

Today's lecture: approximate descriptions of interatomic interactions suitable for descriptions of many organic, inorganic, and other non-reacting systems.

The classical picture

The classical approximation

The extensive computational demands of electronic structure calculations mean that their application to even modest-sized molecular systems is quite limited. Fortunately, to good approximation, we don't need to solve the Schrodinger equation for many systems to accurately reproduce their properties.

Instead, we can use a **classical** description, which ignores the motions of the electrons and describes the time-evolution of the nuclear positions alone. A classical approach uses a **force field** or **classical potential energy function** that approximates the quantum ground-state potential energy surface due to electronic structure and internuclear interactions, as a function of the positions of the nuclei.

Classical descriptions work very well under the following conditions:

- the Born-Oppenheimer approximation is valid
- the electronic structure is not of interest
- the temperature is modest (not close to absolute zero)
- there is no bond breaking or forming
- electrons are highly *localized* (metals and pi-bonded systems are delocalized)

Basic features

In the classical approximation, we describe a system by the positions and momenta of all of the atomic nuclei:

$$\mathbf{r}^N = (x_1, y_1, z_1, x_2, \dots, y_N, z_N)$$
$$\mathbf{p}^N = (p_{x,1}, p_{y,1}, p_{z,1}, p_{x,2}, \dots, p_{y,N}, p_{z,N})$$

Even though technically we deal with nuclei, we can think of the fundamental particle as an atom. Unlike quantum uncertainty, each atom has a definite position \mathbf{r} and momentum \mathbf{p} .

Alternatively, we could consider the velocity \mathbf{v} instead of the momentum, since the two simply differ by a constant mass factor. We will actually use the momentum more often, since in statistical mechanics it is naturally *conjugate* to position.

A **microstate** is just one “configuration” of the system. In a classical system, one microstate is characterized by a list of the $3N$ positions \mathbf{r}^N and $3N$ momenta \mathbf{p}^N , for a total of $6N$ pieces of information. For a microstate m we might use the notation $(\mathbf{p}_m^N, \mathbf{r}_m^N)$ to indicate specific values of these variables.

For any microstate, we can calculate the total, potential, and kinetic energies. The potential energy function depends on the positions $U(\mathbf{r}^N)$ and the kinetic energy function depends on the momenta $K(\mathbf{p}^N)$. The **Hamiltonian** of a classical system is the *function* that gives the energy of a microstate:

$$H(\mathbf{p}^N, \mathbf{r}^N) = K(\mathbf{p}^N) + U(\mathbf{r}^N)$$

The kinetic energy term simply follows

$$K(\mathbf{p}^N) = \sum_i \frac{|\mathbf{p}_i|^2}{2m_i}$$

Interactions between atoms are described by a potential energy function that depends on the positions but not the momenta of all of the atoms, $U(\mathbf{r}_1, \mathbf{r}_2, \dots)$. Fundamentally, the potential energy function approximates ground state energy that would be obtained by solving the electronic structure using Schrodinger’s equation, for fixed nuclear positions.

The time evolution of the system is deterministic and described by Newton’s equations:

$$m_i \frac{d^2 \mathbf{r}_i}{dt^2} = \mathbf{F}_i \quad \text{for all atoms } i$$

Or alternatively,

$$\frac{d\mathbf{p}_i}{dt} = -\frac{dU}{d\mathbf{r}_i}(\mathbf{r}_1, \mathbf{r}_2, \dots)$$

or finally,

$$\frac{d\mathbf{p}^N}{dt} = -\nabla U(\mathbf{r}^N)$$

Here, the LHS gives the time rate of change of momentum of atom i and the RHS gives the force on atom i which, by definition, is the derivative of the potential energy function with respect to atom i ’s position.

The total energy is constant during a system's time evolution according to Newton's laws. This is a statement of the conservation of energy. It is readily shown:

$$\begin{aligned}
 \frac{dH}{dt} &= \sum \frac{dU}{d\mathbf{r}_i} \frac{d\mathbf{r}_i}{dt} + \sum \frac{d}{dt} \left(\frac{\mathbf{p}_i^2}{2m_i} \right) \\
 &= \sum \left(\frac{dU}{d\mathbf{r}_i} \frac{d\mathbf{r}_i}{dt} + \frac{\mathbf{p}_i}{m_i} \frac{d\mathbf{p}_i}{dt} \right) && \text{combining the sums} \\
 &= \sum \left(\frac{dU}{d\mathbf{r}_i} \frac{d\mathbf{r}_i}{dt} + \frac{d\mathbf{r}_i}{dt} \frac{d\mathbf{p}_i}{dt} \right) && \text{using the fact that } \mathbf{p}_i = m_i \frac{d\mathbf{r}_i}{dt} \\
 &= \sum \frac{d\mathbf{r}_i}{dt} \left(\frac{dU}{d\mathbf{r}_i} + \frac{d\mathbf{p}_i}{dt} \right) \\
 &= \sum \frac{d\mathbf{r}_i}{dt} \times 0 = 0 && \text{using Newton's law } \frac{d\mathbf{p}_i}{dt} = -\frac{dU}{d\mathbf{r}_i}
 \end{aligned}$$

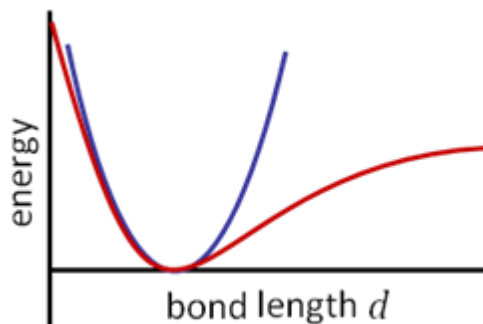
The atomic force field

The most important aspect of the classical description is the potential energy function $U(\mathbf{r}^N)$. This function takes in a set of positions of all of the atoms and returns an energy due to the interatomic interactions. This energy approximates that which one would obtain by solving the Schrodinger equation for all of the electrons in the multi-atom system, for given fixed positions of the nuclei. The approximation stems from examining the various modes by which atoms can interact according to the Schrodinger equation, and patching simple, often first-order theoretical expressions for these together.

U typically has two main components, energies due to **bonded** and **nonbonded** interactions. These arise as different interpretations of the solution of the Schrodinger equation for the electron clouds; it turns out that our notion of bonded interactions naturally manifest as solutions to the wavefunctions when atoms approach at close range.

Bonded interactions

Bond stretching



An accurate description of bond stretching that well-describes quantum-mechanical solutions is the empirical **Morse potential**:

$$u(d) = D_e [1 - e^{-a(d-d_0)}]^2$$

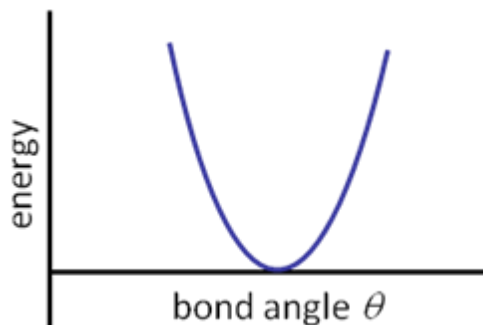
where d is the length of the bond, a a constant, d_0 the equilibrium bond length, and D_e the well depth minimum. However, this form is rarely used. It requires three parameters per bond and is somewhat expensive to compute in simulation due to the exponential term.

Since the energy scales of bond stretching are so high, bonds rarely deviate significantly from the equilibrium bond length. Thus we can use a second-order Taylor expansion around the energy minimum:

$$u(d) = a(d - d_0)^2$$

where a is a different constant from above.

Bond angle bending

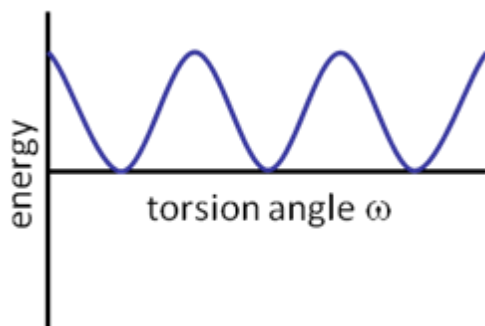


This potential accounts for deviations from the preferred hybridization geometry (e.g., sp^3). Again a common form is a second-order Taylor expansion about the energy minimum:

$$u(\theta) = b(\theta - \theta_0)^2$$

where θ is the bond angle between three atoms and b and θ_0 are constants.

Bond torsions/dihedrals



fig

These interactions occur among four atoms and account for energies of rotations along bonds. Unlike bonding and angle terms, these degrees of freedom are “soft” in that torsional energies are often not so high as to only allow small deviations from an equilibrium structure.

Generic torsional interactions can be approximated by a cosine expansion,

$$u(\omega) = \sum_{n=0}^N c_n \cos(\omega)^n$$

where ω is torsional angle, n is a summation index, c_n are summation coefficients, and N is the number of terms in the expansion. In practice, an alternative but ultimately equivalent expression is often used:

$$u(\omega) = \sum_{n=0}^N c_n [1 + \cos(n\omega - \gamma)]$$

where the c_n are different from above and the parameter γ is an additional offset parameter. This form has the advantage of converging faster for a given number of terms in the expansion. In practice, generally only one or two terms is used.

Nonbonded interactions

The **nonbonded interactions** apply to any atoms that are not closely bonded (generally with a *bond order* greater than 3 or 4), either within the same molecule or between two different molecules. These interactions are described using a **pairwise** decomposition of the energy.

Formally, we can decompose the potential energy function into interactions involving single atoms, pairs of atoms, triples of atoms, and so on and so forth:

$$U(\mathbf{r}^N) = \sum_{i=1}^N u_1(\mathbf{r}_i) + \sum_{i=1}^N \sum_{j=i+1}^N u_2(\mathbf{r}_i, \mathbf{r}_j) + \sum_{i=1}^N \sum_{j=i+1}^N \sum_{k=j+1}^N u_3(\mathbf{r}_i, \mathbf{r}_j, \mathbf{r}_k) + \dots$$

Typically, by symmetry, the two- and three-body terms do not depend on the absolute positions of the molecules but on the relative positions:

$$U(\mathbf{r}^N) = \sum_{i=1}^N u_1(\mathbf{r}_i) + \sum_{i=1}^N \sum_{j=i+1}^N u_2(r_{ij}) + \sum_{i=1}^N \sum_{j=i+1}^N \sum_{k=j+1}^N u_3(r_{ij}, r_{ik}, r_{jk}) + \dots$$


We see that multi-body effects are first included at the level of the pair potential. We would like to truncate this expansion beyond that point for the following reason: the two- and three-body interactions require a loop over atom pairs and triplets, respectively, and thus their relative

computational expense goes as N^2 and N^3 . The cubic scaling might make our simulations prohibitive for modest system sizes.

By truncating the higher order interactions, we typically neglect important effects as these make nonnegligible contributions to the total energy. Instead, we approximate the higher-order energies by renormalizing the pair interactions. At this point, we have an **effective pair potential** rather than the one which would result from this systematic expansion:

$$U(\mathbf{r}^N) \approx \sum_{i=1}^N u_1(\mathbf{r}_i) + \sum_{i=1}^N \sum_{j=i+1}^N u_{\text{eff}}(r_{ij})$$

Electrostatics



In the classical approximation, atoms can have a net charge, which may be a partial or formal charge. Here, the charges are typically determined so as to reproduce the same electrostatic potential that would be given by the true electronic structure and electron density distribution.

Atoms with partial charges interact through Coulomb's law,

$$u(r_{ij}) = \frac{q_i q_j}{4\pi\epsilon_0 r_{ij}}$$

for the two atoms i and j separated by distance r_{ij} . The partial charges are given by q_i and q_j , and ϵ_0 is the free space permittivity. Note that electrostatic interactions are long-ranged and thus require special treatment, which we will discuss in more detail in Simulations of Bulk Phases.

van der Waals attractions

Correlations between the instantaneous electron densities surrounding two atoms gives rise to an attractive energy. This is a general attractive force between all atoms, due to correlation between instantaneous dipoles between electron clouds, and can be derived from a model of interacting dipolar molecules (so called Drude molecules, Leach 4.10.1). Solving the Schrodinger equation shows that the attraction has this functional form:

$$u(r_{ij}) \propto r_{ij}^{-6}$$

where r_{ij} is the distance between the two atoms. The constant depends on the kind of the two atoms (their elements and chemical environments). These forces are called **dispersion, van der Waals, or London forces**.

Excluded volume repulsions

When two atoms make close approach, they experience a steep increase in energy and a corresponding strong repulsion. This occurs because the electron clouds of the two atoms begin to overlap, and the Pauli principle forbids any two electrons from having the same quantum numbers. At moderate internuclear distances, this potential has the approximate form:


$$u(r_{ij}) \propto \exp(-cr_{ij})$$

where c is a constant. However, this repulsion can often be successfully modeled by a simple power law that is more efficient to compute than the exponential:

$$u(r_{ij}) \propto r_{ij}^{-m}$$

where m is greater than 6.

Lennard-Jones interactions



A common way to model both van der Waals and repulsive forces is to combine them into a single expression. Lennard-Jones proposed the following simple approximation for pairwise interactions of atoms:

$$u(r_{ij}) = 4\epsilon \left[\left(\frac{r_{ij}}{\sigma} \right)^{-12} - \left(\frac{r_{ij}}{\sigma} \right)^{-6} \right]$$

where ϵ and σ are constants that depend on the particular types of atoms i and j . The minimum in the potential occurs at:

$$u(r_m = 2^{1/6}\sigma) = -\epsilon$$

The prefactor of 4 ensures that the minimum value of the potential is $-\epsilon$.

Here, the attractive force enters with the r_{ij}^{-6} term, while the repulsive energy is given by r_{ij}^{-12} . Why do we use the power of -12 for the repulsive part? It is mainly for convenience and computational efficiency, since it is just the square of the attractive term. It turns out that the properties of molecules aren't so sensitive to the exact nature of the repulsive energy, so long as it is a steep function when atoms come close.

The LJ interaction can also be written in several equivalent forms using different parameter combinations. Using the minimum energy distance r_m instead of σ we have:

$$u(r_{ij}) = \epsilon \left[\left(\frac{r_{ij}}{r_m} \right)^{-12} - 2 \left(\frac{r_{ij}}{r_m} \right)^{-6} \right]$$

Alternately, we can lump all parameters into two coefficients:

$$u(r_{ij}) = Ar_{ij}^{-12} - Cr_{ij}^{-6}$$

with

$$A = 4\epsilon\sigma^{12} = \epsilon r_m^{12}$$

$$C = 4\epsilon\sigma^6 = 2\epsilon r_m^6$$

Buckingham interactions

A sometimes-used alternative to the Lennard-Jones expression is the Buckingham potential:

$$u(r_{ij}) = \epsilon \left[\frac{6}{\alpha - 6} e^{-\alpha \left(\frac{r}{r_m} - 1 \right)} - \frac{\alpha}{\alpha - 6} \left(\frac{r}{r_m} \right)^{-6} \right]$$

This potential has three parameters: ϵ , r_m , α . Like the Lennard-Jones potential, ϵ and r_m describe the energy well depth and distance. The parameter α controls the overall shape of the well (e.g., narrowness). This potential better-models the repulsive interaction at modest nuclear distances, but is often abandoned in favor of the LJ potential because it requires three (versus two) parameters and because the exponentiation is more computationally expensive.

A minimal force field

Putting all of these energy components together, we arrive at a classical picture of molecular systems described by a potential energy function with the following minimal form:

$$U(\mathbf{r}_1, \mathbf{r}_2, \dots) = \sum_{\text{bonds } i} a_i (d_i - d_{i,0})^2 + \sum_{\text{angles } j} b_j (\theta_j - \theta_{j,0})^2$$

$$\begin{aligned}
& + \sum_{\text{torsions } k} \left[\sum_n c_{k,n} [1 + \cos(\omega_k n + \gamma_k)] \right] \\
& + \sum_{\text{pairs } ij} \frac{q_i q_j}{4\pi\epsilon_0 r_{ij}} + 4\epsilon_{ij} \left[\left(\frac{r_{ij}}{\sigma_{ij}} \right)^{-12} - \left(\frac{r_{ij}}{\sigma_{ij}} \right)^{-6} \right]
\end{aligned}$$

The potential energy is a function of all of the atomic positions \mathbf{r}^N because the bond distances d_i , angles θ_j , torsions ω_k , and pairwise distances r_{ij} are functions of \mathbf{r}^N . Of the four sums here, the most computationally expensive, by far, is that corresponding to the pairwise atomic sum, since its number of terms scales as N^2 rather than N as in the others.

Force field parameterization and transferability

The minimal force field above contains a large number of parameters:

$$a_i, d_{i,0}, b_j, \theta_{j,0}, c_{k,n}, \gamma_k, q_i, \epsilon_{ij}, \sigma_{ij}$$

Notice that there can be different sets of parameters for different types of bonds, angles, torsions, partial charges, and repulsive/dispersive interactions depending on the kinds of atoms involved and their chemical environment (i.e., an oxygen-bound carbon behaves differently than a nitrogen-bound one). This can result in a huge set of adjustable parameters that define a particular force field.

Values for force field parameters are typically taken from a combination of electronic structure calculations on small molecules and experimental data. The inclusion of experimental data tends to improve accuracy because it fits properties to “bulk” phases rather than the very small systems that ab initio methods can treat. As a result, these force fields are **semi-empirical**.

The minimal force field described above is typically fit in the following fashion:

1. **Bond stretching and angle bending parameters** are the easiest to fit because they are associated with “hard” or stiff degrees of freedom. These values can often be fitted to experimental vibrational spectra and structural data, sometimes in combination with electronic structure calculations on small molecular fragments.
2. The **Lennard-Jones parameters** can often be taken from van der Waals radii and energetics extracted from experimental crystal packing data and critical point data for small molecules. Iterative simulations of small molecule fragments are sometimes used to find the parameters that reproduce the correct bulk properties, such as phase envelopes or enthalpies of phase change.

3. Oftentimes, only the **self-self interactions** are determined for Lennard-Jones parameters, that is ϵ_{AA} and σ_{AA} for the interaction between two atoms of the same type A . Then, approximate **combining rules** can be used to estimate the parameters between non-like atom types. The Lorentz-Berthelot rules are widely used:

$$\sigma_{AB} = \frac{1}{2}(\sigma_{AA} + \sigma_{BB})$$

$$\epsilon_{AB} = \sqrt{\epsilon_{AA}\epsilon_{BB}}$$

4. **Partial charges** are assigned on the basis of ab initio calculations of the electron density surrounding small molecular fragments. To perform such calculations, the nuclear conformation is often simultaneously optimized, and if a molecule has multiple relevant conformers (e.g., rotations around a torsional angle), a separate electron density must be computed for each. The continuous electron density can then be parsed into discrete atomic partial charges by a number of methods. The **Restrained Electrostatic Potential Fit (RESP)** algorithm does this in a manner so as to reproduce the electrostatic potential around the molecule.
5. Finally, **torsional potentials** are often used to adjust the energies of rotation around bonds to accommodate any remaining energies not already captured by van der Waals and partial charge interactions that are needed for agreement with electronic structure calculations. Some force fields have no torsional potentials at all, if van der Waals and electrostatic interactions sufficiently model the net energies of bond rotations.

One can imagine that there is a fair degree of flexibility in fitting the “soft” interactions: the van der Waals, partial charge electrostatics, and torsional interactions. These terms have significant overlap. One additional goal and constraint in fitting these parameters is **transferability**, that is, the ability to reuse the same sets of parameters for different molecules and systems.

Transferability is often informed by chemistry: there might be one set of parameters for SP²-hybridized carbons, one for SP³-hybridized, one for aromatic carbons, etc. Thus, force field developers often train the value of parameters to multiple molecules and systems at once, in order to find combinations that are the most widely transferable.

A number of community-developed force field efforts exist. Within each of these efforts, there are multiple versions of the force field, as they tend to be refined over time:

force field	originated by (year)	used for / coverage
AMBER	Kollman, UCSF (1994)	proteins, nucleic acids, carbohydrates
CHARMM	Karplus, Harvard (1983)	proteins, nucleic acids, lipids, organics
GROMOS	U. Groningen (1996)	proteins, nucleic acids, sugars, organics
OPLS	Jorgensen, Purdue (1988)	many liquids
ECEPP	Scheraga, Cornell (1975)	proteins
UFF	Rappe et al. (1992)	approximate for full periodic table
MM2/3/4	Allinger, U Georgia (1977)	small molecules and hydrocarbons

The first three of these are perhaps the most widely used and developed, in part due to the widespread investigation of biomolecules using simulation. These three also refer to simulation **packages** that contain many popular algorithms for evaluating the properties of systems.

For water, there are many specialized force fields that can be used in combination with the above. The most frequently used water models are the SPC/E, TIP3P, and TIP4P models. Many biomolecular force fields (e.g. for proteins) were developed for specific use with some of these models. However, newer and more accurate water models have now been developed, including the OPC family, and TIP4P/EW, TIP4P/2005.

For systems that are non biomolecular in nature or that do not have well-established community-developed force fields available, there are two routes that might simultaneously be pursued. One is to use a **generic force field** such as the Generalized Amber Force Field (GAFF) that is designed to give reasonable representation for a variety of organic compounds. Such models should be carefully tested and benchmarked against experimental data. Second, one should always search the recent literature to identify (and evaluate) current simulation models and approaches for this the particular system of interest, if any such work exists.

Polarizable force fields

The **fixed-partial charge model** described above isn't able to account for induced **polarization**, that is, the rearrangement of a charge distribution in a molecule due to nearby interacting molecules. A number of methods have emerged for incorporating this effect into classical force fields. Two common approaches:

- fluctuating atomic partial charges
- induced dipoles on each atom

In each case, the partial charge or induced dipole must be solved iteratively *for each conformation of the molecule*. The iteration typically reaches a self-consistent state where the electric field due to the charges/dipoles is consistent with the induced response in the molecule.

Because of the need to iterate upon each step of a simulation, polarizable force fields are expensive to simulate and historically have not been widely used. However, with increased computing power and the recognition that polarizability can significantly improve the accuracy of a force field, these approaches are currently being explored by many groups.

Special force fields

A number of other classical force field forms have been developed to better capture specific behavior in atomic systems. These methods all aim to improve agreement with the true quantum-mechanical behavior and hence increase quantitative accuracy. Though we won't discuss them here, we briefly mention some:

- hydrogen bonding potentials (highly directional interactions)
- water models (see http://www1.lsbu.ac.uk/water/water_models.html for an impressive overview of over 30 current models and their capabilities)
- systems involving metals and semiconductors (delocalized electrons, pi-bonding, and directional bonding interactions)
- ionic solids (coupling between polarizability and repulsions)

Force field unit systems


Generally we need to choose an appropriate unit system when setting up a simulation in order to capture length, energy, and time scales consistent with reality. Exceptions are highly simple systems with only a very few force field parameters, for which we can perform dimensionless simulations (see last section below). For all other systems, common unit systems are represented in the following table. Note that once base units for length, energy, and mass have been specified, the intrinsic base unit for time necessarily follows. Temperature is nearly always modeled in Kelvin.

unit system	energy unit ϵ	length unit σ	mass unit m	time unit $\tau = \sqrt{m\sigma^2/\epsilon}$
atomic like AMBER	kcal/mol	Angstrom	atomic mass unit (1.66×10^{-27} kg)	48.9 fs (4.89×10^{-14} s)
atomic like GROMACS	kJ/mol	nm	atomic mass unit (1.66×10^{-27} kg)	ps (10^{-12} s)
MKS	J/mol	m	kg	s

Interactions with solid surfaces

To reproduce the effects of a molecules interacting with solid surface or interface, one option can be to explicitly include several layers of surface atoms that have interactions with each other and with other molecules in the system. However, if the effects of the structure of the surface are not important or the subject of investigation, one can approximate these interactions without explicitly including surface atoms. This approach has the benefit of being far less computationally expensive since it replaces the interactions of a molecule with all of the atoms in a surface by a single interaction. This is one form of **coarse-graining**, which we will discuss in more detail later in the course.

Consider an atom interacting with a surface built from many atoms. Assume the interaction between the atom and any surface atom is described by the Lennard-Jones potential.



The total interaction of the atom with the surface is the sum of the interactions with all atoms out to infinity in the $\pm x, \pm y, -z$ directions (assuming the surface is infinitely thick). We assume the atoms in the surface are **smear**ed out, that is, uniformly distributed with number density ρ_S . Therefore we can compute the total interaction as:

$$\begin{aligned}
u(z) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{-z} \rho_S dZ dx dy \times 4\epsilon \left[\left(\frac{r}{\sigma}\right)^{-12} - \left(\frac{r}{\sigma}\right)^{-6} \right] \\
&= \int_0^{2\pi} \int_0^{\infty} \int_{-\infty}^{-z} \rho_S dZR dR d\theta \times 4\epsilon \left[\left(\frac{r}{\sigma}\right)^{-12} - \left(\frac{r}{\sigma}\right)^{-6} \right] \\
&= \int_0^{2\pi} \int_0^{\infty} \int_{-\infty}^{-z} \rho_S dZR dR d\theta \times 4\epsilon \left[\frac{\sigma^{12}}{(Z^2 + R^2)^6} - \frac{\sigma^6}{(Z^2 + R^2)^3} \right] \\
&= \frac{4\pi\epsilon\rho_S\sigma^3}{3} \left[\frac{1}{15} \left(\frac{z}{\sigma}\right)^{-9} - \frac{1}{2} \left(\frac{z}{\sigma}\right)^{-3} \right]
\end{aligned}$$

This **9-3 potential** describes the interaction of a single atom with a virtual wall. This is a single-body term, not a pairwise or higher order interaction, and thus its computational expense scales as N . Here, the z coordinate is given by the z -component of the atomic position of the atom, which corresponds to the wall being located at $z = 0$. If the wall is at some other location, say z_0 , then the potential becomes

$$u(z) = \frac{4\pi\epsilon\rho_S\sigma^3}{3} \left[\frac{1}{15} \left(\frac{z - z_0}{\sigma}\right)^{-9} - \frac{1}{2} \left(\frac{z - z_0}{\sigma}\right)^{-3} \right]$$

Note that the potential involves the wall density. Typical crystalline close-packing densities can be used for this parameter, $\rho_S\sigma_{\text{surface}}^3 = 0.74$, where σ_{surface} gives the size of atoms in the surface and can be distinct from the atom-surface interaction parameter σ . Another way of writing the potential above involves the substitutions,

$$\sigma_0 = \left(\frac{2}{15}\right)^{1/6} \sigma \quad \epsilon_0 = \frac{8\pi}{3^{5/2}} \epsilon\rho_S\sigma^3$$

such that,

$$u(z) = \frac{3\sqrt{3}}{2} \epsilon_0 \left[\left(\frac{z}{\sigma_0}\right)^{-9} - \left(\frac{z}{\sigma_0}\right)^{-3} \right]$$

For this formulation of the potential, the minimum energy is $-\epsilon_0$ and it occurs at distance $z = 3^{1/6}\sigma_0$.

Implicit solvation

In systems of solvated molecules, such as proteins or DNA, the requirement of a very large number of solvent molecules can drastically increase the simulation expense (which goes as the number of atoms squared). When the details of the solvent are not of interest, it is possible to create a continuum approximation that correctly reproduces the thermodynamic properties of a system. This is another form of **coarse-graining**.

Statistical-mechanical formulation

Consider N solute atoms (perhaps part of a macromolecule) embedded in a solvent consisting of M solvent atoms. Assuming pairwise interactions, we can rigorously break the total potential energy into solute-solute (X-X), solute-solvent (X-S) and solvent-solvent (S-S) interactions:

$$U(\mathbf{r}_X^N, \mathbf{r}_S^M) = U_{XX}(\mathbf{r}_X^N) + U_{XS}(\mathbf{r}_X^N, \mathbf{r}_S^M) + U_{SS}(\mathbf{r}_S^M)$$

The total configurational part of the canonical partition function can be written as:

$$Z = \int \int e^{-\beta U_{XX}(\mathbf{r}_X^N) - \beta U_{XS}(\mathbf{r}_X^N, \mathbf{r}_S^M) - \beta U_{SS}(\mathbf{r}_S^M)} d\mathbf{r}_S^M d\mathbf{r}_X^N$$

This can be factored as:

$$Z = \int e^{-\beta U_{XX}(\mathbf{r}_X^N)} \left[\int e^{-\beta U_{XS}(\mathbf{r}_X^N, \mathbf{r}_S^M) - \beta U_{SS}(\mathbf{r}_S^M)} d\mathbf{r}_S^M \right] d\mathbf{r}_X^N$$

Define

$$F_S(\mathbf{r}_X^N; \beta) = -k_B T \ln \int e^{-\beta U_{XS}(\mathbf{r}_X^N, \mathbf{r}_S^M) - \beta U_{SS}(\mathbf{r}_S^M)} d\mathbf{r}_S^M$$

This is the free energy of solvation for a particular solute conformation \mathbf{r}_X^N . Notice that this free energy depends on the temperature. We can now rigorously rewrite the total partition function as:

$$Z = \int e^{-\beta U_{XX}(\mathbf{r}_X^N) - \beta F_S(\mathbf{r}_X^N; \beta)} d\mathbf{r}_X^N$$

This becomes an integral over the solute conformations only. Thus, we can define an effective potential energy function for the system, in which we have averaged out the solvent degrees of freedom:

$$U_{\text{eff}}(\mathbf{r}^N) = U(\mathbf{r}^N) + F_S(\mathbf{r}^N; \beta)$$

Here, for simplicity we removed the X subscript from the solute degrees of freedom.

The challenge is to determine a good functional form for $F_S(\mathbf{r}^N; \beta)$. If we find an accurate one, by the derivation above, our potential will rigorously return the same thermodynamic properties as a fully atomic treatment of the solvent. Any kinetic properties associated with U_{eff} , however, will be different since now we have averaged away any solvent degrees of freedom that would contribute to a viscosity.

Simple approaches to continuum solvents

The simplest approach is to assume the solvent is a continuous dielectric medium with a constant dielectric. In this case, all solute partial charge interactions will be reduced relative to the vacuum Coulomb case:

$$u(r_{ij}) = \frac{1}{4\pi\epsilon_0\epsilon} \frac{q_i q_j}{r_{ij}}$$

Here, ϵ is the dielectric constant for the medium of interest. For water, $\epsilon = 80$. For vacuum, $\epsilon = 1$.

A slightly better approximation that interpolates between a dielectric of $\epsilon = 1$ at short distances and $\epsilon = \epsilon_{\text{solvent}}$ at large distances is a distance-dependent dielectric:

$$u(r_{ij}) = \frac{1}{4\pi\epsilon_0\epsilon_{\text{eff}}(r_{ij})} \frac{q_i q_j}{r_{ij}}$$
$$\epsilon_{\text{eff}}(r) = \epsilon_{\text{solvent}} - \frac{\epsilon_{\text{solvent}} - 1}{2} (r^2 S^2 + 2rS + 2)e^{-rS}$$

Here, S is a parameter, typically between 0.15 \AA^{-1} and 0.3 \AA^{-1} .

Another approximate form suitable for use in dilute **electrolyte solutions** (systems with mobile charges) is the **screened Coulomb potential**:

$$u(r_{ij}) = \frac{q_i q_j}{4\pi\epsilon_0\epsilon r_{ij}} e^{-\frac{r_{ij}}{\lambda_D}}$$

Here, λ_D is the Debye length, dependent on the salt concentration of the solution, and ϵ is the dielectric constant of the medium.

None of these methods, however, is particularly accurate at reproducing the net effect of the potential U_{eff} in the general case. What they neglect is the complex, multi-body rearrangement of the solvent in response to the solute. This effect is typically dominated by electrostatic interactions between the two.

More accurate approaches

A common approach to modeling F_S is to split it into several parts for which good analytical approximations can be made:

$$F_S = F_{\text{elec}} + F_{\text{vdW}} + F_{\text{cav}}$$

This is an approximate separation into interactions between the solute and solvent involving: electrostatic, attractive van der Waals, and repulsive or cavity interactions. Typically for many macromolecules, the cavity term can be neglected and the van der Waals interactions be successfully modeled using a surface area approach:

$$F_{\text{nonelec}}(\mathbf{r}^N) = F_{\text{vdW}} + F_{\text{cav}} \\ \approx \gamma A_{\text{SAS}}(\mathbf{r}^N)$$

Here, γ is the surface tension of the molecular boundary with the solvent. It is typically treated as a universal constant, although some approaches choose a different value of surface tension for each atom type at the boundary.

A_{SAS} is the **solvent-accessible surface area**. It is the area one would compute by tracing out the surface formed by rolling a solvent molecule over the solute(s). The surface area depends on the conformation of the solute(s) and there exist fast, approximate methods for computing it.

The electrostatic component is typically treated by considering the solvent to be a dielectric described by continuum electrostatics. The fundamental equation that must be solved for the solvent dielectric is the **Poisson-Boltzmann equation**:

$$\nabla \cdot \epsilon(\mathbf{r}) \nabla \phi(\mathbf{r}) - \kappa' \sinh[\phi(\mathbf{r})] = -4\pi\rho(\mathbf{r})$$

Here, ϵ is the dielectric constant, which varies as a function of position in space. ϕ is the electrostatic potential and also varies in space. κ' is a constant that depends on the ionic strength of the solution. ρ gives the charge density, which can be extracted from the atomic partial charges on the solute.

How does this equation work? In broad terms,

1. The dielectric $\epsilon(\mathbf{r})$ is assumed to be equal to ϵ_{solute} everywhere inside the solute (often 1) and equal to $\epsilon_{\text{solvent}}$ everywhere outside the solute (e.g., beyond the solvent accessible surface area). In practice, the transition is often made continuous over the solute boundary for numerical stability.
2. The charge density $\rho(\mathbf{r})$ is constructed from the atomic charges in the solute. In practice, these charges are often smeared over a grid in space.
3. One then calculates the electrostatic potential $\phi(\mathbf{r})$ everywhere in space by solving this nonlinear, second-order differential equation. This is the hard and expensive part.
4. With the solution for $\phi(\mathbf{r})$, one can compute the electrostatic component of the solvation free energy from a sum over partial charges in the solute: $F_{\text{elec}} = \frac{1}{2} \sum_i q_i \phi(\mathbf{r}_i)$

Solutions to the Poisson-Boltzmann approximation often give good estimates of the electrostatic component of the solvation free energy. The quality of these results relies on the extent to which the solvent behaves like a dielectric continuum. In cases where molecular granularity is important—for example, in which a single water molecule bridges two ion’s interaction—this approach fails.

The Poisson-Boltzmann equation is fairly computationally demanding to solve. There are several approximate solutions that work well; however, a complete solution needs to be found every time the conformation of the solute changes.

There has been much interest in using fast, highly approximate solutions to the PB equation that don’t require solution of a differential equation. One that has gained much popularity is the so-called **Generalized Born** method. In this approach,

$$F_{\text{elec}} \approx -\frac{1}{4\pi\epsilon_0} \left(\frac{1}{\epsilon_{\text{solute}}} - \frac{1}{\epsilon_{\text{solvent}}} \right) \sum_{i=1}^N \sum_{j=i+1}^N \frac{q_i q_j}{\sqrt{r_{ij}^2 + a_i a_j e^{-r_{ij}^2/4a_i a_j}}}$$

Here, the summations proceed over all pairwise interactions of partial charges in the solute. This equation bears some resemblance to the Coulomb expression, most notably except for the square root term. The variables a_i are called the Born radii of each atom in the solute. Generally speaking, they measure the radius from an atom to the solute-solvent boundary, and they must be computed as a function of solute conformation.

Generalized Born models are most frequently used in the simulation of large solvated biomolecules, although they have seen a decline in popularity in the past decade given hardware advances that allow for explicit water modeling. While the computation of F_{elec} is typically the most expensive calculation in the evaluation of the potential energy of these systems, the GB approach gives significant time savings relative to including large numbers of explicit water molecules.

Keep in mind that the PB and GB methods must be coupled with an approach to determine the nonelectrostatic component of the solvation free energy. The complete approaches that we have discussed here are termed the **Poisson-Boltzmann Solvent Accessible Surface Area Method (PBSA)**:

$$U_{\text{eff}}(\mathbf{r}^N) = U(\mathbf{r}^N) + \gamma A_{\text{SAS}}(\mathbf{r}^N) + F_{\text{elec,PB}}(\mathbf{r}^N)$$

and the **Generalized Born Solvent Accessible Surface Area Method (GBSA)**:

$$U_{\text{eff}}(\mathbf{r}^N) = U(\mathbf{r}^N) + \gamma A_{\text{SAS}}(\mathbf{r}^N) + F_{\text{elec,GB}}(\mathbf{r}^N)$$

In both cases, $U(\mathbf{r}^N)$ simply corresponds to the kind of semi-empirical force field we described earlier for the intramolecular interactions of atoms in the solute with each other.

Simpler perspectives

Simplified and coarse-grained models

Oftentimes, we are not interested in *quantitative* accuracy and do not need a highly detailed force field. We can examine a large range of behavior using very simplified models. This includes:

- scaling laws (e.g., dependence of properties on system size, chain length, molecular size or energy scales)
- microscopic mechanisms (e.g., of diffusion, binding, or conformational changes)
- relative magnitudes of different driving forces (e.g., electrostatic, dispersive, excluded volume/repulsive, hydrogen bonding, hydrophobic interactions)
- microscopic structure (e.g., conformational fluctuations, molecular packing, structure in bulk liquids, degree of geometric ordering)
- functional form of the dependence of these and other properties on state conditions (e.g., temperature, pressure, density, composition)

The simplified models themselves can be used to rationalize many universal behaviors in molecular systems. For example, simple bead-spring models are able to capture a huge array of polymer thermodynamics. Similarly, very simple Lennard-Jones type potentials can describe many features of critical phenomena for a wide range of single- and multi-component liquid-state systems.

There is an important philosophy here that goes well beyond practical simulation needs. To develop a basic *physical* understanding of a particular system, we would like to identify a minimal model that best describes the basic qualitative features of its behavior.

In that respect, we ultimately want to pick the simplest model possible, use simulations to evaluate its behavior, and determine by comparison with experiment which properties are correctly reproduced and which are not. When then begin systematically building greater detail into the model, evaluating at each point which new properties improve agreement. By taking such a hierarchical approach, we can pinpoint exactly which aspects of the physics explain various features of a system's behavior, in terms of driving forces and interactions.

What are simple models? These are often called **coarse-grained** or reduced models. They involve large **pseudoatom** sites that are designed to represent combined groups of multiple atoms. For example, a bead-spring polymer model might involve one pseudoatom per monomer.

Basic considerations to make when constructing such models is to ask:

- What degrees of freedom need to be included? In other words, what is the basic chemical architecture of the molecules that drives the relevant physics? For example, is it desired to capture motions along a polymer backbone? If so, one needs several large pseudoatom sites along it. If the motions of side-chain groups hanging off of the backbone affect the physics, one will have to include pseudoatoms for those as well.
- What are the dominant length scales in the system? Excluded volume interactions will have to be modeled with repulsive terms in the pairwise energy functions for these.
- What are the dominant energy scales in the system? Will a Lennard-Jones functionality (repulsion + attraction) suffice to model these? Are other functionalities more appropriate? A hard-sphere term or a square-well model might also be appropriate here, although discrete potentials can require special molecular dynamics methods if kinetics are to be investigated.
- How many different kinds of interaction potentials and pseudoatom types are needed? Ultimately, one wants to pick a minimal number of types that will allow an understanding of the basic physics.

Dimensional analysis

With highly simple models that involve only a very few length and energy scales, it is sometimes possible to perform simulations in reduced, dimensionless parameters. This approach is *should NOT be used* for simulations of detailed atomic-scale systems in which multiple atom types and thus many parameters are present.

Consider an atomic liquid modeled by the Lennard-Jones interaction. Such a model might represent argon or other noble gasses at a quantitative level, and many other molecular liquids at a qualitative level. Here, there are N atomic molecules. The total potential energy is:

$$U(\mathbf{r}^N) = \sum_{i < j} 4\epsilon \left[\left(\frac{r_{ij}}{\sigma} \right)^{-12} - \left(\frac{r_{ij}}{\sigma} \right)^{-6} \right]$$

Imagine that we simulate this system in a cubic box of length L (volume V) at a given temperature T . If we let the fundamental length unit in our system be σ and the fundamental energy unit be ϵ , we can rewrite the potential in dimensionless units as:

$$\begin{aligned}
 U^*(\mathbf{s}^N) &\equiv \frac{U}{\epsilon} \\
 &= \sum_{i < j} 4[s_{ij}^{-12} - s_{ij}^{-6}]
 \end{aligned}$$

where \mathbf{s}^N now gives the dimensionless reduced positions:

$$\mathbf{s}^N = \frac{\mathbf{r}^N}{\sigma}$$

Notice that the reduced potential no longer contains the parameters ϵ and σ . This enables us to simulate a *generic* Lennard-Jones system in reduced parameter space. That is, we perform a simulation where our positions are the reduced \mathbf{s}^N and our energies are U^* .

We can define reduced simulation conditions. The reduced temperature:

$$T^* = \frac{k_B T}{\epsilon}$$

The reduced volume:

$$V^* = \frac{V}{\sigma^3}$$

The reduced number density (number of molecules per volume, N/V):

$$\rho^* = \rho \sigma^3$$

The reduced pressure:

$$P^* = P \frac{\sigma^3}{\epsilon}$$

To non-dimensionalize any quantities involving time, such as a diffusivity or viscosity, we need to incorporate the mass m of each of the particles. The reduced time is:

$$t^* = t \sqrt{\frac{\epsilon}{m \sigma^2}}$$

The reduced velocities and momenta are:

$$p^* = p \sqrt{\frac{1}{m \epsilon}} \quad v^* = v \sqrt{\frac{m}{\epsilon}}$$

The reduced kinetic energy is:

$$\begin{aligned}
K^* &= \frac{K}{\epsilon} \\
&= \frac{1}{2} \sum_i |\mathbf{p}_i^*|^2 = \frac{1}{2} \sum_i |\mathbf{v}_i^*|^2
\end{aligned}$$

The reduced self-diffusivity is:

$$D^* = D \sqrt{\frac{m}{\epsilon \sigma^2}}$$

The advantage of the dimensionless approach is the following. We can map out the entire behavior of this system in dimensionless state space (e.g., as a function of T^* and ρ^*) without specifying values of the parameters ϵ and σ . Our simulation results can then be scaled to specific values of ϵ and σ for a particular system or set of systems using the relations above.

If our potential contained three parameters, we would have to evaluate our dimensionless system as a function of three different state conditions. Beyond three or four parameters, however, this approach grows unwieldy and is not particularly informative.

Another example

Generally, there are **only three scales that can be chosen** in the system: a **length scale** σ_0 , an **energy scale** ϵ_0 , and a **mass scale** m_0 . The mass scale does not affect any of the thermodynamic properties, but rather defines the time scale via the relation:

$$t_0 = \sqrt{\frac{m_0 \sigma_0^2}{\epsilon_0}}$$

Once the energy and length scales are chosen, there is a unique form of the dimensionless interaction potential.

It is important to choose scales that are representative of the length and energy scales that dominate the physics of interest in the system. Consider a system of particles that interacts through a **screened Coulomb potential plus a soft-sphere repulsion**:

$$u(r_{ij}) = \epsilon_{SS} \left(\frac{r_{ij}}{\sigma_{SS}} \right)^{-12} + \frac{1}{4\pi\epsilon_0\epsilon} \frac{q_i q_j}{r_{ij}} e^{-\frac{r_{ij}}{\lambda_D}}$$

If the charges are specified in terms of the unit electron charge e_0 , then this expression becomes

$$u(r_{ij}) = \epsilon_{SS} \left(\frac{r_{ij}}{\sigma_{SS}} \right)^{-12} + \frac{e_0^2}{4\pi\epsilon_0\epsilon} \frac{q_i q_j}{r_{ij}} e^{-\frac{r_{ij}}{\lambda_D}}$$

To non-dimensionalize we first choose a length scale. There are two in this system, the soft-sphere diameter σ_{SS} and the Debye length λ_D . The dimensionless density and volume fraction of our system is more naturally expressed in terms of the particle diameters, so we typically choose the particle size parameter as the length scale:

$$r_{ij}^* = \frac{r_{ij}}{\sigma_{SS}} \quad ; \quad V^* = \frac{V}{\sigma_{SS}^3}$$

The potential then becomes:

$$u(r_{ij}) = \epsilon_{SS}(r_{ij}^*)^{-12} + \frac{e_0^2}{4\pi\epsilon_0\epsilon\sigma_{SS}} \frac{q_i q_j}{r_{ij}^*} e^{-r_{ij}^* \left(\frac{\sigma_{SS}}{\lambda_D}\right)}$$

We can now define two new parameters,

$$\epsilon_{Coul} = \frac{e_0^2}{4\pi\epsilon_0\epsilon\sigma_{SS}}$$

$$\lambda_D^* = \frac{\lambda_D}{\sigma_{SS}}$$

The latter is the Debye length expressed in units of particle diameter. Thus,

$$u(r_{ij}) = \epsilon_{SS}(r_{ij}^*)^{-12} + \epsilon_{Coul} \frac{q_i q_j}{r_{ij}^*} e^{-r_{ij}^* \left(\frac{\sigma_{SS}}{\lambda_D}\right)}$$

Our second choice is the energy scale in the potential. We have two from which we can choose: ϵ_{SS} and ϵ_{Coul} . Typically, we want to pick the one that will dominate the physical interactions of our system. Here, the soft-sphere term is relatively insensitive to ϵ_{SS} because the exponent is so great. Consider the limit that the repulsive exponent becomes negative infinity; in this case, this term has the behavior:

$$\lim_{n \rightarrow \infty} \epsilon_{SS}(r_{ij}^*)^{-n} = \begin{cases} \infty & r_{ij}^* < 1 \\ 0 & r_{ij}^* > 1 \end{cases}$$

In other words, we recover the hard-sphere interaction in this limit and ϵ_{SS} becomes completely irrelevant to the system behavior.

As a result, we choose ϵ_{Coul} as the energy scale in the system. This means that

$$u^*(r_{ij}^*) = \frac{u(r_{ij}/\sigma_{SS})}{\epsilon_{Coul}}$$

$$T^* = \frac{k_B T}{\epsilon_{Coul}}$$

The dimensionless interaction potential is thus:

$$u^*(r_{ij}^*) = \frac{\epsilon_{SS}}{\epsilon_{Coul}} (r_{ij}^*)^{-12} + \frac{q_i q_j}{r_{ij}^*} e^{-r_{ij}^*/\lambda_D^*}$$

We now have two dimensionless parameters that we can tune in our system in addition to the reduced temperature T^* and density ρ^* :

$$\lambda_D^* ; \left(\frac{\epsilon_{SS}}{\epsilon_{Coul}} \right)$$

Since we might expect our system behavior to be relatively insensitive to ϵ_{SS} for large negative values of the repulsive exponent, we might set

$$\frac{\epsilon_{SS}}{\epsilon_{Coul}} \approx 1$$

And a simpler potential would be

$$u^*(r_{ij}^*) = (r_{ij}^*)^{-12} + \frac{q_i q_j}{r_{ij}^*} e^{-r_{ij}^*/\lambda_D^*}$$

Simulations in the canonical ensemble could now be performed for different values of T^* , ρ^* , and λ_D^* to map out the complete behavior of this simple system.