

Equilibrium and Statistical Mechanics

Conor Rowan

Winter 2024

1 Introduction

The equilibrium configuration of a linearly elastic body is found by computing the displacement field which minimizes the “total potential energy” Π . This energy functional is defined as

$$\Pi(u(x)) = U - W = \int_V \Psi dV - \int_v b_i u_i dV - \int_S t_i u_i dS$$

The functional U is the internal strain energy of the system, and W is the work done by external body forces and tractions. The total strain energy of the system is computed by integrating the strain energy density Ψ over the volume of the body. In linearly elasticity, we know that

$$\Psi = \frac{1}{2} \sigma_{ij} \epsilon_{ij} = \frac{1}{2} C_{ijkl} \epsilon_{ij} \epsilon_{kl}$$

This also demonstrates that the stress and strain tensors are related through the strain energy density with

$$\sigma_{ij} = \frac{\partial \Psi}{\partial \epsilon_{ij}}$$

It is often said that the stress is “conjugate” to the strain when they are related through gradients of the strain energy. There are a number of concepts at play here which are relatively simple to understand in the case of linear (small strain) problems, but become less clear in non-linear elasticity. The first set of questions relate to the energy formulation: why does the equilibrium configuration of a body correspond to minimizing an energy functional? Why is this energy functional the difference between internal energy and external work? The second set of questions relates to the conjugate relationship of stress and strain: what does it mean for stress and strain measures to be conjugate to one another? How does the same strain energy density spit out different stress measures when differentiated with respect to various strains? Where does this strain energy density come from?

One answer to concerns around the energy formulation of linear elasticity is that one can easily verify that the Euler-Lagrange equations corresponding to the energy functional lead to the strong form of the governing equations of elasticity, which are derived based on force balances alone. Thus, equilibrium and energy minimization simply say that the displacement field satisfies Newton’s second law at every point in the body. This suggests that the energy is appropriately defined, but does not justify it is a fundamental principle. Regarding the conjugate relationship between stress and strain, it is simple to see by differentiation that the above relationship holds, but once again, verifying that it is correct does not explain what it means.

I believe that answering these questions involves deeply understanding thermodynamics. In non-linear elasticity, thermodynamic principles are more freely evoked, though derivations of finite strain models do not tend to carefully explain their foundation in thermodynamics. One example of this is the constant invocation of the “free energy” as the appropriate analogue of the strain energy density in linear elasticity. This is a fundamentally thermodynamic quantity, especially because it involves temperature and entropy (as we will see). In my experience, the meaning and role of the free energy in finite strain elasticity is totally perplexing. In this short report, I will attempt to explore how concepts like equilibrium, energy minimization, and free energy are fundamentally tied to thermodynamics. Namely, I will demonstrate from a statistical thermodynamic perspective how the equilibrium configuration of a box of gas is obtained by minimizing the Helmholtz free energy. Just as principles of elasticity appear impossible to fully grasp without appeal to thermodynamics, it seems that classical thermodynamics is impossible to understand without appeal to the underlying atomic and statistical nature of matter. This is the domain of statistical thermodynamics. I have not yet fully charted the connection between statistical thermodynamics, classical thermodynamics, and continuum mechanics. For example, I am not sure how to derive the total potential energy for an elastic system starting from a statistical thermodynamic perspective. But, in showing that the equilibrium state of a box of gas corresponds to the minimum of a free energy function, I feel that I have gotten a little bit closer.

2 Micro- and Macro-states

Imagine a collection of N fair coins. When we flip every single one of them, each coin is in either a state of heads or tails with equal probability. By definition, the “microstate” of the system is the state of all the coins. This could be stored as a vector

$$\underline{s} = [h, t, t, \dots]^T$$

where h and t represent “heads” and “tails” respectively. Note that we can distinguish between coins, meaning that each coin is thought of as being labeled so

that we know which slot in the microstate vector it occupies. This demonstrates that the order of the entries in the microstate vector matter:

$$[h, t, \dots] \neq [t, h, \dots]$$

Conversely, the “macrostate” of the system is the total numbers of heads and/or tails contained in the microstate vector. Of course, the order of heads and tails in the microstate vector does not matter when we compute a total number of heads or tails. The macrostate is a “coarse-grained” representation of the microstate. If we choose the macrostate to be the number of tails in the microstate, the coarse-graining operation replaces an N -dimensional vector with a scalar.

The coin flip system is just an example—we are ultimately interested in physical systems. For a gas, the microstate is the position and momentum of every particle making up the gas. In three dimensions, there are six coordinates corresponding to every gas particle, thus the microstate vector has dimension $6N$ where N is the number of particles. This requires justification, but it makes intuitive sense that the macroscopic state of a collection of gas particles might involve variables such as temperature and pressure. Though we know that thermodynamic phenomena are the result of the evolution of microstates of material systems (after all, apparently fundamental thermodynamic concepts like heat and temperature are really emergent properties of these discrete systems), we would like to neglect the tremendously high-dimensional microstate of the thermodynamic system to whatever extent is possible. The coarse-graining operations which relate position/momentum microstates to temperature/pressure/entropy macrostates are an attempt to ignore this complex microscopic behavior with minimal loss of information. Statistical mechanics should justify how macroscopic variables like these accurately characterize the relevant properties of the system.

It is not necessarily clear that coarse-grained variables should be able to precisely characterize the states of large statistical systems. Of course, we know that the coarse-grained approach will not allow us to make claims about what individual components of the system are doing. We are interested in summary quantities of the microscopic system, such as the total number of tails (coin flip), or the total kinetic energy of the system of particles (gas). We envision both of these systems as statistical, even though it may seem odd to treat a classical physical system like a collection of particles as random. Of course, these particles evolve according to the laws of mechanics. But, because we do not know the state of all of the particles in a gas (and we don’t want to know it!), we model microstates as occurring randomly. All microstates are equally probable, so there is some high-dimensional uniform distribution over possible microstates of the gas system. Note that admissible microstates can be constrained to have a specific total energy (this is what we observe, after all). This is not important

right now—the point is that we think of microstates as occurring randomly, even in a gas.

Because the underlying microscopic systems are statistical, repeating an experiment will yield different measurements of the microstate. We need to demonstrate that a coarse-grained variable is an accurate characterization of the system even when the microstates vary randomly from experiment to experiment. In the coin flip example, does the average number of heads accurately represent the macrostates we might encounter by repeating the flipping experiment multiple times? Say that we have N coins. There are 2^N microstates available to the system when we can distinguish between coins. To explore whether the average is a useful method of coarse-graining the system, we need to more precisely describe the statistics of the system. The first question we ask is: how many different microstates are there corresponding to the macrostate of m observations of heads? Call this $\Omega(m)$. We can explore this by looking at a few examples. If $N=3$, we have

$$\Omega(0) = 1, \quad \Omega(1) = 3, \quad \Omega(2) = (3 * 2)/2, \quad \Omega(3) = 1$$

Using a general form of this counting variable Ω for a system of N coins, a probability distribution is obtained by normalizing by the total number of states:

$$p(m) = \frac{1}{2^N} \Omega(m; N) = \frac{1}{2^N} \frac{N!}{m!(N-m)!}$$

It can be shown that the mean of this distribution is $N/2$, and its standard deviation is $\sqrt{N}/2$. Thus, the ratio of the standard deviation to the mean is

$$\frac{\sigma}{\mu} = \frac{1}{\sqrt{N}}$$

This means that relative fluctuations from the mean become arbitrarily small as the number of “particles” in the system becomes large. It can be shown with the central limit theorem that this is a very general result. In other words, when we sum up many probability distributions, the resulting distribution will become normal, and the standard deviation becomes very small compared to the mean as the size of the system grows. This means that when measuring the state of a system of independently distributed particles whose sum constitute the measurement, it is very unlikely to take a measurement which differs significantly from the mean. This justifies the use of coarse-grained macroscopic variables like an average for large systems of independently distributed particles. Note that an important assumption is that each particle contributes to the coarse-grained variable additively, and that there are no correlations between microstates. This limits the kinds of systems which the central limit theorem applies to. If the gas particles interacted strongly with one another, their states would be correlated, thus the probabilities of finding them in certain states would not be independent. Similarly, if we were interested in a macroscopic

variable other than energy, it may not be clear that each particle contributes additively to this macrostate.

3 Energy and Counting Microstates

Unlike the coin flip example, it is not possible to directly count the number of microstates a gas can take on because the position and momentum vary continuously. However, what we can do is say that the number of microstates available to a gas is proportional to the volume in a state space defining position/momenta for all of the particles. If there are no constraints imposed on the gas, meaning that it could have any energy, volume, etc., there are an infinite number of microstates available to the system. This is because we integrate over the entirety of state space. However, if we impose a constraint, such as that the total kinetic energy of the collection of particles is equal to some constant E , we constrain ourselves to a subset of state space with finite volume. We know that the energy of a single particle is $p^2/2m$. This means that a collection of N particles has a given energy when

$$\sum_{i=1}^N \sum_{j=1}^3 \frac{1}{2m} p_{ij}^2 = E$$

This equation defines a hypersphere in state space. The number of microstates available to a system with given energy is thus proportional to the measure of this hypersphere, whose radius is \sqrt{E} . Using the notation $\Omega(E)$ to mean “the number of states with total energy E ,” we can write that

$$\Omega(E) \propto E^{3(N-1)/2}$$

Thus, we see that increasing the total energy of the system increases the number of microstates which are consistent with that energy. Similarly, increasing the number of particles increases the number of available microstates.

4 Maxwell-Boltzmann Distribution

We now want to explore how a system of particles with position and momenta interacts with its environment. Figure 1 shows the “system,” which we will denote as A , embedded in an environment, which we denote as B . The total system of $A+B$ is isolated in the sense that its energy and volume are conserved. Denoting the total system as T , we can write basic relations such as

$$E_T = E_A + E_B, \quad V_T = V_A + V_B$$

Both the system and its environment are conceptualized as gases with corresponding micro- and macrostates. If we assume that the system and its surroundings only exchange energy but not volume, the number of states will only depend on the energy. Thus, the number of states available to the system of




Figure 1: A thermodynamic system interacts with its surroundings, which are considered to be large. The sum of the two can be treated as an isolated system, which does not interact with any environment and thus conserves energy and other quantities such as volume.

interest is $\Omega_A(E_A)$, and the number of states available to the environment is $\Omega_B(E_B)$. The previous section demonstrated how the number of states is a function of the energy. Now, we make the additional assumption that

$$\Omega_T = \Omega_A \Omega_B$$

which says that the microstates of the two systems are independent from one another. System A being in a particular state does not influence the microstates available to system B except through the energy relation $E_B = E_T - E_A$. Now we make a somewhat strange assumption that will prove to be useful. We argue that because the environment is large compared to the system of interest, its energy E_B does not change appreciably as the two systems interact. In other words, the relative changes in the energy of system A will be much larger than those of system B . We can then take the log of the unknown functional relation Ω and expand it as a Taylor series:

$$\log(\Omega_B(E_B)) \approx \alpha + \beta E_B$$

$$\implies \Omega_B(E_B) = e^{\alpha + \beta E_B}$$

It is not obvious why this is useful, and we do not know the coefficients α and β at this point. The question we will ask is: what is the probability that system A is in a particular microstate $\underline{s} = [x_1^1, x_2^1, x_3^1, p_1^1, p_2^1, p_3^1, \dots, x_1^N, x_2^N, x_3^N, p_1^N, p_2^N, p_3^N]$ with a given energy E_A ?

The number of microstates of the total system for which A is in a particular microstate is simply $\Omega_B * 1$. This is because system A and B are independent from one another except through energy exchange. So when the energy E_A goes down, there are less microstates available to A , thus the chance of finding it in a particular one increase. This is the same as saying that as E_A gets small, the energy E_B increases, which increases Ω_B . Thus, we can write

$$p(\underline{s}|E_A) \propto \Omega_B(E_B) = \Omega_B(E_T - E_A) = e^\alpha e^{\beta E_T} e^{-\beta E_A} \propto e^{-\beta E_A}$$

This means that the probability of finding system A in a particular microstate \underline{s} with energy E_A is proportional to $e^{-\beta E_A}$. To make this a true probability, we need a normalization constant. This normalization constant is supplied by the “partition function” Z . It is not exactly clear what we sum over in order to find this normalization constant. One thought is that the probability that the system is in some state at some energy is unity. This reads

$$\sum_E \sum_{\underline{s}|E} = \frac{1}{Z} e^{-\beta E_A} = 1 \implies Z = \sum_E \Omega_A(E) e^{-\beta E}$$

This is strange because the energy states E are not discrete, and it doesn't seem intuitive that system should be able to take on all energy states with equal probability. At the end of the day, the partition function Z is just some number and often turns out to not be important. We can now write

$$p(\underline{s}, E_A) = \frac{e^{-\beta E_A}}{Z}$$

This is the probability that A is in a specific microstate with energy E_A . What about the probability that A is in *any* microstate with energy E_A ? This is equivalent to asking what is the probability that system A has some energy, regardless of what microscopic configuration gives rise to that. This can be written by a simple counting argument as

$$p(E_A) = \Omega_A(E_A) \frac{e^{-\beta E_A}}{Z}$$

If we define a new quantity $S = \kappa \log(\Omega_A)$, we can write this probability as

$$p(E_A) = \frac{1}{Z} e^{-(\beta E_A - S/\kappa)}$$

Now we need to relate the unknown parameters β and κ to properties of the system of particles that are measurable.

5 Determining Parameters

We will ultimately relate the parameter β to the macroscopic quantity of temperature, and choose a convenient value of κ once β is determined. Before we do this, we need to establish a certain result: that the temperature is related to the average kinetic energy of a particle. We will do this for the simplified model of an ideal gas, though this should be a general result. The pressure in a gas is related to the momentum transfer from particles colliding with the walls of the container. The pressure the i component of velocity of a given particle exerts on the wall is proportional to $m\Delta v_i = 2mv_i$ because the collision reverses the sign of the velocity component. We do not know the time over which this

reversal takes place, but we can get around this. For particles only moving in the x direction and colliding elastically with a wall in the $y - z$ plane, the net force they exert on the wall is

$$F = pA = \int_A 2(m\bar{v}_x)(\bar{v}_x) \frac{N}{V} dydz$$

where V is the volume of the container and N is the number of particles. This expression says that an average collision imparts momentum $2m\bar{v}_x$ to the wall, and this happens at a rate proportional to the average speed, and the density of particles. This allows us to define the pressure as

$$p = \frac{N}{V} m\bar{v}_x^2$$

It is not clear where the factor of 2 goes. For general three-dimensional velocities, we define the mean square speed \bar{v}^2 such that $\bar{v}_x^2 = \bar{v}_y^2 = \bar{v}_z^2 = \frac{1}{3}\bar{v}^2$. This allows us to write the pressure as

$$p = \frac{N}{V} \frac{1}{3} m\bar{v}^2$$

Now we can recognize the average kinetic energy of the particle as $e = \frac{1}{2}m\bar{v}^2$. Rearranging and substituting this, we obtain

$$PV = \frac{2}{3}Ne$$

But the ideal gas law states that $PV = NkT$ where k is the Boltzmann constant. Thus, we have found that

$$e = \frac{3}{2}kT$$

The average kinetic energy of a particle with a three-dimensional velocity vector in the gas is related to the temperature. The energy is shared equally between the velocity components, so that the energy associated with one of these components is $kT/2$. We will use this result shortly.

We will now ask the same question of our statistical mechanical model of the gas: what is the average kinetic energy associated with one velocity degree of freedom in the gas? Note that for the gas, the total energy is

$$E = C \sum_{i=1}^N \sum_{j=1}^3 v_{ij}^2$$

where $C = m/2$ and i indexes the particle number while j indexes the velocity component. We want the average energy associated with a single degree of freedom v_{ij} . We can write the total energy as

$$E = E_1 + E_2 = Cv_1^2 + E_2$$

where this notation is meant to suggest that Cv_1^2 is the energy associated with the degree of freedom we are interested in, and E_2 is the remaining energy from all other velocity degrees of freedom. This expression parameterizes the previously derived probability distribution for the state and energy of the system. We can then say

$$p(\underline{s}, E(v_1, \dots)) = \frac{1}{Z} e^{-\beta(Cv_1^2 + E_2)}$$

Using this distribution, we now compute the expected value of E_1 :

$$\langle E_1 \rangle = \frac{1}{Z} \int_{\underline{s}} \int_{v_1} \int_{E_2} (Cv_1^2) e^{-\beta Cv_1^2} e^{-\beta E_2} dE_2 dv_1 d\underline{s}$$

This notation is fairly unclear. We have to integrate over all of the degrees of freedom defining the energy, and all of the states at a given energy. The integral over E_2 is meant to represent integration over all velocity degrees of freedom besides v_1 . The integral over states \underline{s} is interpreted as adding up all of the states associated with a given energy level. Abstractly, we think of the integral over the states as depending on the energy only through the limits of integration. The integrand is a uniform probability distribution by the assumption of all states being equally probable. The partition function when written with integrals takes a similar form:

$$Z = \int_{\underline{s}} \int_{v_1} \int_{E_2} e^{-\beta Cv_1^2} e^{-\beta E_2} dE_2 dv_1 d\underline{s}$$

We need to argue that the three integrals factorize. If somehow $\int_{\underline{s}} d\underline{s} = \Omega$, i.e. the number of states is independent of the energy level, we obtain the result that

$$\langle E_1 \rangle = \frac{\Omega \left(\int_{E_2} e^{-\beta E_2} dE_2 \right) \int_{v_1} Cv_1^2 e^{-\beta Cv_1^2} dv_1}{\Omega \left(\int_{E_2} e^{-\beta E_2} dE_2 \right) \int_{v_1} e^{-\beta Cv_1^2} dv_1} = \frac{\int_{v_1} Cv_1^2 e^{-\beta Cv_1^2} dv_1}{\int_{v_1} e^{-\beta Cv_1^2} dv_1}$$

It is not clear how exactly this result is obtained given that the number of states should depend on the energies E_1 and E_2 . This would mean that the integral over states does not neatly divide out. But this is the expression we are supposed to obtain. Note that this can be written compactly as

$$\langle E_1 \rangle = -\frac{\partial}{\partial \beta} \left(\log \left(\int_{v_1} e^{-\beta Cv_1^2} dv_1 \right) \right)$$

A simple change of variables integration ($y = \sqrt{B}v_1$) allows this expression to be dramatically simplified:

$$= -\frac{\partial}{\partial \beta} \left(\log \left(\beta^{-1/2} \int_y e^{-Cy^2} dy \right) \right) = -\frac{\partial}{\partial \beta} \left(\log(\beta^{-1/2}) + \log \left(\int_y e^{-Cy^2} dy \right) \right)$$

$$= \frac{1}{2\beta}$$

But we know from the ideal gas law that the average kinetic energy in a degree of freedom of a particle is also equal to $kT/2$. Thus we find that

$$\beta = \frac{1}{kT}$$

6 Free Energy

Previously, we derived that the probability of finding system A in a macrostate with energy E_A is

$$p(E_A) = \frac{1}{Z} e^{-(\beta E_A - S/\kappa)}$$

with $S := \kappa \log(\Omega_A(E_A))$. We have now determined the parameter β in terms of two familiar quantities: the Boltzmann constant k (empirically determined) and the temperature T (measurable macroscopic variable). This allows us to rewrite this probability as

$$p(E_A) = \frac{1}{Z} e^{-(E_A - SkT/\kappa)/kT}$$

Because S is a definition, we have the freedom to choose the constant κ . This expression makes clear that $\kappa = k$ is a convenient choice. It turns out S is the entropy, and it is conventional to choose the constant in the definition as the Boltzmann constant. With this choice, the probability of finding the thermodynamic system in a macrostate with given energy is

$$p(E_A) = \frac{1}{Z} e^{-(E_A - TS)/kT}$$

Note that because the parameter β was defined in terms of the environment (system B), the temperature T corresponds to the temperature of the environment (i.e. $T = T_B$). However, the entropy S is defined with respect to the number of microstates of system A , thus $S = S_A$. Invariably, it is confusing to keep track of the systems that the various quantities of interest correspond to. The temperature of the environment is fixed because it is large compared to the system of interest. Similarly, Z involves integration over all states and energies, and functions simply as a normalization constant. If we assume that the system has a specified energy E_A , it is natural to assume that the state we observe is the one which maximizes the probability $p(E_A)$. This is quite an interesting idea: *that the observed states of thermodynamic systems are the most probable ones. These maximally probable states define equilibrium configurations.* Because the temperature of the environment is fixed, maximizing the probability amounts to minimizing the term in parentheses within the argument of the exponent. Thus, the “equilibrium” configuration of the system is the minimum of

$$\Phi = E_A - TS$$

This is a “thermodynamic potential” called the Helmholtz free energy. Minimizing it involves a competition between the energy E_A and the entropy. When the energy E_A is small, the number of available states for the environment Ω_B is large, and that means there are many opportunities for A to take on a certain microstate. But because we are looking at macrostates, this effect is countered by the entropy, which tells us that there are fewer ways for A to have a given energy as that energy decreases. The contribution TS to the Helmholtz free energy corrects for the observation of macro-, as opposed to micro-states. To see that minimizing the free energy corresponds to equilibrium, observe that

$$\frac{\partial \Phi}{\partial E_a} = 1 - T_B \frac{\partial S_A}{\partial E_A} = 1 - T_B \frac{\partial}{\partial E_A} k \log(\Omega_A(E_A)) = 1 - \frac{T_B}{T_A}$$

where we are now using a general result that $\partial S_i / \partial E_i = 1/T_i$. Entropy and temperature are said to be thermodynamically conjugate variables. Equilibrium, interpreted as the maximum probability state of the system, is thus given by

$$T_A = T_B$$

7 Conclusion

We have seen that minimizing an energy to find the equilibrium configuration of a system is equivalent to maximizing the probability of observing a certain macrostate. This quite a profound result! By starting from high-dimensional random microscopic states, we have obtained macroscopic principles which govern the kinds of behavior we will observe thermodynamic systems exhibit. In order to extend this analysis to elasticity, we need to change our conceptions about the energy associated with the underlying material system. Solids are not collections of freely moving particles, rather these particles are bonded together. Thus, it may be necessary to include potential energy of bond stretching in addition to kinetic energy. Furthermore, solids deform and thus “exchange volume” with their surroundings. It can be seen that the “Gibbs free energy” characterizes a gaseous system which can vary in volume. In this case, work done on the system takes the form $W = \int p dV$. Internal energy of the system simply has to do with the temperature of the gas. In solids, however, we need energy storage in the form of elastic deformation. This means that the internal energy should include strain energy $\int \Psi dV$, which will likely be related to the potential energy of the bonded atoms. External work done on the system should take the form $W = \int_V b_i u_i dV + \int_S t_i u_i dS$. I suspect that a Gibbs free energy type thermodynamic potential should govern the equilibrium of an elastic body when a) we account for the fact that particles in a solid have kinetic and potential energy, b) we correctly prescribe energy storage in the form of strain energy in addition to temperature, and c) the form of external work is adjusted

to align with that of an elastic body. I see it as very profound that minimizing energy to find equilibrium is equivalent to, and maybe even derived from, the extremization of a high-dimensional probability distribution.