, and the key aspect is that the states in the two ensembles are the same. The ensembles are the same, so the accessible states are the same, it's just that you weigh them differently.

So this is almost just like non-Boltzmann sampling. I Boltzmann sample for ensemble one, but you could say I non-Boltzmann sample for ensemble two and I correct the probability and this is how I get the free energy. So you can already see when this is not going to work. This is not going to work when your states are too far apart. OK.

By far the most used method to get free energy is sort of a trivial one in some sense, it's free energy integration. But it's the one that always works if you put enough time in it and it starts from this kind of trivial idea that the difference in a quantity is the integral of the differential, which this is why you come to MIT to learn this. So if I can actually sample this, I may be able to get at the quantity simply by integrating that.

And why is that important? Because if A is a free energy or an entropy, the derivatives of free energies and entropies are things that can be sampled because they tend to be either averages or fluctuations. For example, for the entropy, you wanted the entropy difference between two states so you integrate the derivative of the entropy. Well, the derivative of the entropy is the heat capacity and the heat capacity you can get from Monte Carlo. The heat capacity is essentially the fluctuation of the energy.

So if you want to know the entropy at a given temperature, you start from some reference temperature where you know the entropy or you just fix it to some value if you want to reference everything to the same thing and you integrate the heat capacity. What do you take as reference states? Well, again, you could just integrate between two states and get the entropy difference. Often you start from 0.

If you start from 0 temperature in models with discrete degrees of freedoms like the Ising model, the spin model, the entropy is 0 at 0 Kelvin so that's an easy integration state. In some cases, you can also find the entropy at infinity because in infinity, your phase space is random, the probability distribution is flat, and so sometimes you can get analytically the entropy there. So these are all proper reference states.

So here's an example. This is, again, our very simple 2D square magnetic Ising model. You would essentially integrate C over T , so you'd essentially integrate under this curve from 0 up. The way you practically do it is that you wouldn't start from 0, because the reason is you're integrating. So you're integrating C over T and that integral, in reality of course, converges as T goes to 0. But numerically, it will never in your simulation because T goes to 0 and you force T to 0 in your simulation.

That's a well defined number, but the heat capacity will not go to 0 just because of numerical noise. In reality, it should be really 0, but what will actually happen is that because of some minor noise and fluctuations, we will get non-zero heat capacity. So you divide by a number that gets exceedingly small as you go to 0, and so your integral will blow up numerically.

So what you typically do is you look at your heat capacity and say, well, below 0.5 or 0.75 I essentially have no heat capacity. That means all the way from 0 to there I have essentially no entropy, and you just start integrating from 0.5 or 0.75. If you want to be more accurate, there are ways of analytically writing the heat capacity from 0 to low temperature by things like low temperature expansions.

Essentially from here to about here, all that would ever happen is single spin flips. So you'd have this ferromagnet sitting there, once in a while one spin would go, [FAST-SOUNDING WHISTLE]. So they're single excitation so you can write out what the partition function pretty much looks like. It's the groundstate energy plus just the first excited states. So you can write out analytically what the entropy is up to that point and then integrate from there.

Practically you can integrate a whole bunch of other variables. The one that if you integrate from infinity that's very practical is to use the Gibbs-Helmholtz relation. Gibbs-Helmholtz relation essentially tells you that the average energy is the derivative of the free energy over T with respect to $1/T$. This is actually a generic relation. If you think here the average of any Hamiltonian is actually the derivative of its corresponding free energy over T with respect to $1/T$.

OK. So if you do this in a canonical system, then you're Hamiltonian's like the energy minus μ times some concentration. So then that's the average quantity you would get here. Why is this useful? Because essentially now you're integrating in $1/T$, so in β , and that quantity is 0 when you're at infinite temperature. So β is 0 and T equals infinity.

So now this is an easy way to integrate from infinity. Why do you want integrate from infinity? Like I said before, sometimes that infinite temperature, you know the free energy analytically, because everything is totally random and so something that's a useful integration state. Another one that's quite practical is simply integrating in composition, and that's why I sort of shown this on this phase because it shows the three major ways of integration.

If you come from low temperature, you integrate the heat capacity. If you come from high temperature, you use the Gibbs-Helmholtz relation to integrate the average Hamiltonian, so the energy most cases. If you come from the sides, you integrate in composition space. So you integrate essentially $\int \sigma d\mu$, and the reason is that the composition-- let me write in regular thermodynamic variables-- is essentially derivative of free energy with respect to the chemical potential.

So when you integrate that-- I've written composition here as the average spin in a spin model, so you integrate $\int C d\mu$, and that gives you dF and so that's essentially what this here is. Why can you integrate from the sides? Well, if you have pure 1-- so you call this A or B. If you have pure A, you know the free energy in many cases because if you have a model with only configuration entropy when you have pure A, you have no configuration entropy so the free energy there is just the energy. OK? So in some sense, if you think of this phase diagram as this line going to infinity, you essentially know the thermodynamic properties at all four edges and then you can integrate.

This is by far the most accurate way of determining first order transitions, by far. The reason is that you can get the answer as accurately as you want it. So you have to integrate now along a path, that's the painful thing. So if I like the free energy here, I don't have to just simulate here, I have to simulate pretty much all the way up from here.

So rather than simulating at one set of conditions, I got to simulate along a whole path. But the nice thing is that-- so where does your error come from? It comes from things you all control. You know, how long I've sampled at each state to get the quantity I'm integrating, like heat capacity or energies. You can control that. If you want it better, you sample longer.

How many steps I take along the integration path, so now you're numerically integrating along a part. If I want that more accurate, I take more steps. So while it looks difficult, you have all the properties under control. So that's the nice thing, that you know how to make it better if you don't like the answer.

The only thing to keep in mind when you free energy integration, that you have to iterate through equilibrium states. These thermodynamic relations do not hold when you're away from equilibrium states. So if you integrate through transitions, then you want to be absolutely sure that that transition is occurring in equilibrium, and that's often the problem, and that's why we combine all these schemes.

Like if I want to get, say, this phase boundary, I would probably integrate one phase up from here to here and then I would integrate either the solid solution from the side or from infinity to that point. OK? So I would never try to get the solid solution by integrating from the ordered phase through the transition into the solid solution because I get way too much error from non-equilibrium phenomena at the transition. OK.

Then we come to the last form of free energy integration, which is a form of thermodynamic integration that has a bit of a science fiction component to it. Come on. OK. Anyway, I think I've said all these things. What the advantages and disadvantages are, thermodynamic integration. You can actually do something a little more fancy or esoteric in thermodynamic integration, which tends to go by the name of lambda integration.

What I showed you before was you were integrating with respect to derivatives of physical parameters, like temperature or concentration or 1 over temperature. You can actually integrate with respect to nonphysical parameters. For example, I may want to get the free energy difference between systems that have two different Hamiltonians.

What would that mean? Maybe I want to know how to free energy change if I turn on a certain interaction in the Hamiltonian. Like maybe I want to know, I turn on Coulombic interactions and I want to know, how does this really affect the free energy of my system? Maybe I want to add a particle.

You could think of actually changing the temperature as a way of changing your Hamiltonian, because the thing you put in the exponential is beta times Hamiltonian, so it's really kind of one unit. It's when you change that product that you're changing something to the probability density. I'll show you some cool examples of this in a second, but let me show you how it works.

So now you're going to integrate along a path of lambda that essentially describes how you go from Hamiltonian 1 to Hamiltonian 2. And Hamiltonian 2 could be totally different physics, it could be different chemistry, whatever. So I'm going to write the Hamiltonian as a linear combination of the Hamiltonians of 1 and 2, OK?

So now you see, as lambda goes from 0 to 1, if I'm 0 I have Hamiltonian 1, if I'm one I have Hamiltonian 2. OK? So the integral 0 to 1 for lambda defines the path which I integrate. So what I want to know is essentially the free energy between lambda is 1 and 0-- but I'll get the free energy along the whole path, as you'll see in a second. You can see how you can get that.

If you look at the derivative of the free energy with respect to lambda, well, the free energy is the log of the partition function. So if I take the derivative of that, you can sort of do the math here. If I take a derivative of these exponentials I get the exponential back, so that's going to give me a partition function. Because remember, if I take derivative of the log I get 1 over this thing.

So you know, I get $\log z$. When you take the derivative of that, I'm going to get 1 over z and then times d, z, d , whatever I'm taking the derivative of with respect. So that gives me that partition function here. And then I take the derivative of what's inside the log and that gives me this.

And essentially what shows up is the derivative of the Hamiltonian with respect to λ , and this is actually classic. This always happens. If you take any derivative of the partition function, you essentially always end up with a weighted derivative of the Hamiltonian. And so what you see is this is the probability, this exponential weighted by q is the probability.

So essentially what this free energy derivative is, it's the average of the derivative of the Hamiltonian with respect to λ . OK? This is actually quite generic in statistical mechanics. So the quantity that you need to integrate is this derivative. Now if we've linearized our Hamiltonian, that derivative is just Hamiltonian difference. OK? But this is actually more generically true. But if you linearize it, all you need to do is average the Hamiltonian difference.

So I'm going to show you an example which I got out of this paper, which is hidden here by-- kind of move this thing. There we go. And this was a study where they wanted to look at the effects of a water dipole on the free energy of water. And so the reason you have a dipole in water is, of course, you have H₂O-- how does this work again?

So the hydrogens are slightly positively charged and so this is then minus 2 times that quantity. And so because of that, you of course, have a dipole which-- does a dipole point from positive to negative or, yeah, from negative to positive dipoles? Well, whatever. We'll define it this way. So you have a dipole moment and in a simple simulation, you can essentially just represent the water by its dipole, nothing else. No atoms, no molecules.

So you want to look at the dipole interactions between water. Now if you want to look at how does the free energies change with the strength of that dipole, for example, you could write dipolar strength in terms of some parameter λ , and that's what I've done here. So the positive and the negative charge in the dipole depend on the parameter λ , and so essentially what I'm going to do is look at the free energy as a function of λ . And let me show you how it works.

The green line here, there's a few too many lines on this plot unfortunately. This is the exact result-- and actually, this was well parameterized because at 0 and 1 you should get the same answer because when λ is 1, the dipole is the same as at 0, it's just inverted the plus minus charge. OK? So the purple line is what they got with λ integrations paper, so it does pretty well. I mean it's this line here, sorry.

But what you see is there's of course a bit of error. This point should be the same as that point, OK? And so you clearly see that they accumulated some error along the integration path.

And if you're really hard core computational and you think this is fun, this is one way you can check your integration errors is essentially do a circular integration in your phase space. So not only go from Hamil state 1 to 2, but come back and you essentially have some idea of the error you've accumulated along the integration path.

This one, the next one I like as an example. This is starting to get very close to alchemy. You know, in the old days people tried to turn lead into gold. This is getting pretty close. You can look at three energy changes literally by changing chemistry.

I know nothing about organic chemistry, but I think that ring there is called the phenyl and so you can put different groups on the phenyl and if you put chlorine on it, it's chlorophenyl. If you pull a methyl group on its methylphenyl, and if you put a cyan group it's cyanophenyl. So there's clearly some logic in chemical names. But for example, you could write a Hamiltonian that slowly changes one of these species into another one.

And what does that mean, slowly change? That means that you mix the interaction. Let's say you did this with potential, you would essentially write a potential from that group. Potential that comes from this group here that you attach has a weighted average of-- let's say we go from, like we did here, methylphenyl to chlorophenyl. So you would weigh the potential with the parameter λ and integrate in that space and you'd get the free energy difference between these two groups attached.

OK. So here's my take on Monte Carlo. What I really like about it, it's conceptually simple. Of all the simulation methods, it's probably the one that has the least amount of frills. It tends to be very easy to implement. And maybe the most important thing, it's as accurate as your Hamiltonian can be so you can push the sampling as far as you want it to be.

So the sampling part can be done as accurately as you have time for, essentially. Time and money for. Which is not always true about models. If I set up a potential model, I can't necessarily improve that infinitely better to model the energetics of my system. But at least with Monte Carlo, the sampling part, so the finite temperature part can be done with as little error as you'd like it to be.

I think there's sort of two or three major disadvantages is that it's not a dynamical method, so it doesn't give you any kinetic information like molecular dynamics does. Sometimes that's an advantage because it means you don't need a kinetic mechanism to study how the system goes through it's phase space. But the second one is the major hit most of the time. It's an extremely wasteful method in terms of energy evaluations.

You do a lot of kind of random excursions in phase space and every time you need to get the energy, and so that's why it's really great to implement it with fast energy methods. You know, [INAUDIBLE] small Hamiltonian G . That's just adding a few numbers and multiplying a few numbers and you got the energy, bam! Or something like potential models if you do it in a continuous space or embedded atom works great with those methods.

It's essentially impossible to implement it with quantum mechanics directly, you know? I mean, that embedded. The method, the copper nickel one I showed you did millions of energy evaluations to equilibrate the system. So can you imagine you're going to millions of direct DFT calculations? No.

So that the fact that you need typically a fast energy method is sort of a limitation. The stochastic nature of it can be a limitation. It's less and less so, but when you run two Monte Carlo simulations you don't get the same answer, so because of that, there is noise in data. It's a lot harder to write methods on top of it, in some sense, drivers for it. If you do DFT, if you converge, you get the same answer every time. So it's not stochastic.

So you can now develop algorithms that use that input and do stuff with it. It's a lot harder to work with stochastic input, let me tell you, because you kind of have to average noise away. And I think the fourth one-- you know, I should update. This is becoming less and less an issue. It used to be, but as computers get faster it's really, I would say we can get free energy when we want it, especially on models with discrete degrees of freedom.

Models with continuous degree of freedom you can do it. There's still more work, but it's not much of a disadvantage anymore. OK, I've gone through the references Before. OK. So what I want to start maybe now-- and I'm probably not going to get this finished-- is start to talk about coarse-graining methods.

And why do you need coarse-graining methods? it's essentially a nice follow up on what we just discussed. Monte Carlo allows you to get full phase space sampling, but you cannot do it on an accurate Hamiltonian like density functional theory. So what if you actually need highly accurate energetics and you need to sample phase space well? Then you're in trouble because the two are very hard to combine.

Sampling phase space, well, means a lot of energy evaluations and a lot of energy evaluations precludes using a very highly accurate Hamiltonian, but there are problems for which you need both anyway. And then you sort of need to go to coarse-graining methods. And the idea in coarse-graining methods is that you try to either remove spatial degrees of freedom or temporal degrees of freedom or your system very systematically so that your model becomes simpler and simpler, giving up as little accuracy as possible. OK? So accumulating as little error on the way as possible.

I'll talk first about temporal coarse-graining since it's easier, but if we have time in the end I may say a little about spatial coarse-graining, which is a much more difficult problem. OK, let me skip this. OK. Well, actually let me show it to you. Something wrong with the color here.

Here's an example of why you need coarse-graining. This is the copper, aluminum phase diagram. All the stable phases in here are in many cases within something like 5 to 10 milli electron volt of several other phases. So that means that to actually know that these are the stables facing that system, you need highly accurate energetics. You can't afford more than a few million electoral votes to an error. That's a very small error. So an election volt is about 100 kilojoules, so a million electron volt is about 100 joules.

All right. Yeah. So I don't know what that in calories is, but anyway, so highly accurate energetics and then to get the temperature behavior you need to sample phase space because you need to get the entropy and the excitations. For some problems like these, essentially we figured out how to do this. It's still hard work, but essentially the road map is laid out. The idea is that you successively integrate over slower and slower timescales.

So you look at what excursions, what excitations occur in the system, first at the really fast time scales and then you try to either variationally remove them or you try to integrate over them, and those two are different things. If you variationally remove a degree of freedom, then you're really finding, essentially, what gives you the lowest energy for that degree of freedom. If you integrate over the excursions of that degree of freedom, then essentially you capture its full entropic component, and I'll come back to the distinction between the two.

But let's take, for example, a simple binary solid, like that aluminum copper I showed you. What are the excitations in that system? Well, it's a metal so at the highest level there's probably electronic excitations. If the Fermi level of the system cuts through a band so you have density of states at the Fermi level, then electrons can get excited across the Fermi level. OK?

So that's a form of entropy right there already. Most of the time we don't worry about it. That's one that we variationally remove. You do the FT and you find the lowest energy state. You don't say, I'm going to integrate over all the accessible electronic states. You find the lowest one.

Then you have vibrational excitations. These are ones you can get around. They're essentially present in every material by definition. These live on timescale 10 to the minus 11 , 10 to the minus 12 , 10 to the minus 13 seconds. And then typically the slower one is configurational ones. If you have just A and B that they interchange positions, they start giving you disorder.

Again, the reason you can sort of do this problem easily is that this would be the perfect Monte Carlo problem, just like in that copper nickel example segregation. If you could do fast energy evaluations, you would just do small displacements to capture the vibrations and then you would do big exchanges to capture the configurational excitations. But because you need the accuracy, you'd almost need to do it on a DFT Hamiltonian.

You could say the vibrational ones you could capture very well with molecular dynamics. You just track the displacements of the atoms. Think about it, the sort of slower phonons go on a timescale of maybe 10 to the minus 11 . So what is that? That's 10 picoseconds. So if you simulate 100 picoseconds in nanoseconds, you're going to start fairly allegorically sampling the vibration. So you'd have a pretty good result for vibrational free energy, but you'd never get down to the configurational timescale.

So how do you solve that problem? Again, the idea is that we integrate over the fast degrees of freedom and try to define a Hamiltonian that's only defined in the phase space of the slow degrees of freedom. OK? And the question is how accurate can we do this?

So I'm going to focus on this alloy problem just so that we keep our focus. So what I want to get to is a Hamiltonian that has integrated away the electronic excitations, the vibrational excitations, and therefore that just lives in the phase space of the substitutional excitations, which is a much smaller phase space.

OK. Let me show you the math. If you think of a crystal of A and B atoms, they may live on a lattice, but of course they can be displaced from the lattice just through static relaxation but also through vibrational excursions. So normally you could define that system just by coordinate vectors.

If I have n atoms, I need n coordinate vectors. I'm going to change the way I characterize that system by first, a lattice index. So this is a topological index. This is essentially if I have a crystalline material, I could start indexing the possible sites in that material. So I will be the index of these possible lattice sites and Δr will be the displacement from these lattice sites. OK?

So do you agree that the combination of these two is essentially the same as having a full coordinate? OK. Now the set of indices i I'm going to represent essentially by a lattice model. So the set of indices i is essentially saying at each lattice point or around there-- because the atom doesn't exactly have to sit there, but around there-- is it an A or a B there?

So the variables to describe that are the same variables as a spin model or lattice model, it's a binary problem now. OK? So the question is, at I is it A or B? But again, I don't need to specify that the atom exactly sits at A or B. That I specify by the displacements, the Δr 's, OK?

So what I've done is I've separated variables that give me the configurational topology from the variables that gives me the excursions away from the ideal lattice site. And now you see where I'm going. I'm going to integrate over the excursions and then retain something that only exists in the space of the configurational degrees of freedom. OK, here we go.

So the partition function is the sum over all states. So again, the sum over all states you would have normally be the sum over all vectors. I'm going to write that as the sum over configurational states and then the sum over all displacement states, which have called ν .

So ν is the set of δr_i , and so the energy depends on both variables. OK? Now what I'm going to do is I'm going to essentially assume I can do this integration. We'll talk in a second about how you can do that. So I'm only going to do the sum over the displacement, but for a given configurational state. So what does that practically mean?

That means that I'm integrating the phase space of a fixed topology, but allowing the atoms to sort of vibrate around their average positions. That's essentially what I'm doing. So essentially I'm capturing, in that integral, the vibrational free energy component for a given configuration. And so I'm going to define.

So this is that integral we talked about. So I'm going to give that a free energy, which is just the logarithm of that. And then just if I substitute that in here, you see that what I end up with is this and this is kind of important. What do I have? I have the partition function of a lattice model.

I have a partition function that only sums over different topologies, it only sums over different configurational states, OK? So I've reduced the phase space. I'm not summing over vibrational states. Remember, I've integrated them. But what is the quantity I'm summing? This is the important thing. The quantity I'm summing is not the energy, it's the free energy of the vibrations essentially. OK?

So as you coarse-grain in time, your Hamiltonian at the slower timescale is essentially the free energy of the faster timescale. OK? Why? Because you successively integrate. So in essence, if I sum over lattice model stage and I put A's and B's on different lattice positions, what this is telling you is that the quantity-- that's my Hamiltonian-- is not the energy of that state. It's the vibrational free energy of that state.

So if I take that as my Hamiltonian and then I do this partition function, I will essentially have an exact result. I will have the exact partition function of the system. OK? And that's pretty amazing because I've really sort of separated the timescale, integrate over them separately, but I get what's essentially still-- it's an almost exact result, because let's say I'm going to do it this way.

So let's say on this I do Monte Carlo now. There is a small assumption that I've made in all of this. Think of the physical picture here. Essentially what I'm saying is that I got these A's and B's sitting on-- you can't call them exactly lattice site, but they're associated with a given lattice site.

They may be displaced from it, and they sort of vibrate around and I integrate that those vibrations to get the vibrational free energy. And then once in a while they hop, exchange, and if I sample that that gives me the free energy coming from that slower timescale. There's one assumption I've made in all of this.

It's sort of a subtle one, but I've essentially assumed that the time scales are uncoupled, and here you have to be careful with what I say. I say the timescales are uncoupled. I don't mean the energetics is uncoupled. Obviously the vibrational free energy depends on the configuration.

If I arrange the atoms differently over the lattice sites, I get a different vibrational free energy. That's not the problem. I've assumed that you can define a free energy. OK? That this thing exists. So what does that physically mean? What I assume is that for a given lattice model state, the system actually waits long enough before it goes to the next one.

If it actually only did one vibration and bam, it goes to the other one, then the system is not ergodic in its vibration. And what that means is essentially it doesn't sample all its vibrations before it goes onto the next one. So it's the fact that I can separate the excitation that really allows me to do this. Now in most materials, this is no problem whatsoever.

So vibrations, again, these are timescale 10 to minus 11 , 10 to the minus 13 kind of range. The exchanges between atoms on lattices depends on the diffusion constant, but you're going to be hard pressed to find any solid where that happens faster than of a rate of, say, 10 to the 4 per second, for example. 10 to the 4 , 10 to the 5 per second. That's really fast diffusion. That gives you diffusion constants of like 10 to the minus 7 , 10 to minus 8 , which are very high. Extremely high.

So at room temperature, for a lot of-- say for metals, this is well below one hop per second. Fast conductors you start to get to a few hops per second at room temperature. Then of course, if you go higher, temperature goes faster, but almost always are these timescales extremely well separated.

Places where they might not be like fast proton motion. They might not be very well separated and then this stuff breaks down. But remember, if they move that fast, you should just do molecular dynamics because then it's within the range. It's well within the scope of molecular dynamics, and that's the perfect approach at that point. OK.

OK, let me do one more thing and then we'll stop. There's a variety of approximations you can do. Remember that your Hamiltonian in your lattice model should be the free energy of the higher order states, which are essentially the vibrational and the electronic excitations that you've removed.

Now in some cases people say, I don't want to do all that work of integrating over the electronic states and integrating over the vibrational states. I'm going to variationally remove them, which means I'm going to find the lowest energy electronic states and the lowest energy displacement state, δr_i state. OK? So I do that practically while you take an arrangement of atoms and you just relax them both the electronic states and the positions to the minimum energy, and that gives you some E value.

So what it essentially is is if you think of the F as the free energy of an ensemble, the E you take is the lowest energy value in that ensemble. So if you do that and then you stick that in Monte Carlo, what do you have? You have proper energetics and you only have configurational entropy because rather than integrating all the vibrations and the electronics states, you've taken the minimal state out of that sub-ensemble.

You'll integrate over-- that means essentially your system hasn't sampled excursions for those variables, so you don't have entropy from those variables. OK? And then you can do all kinds of approximation. You can say, well, I don't care about the vibrations. I think that's too much work. I'm going to just get the electronic entropy.

And one of the reasons electronic entropy and metals is easy, any time you have delocalized states, you can write as a simple integral over the density of states, which you often have [INAUDIBLE] density functional theory. So if you say I'm going to take the minimal energy and the electronic entropy, well then after you've done your Monte Carlo, you have configurational entropy and electronic entropy. And then if you want to go all the way, rather than minimizing the energy you integrate over to displacement, then you're going to get the vibrational entropy component as well.

OK. I'm going to stop here because the rest is going to take us a little too long. And so I'll pick some of this up again. It's unfortunately almost two weeks away from now, I think, because like I said, so Tuesday's no lecture. Thursday-- oh, wait Thursday's you, no?

Yeah, so actually Thursday's a lecture by Professor Marzari and then it's the next Tuesday's the lab and then Thursday after that I pick this up again and then I think we're in May or something and we're almost over. Anyway, so have a good-- what is this holiday again? [INAUDIBLE] Day? No? [INAUDIBLE] Day? OK. Watch the marathon.