

1. **Wait, are we talking about magnets or fluids? Or both?!** In class we introduced the Ising model. We considered a lattice of N spins, each pointing either up ($s = 1$) or down ($s = -1$). The spins interact with an external magnetic field and with their neighboring spins such that the total energy is given by

$$E_{\text{Ising}} = -h \sum_i s_i - J \sum'_{i < j} s_i s_j, \quad (1)$$

where the prime on the summation sign is shorthand to indicate that we are only including spins for which i and j are at neighboring sites. Unlike our models with independent (decoupled) degrees of freedom, J causes the various spins to be coupled together. A consequence is that the partition function no longer factorizes into a product of single-particle terms. Rather,

$$Q(\beta, N, h) = \sum_{s_1=\pm 1} \sum_{s_2=\pm 1} \dots \sum_{s_N=\pm 1} e^{\beta h \sum_i s_i + \beta J \sum'_{i,j} s_i s_j}. \quad (2)$$

The Ising model is a lattice model, and in this class we have also talked about lattice models for gases. In particular, we have split up space into cells which are so small that they house zero or one particle. Unlike the Ising spins which take values ± 1 , each of our lattice gas cells took the value $n = 0$ or $n = 1$. You have always thought it was crazy that we acted like neighboring cells wouldn't influence each other. After all, if there is a van der Waals attraction between our gas molecules, then the presence of a molecule in one cell would decrease the energy if another molecule is in the neighboring cell. Let us model that attraction as an attractive energy

$$E_{\text{lattice gas}} = -\epsilon \sum'_{i < j} n_i n_j. \quad (3)$$

- (i) Show that this form of $E_{\text{lattice gas}}$ corresponds to a decrease in energy of ϵ when two neighboring cells i and j are occupied, but that it contributes nothing to the energy in the other three cases—when neither cell is occupied, only cell i is occupied, or only cell j is occupied.
 - (ii) A microstate ν of the lattice gas corresponds to a vector $(n_1, n_2, \dots, n_N) = (1, 0, \dots, 1)$ that specifies the state of every cell. Let us assume that the gas can exchange energy and particles with the outside world but that the volume is fixed. In that case the possible microstates can have a fluctuating value of $E_{\text{lattice gas}}$ and a fluctuating value of $\mathcal{N} = \sum_i n_i$. In terms of the grand canonical partition function $\Xi(\beta, N, \mu)$ for this lattice gas, what is the probability of microstate ν ?
 - (iii) Develop an explicit mapping between the Ising problem and the lattice gas problem to relate Q and Ξ . In particular, derive the precise relationship between the Ising parameters h , and J and the lattice gas parameters μ and ϵ . Assume the spins/cells are in d -dimensional space.
2. **The Ising model in one dimension.** One of the core predictions of the theory of phase transitions is that the ordering of microscopic degrees of freedom changes dramatically as a function of temperature. In the Ising picture, this dramatic change corresponds to a transition from aligned spins at low temperature to disordered spins at high temperature. In the lattice gas picture, it is a change from a liquid with $\langle n \rangle \approx 1$ to a gas with $\langle n \rangle \approx 0$. Here we will focus on the Ising magnet in one dimension (you could map these results over to the liquid/vapor context using problem 1).

- (i) First imagine that you already know how to compute $Q(\beta, N, h)$. Express $\langle s \rangle = N^{-1} \langle \sum_i s_i \rangle$ as a derivative of $\ln Q$.
- (ii) For the one-dimensional Ising model it is possible to perform that computation of Q via an accounting trick with the so-called transfer matrix

$$T = \begin{pmatrix} e^{\beta(J-h)} & e^{-\beta J} \\ e^{-\beta J} & e^{\beta(J+h)} \end{pmatrix} \quad (4)$$

Take, as a simple example, an Ising model with periodic boundary conditions and with $N = 3$ sites. Then Q is a sum over eight possible microstates: $(1, 1, 1), (1, 1, -1), \dots, (-1, -1, -1)$. Explicitly show that Q computed from Eq. (2) is equal to $\text{Tr}(T^3)$, where Tr is the trace of the matrix.

(iii) Let D be a diagonalizing matrix such that $\Lambda = D^{-1}TD$ is the diagonal matrix with the eigenvalues of T along the diagonal and zero off the diagonal. Using your earlier work and the cyclic property of the trace (which you need not explicitly prove here), show that $Q(\beta, N, h) = \text{Tr}(\Lambda^N)$.

(iv) The trace from (iii) results in a sum of two matrix elements. Argue that one of those terms will vanish in the limit of large N , to yield the large- N limit of $Q(\beta, N, h)$.

(v) Combine your answers to (i) and (iv) to determine $\langle s \rangle$ as a function of β in the absence of an external magnetic field ($h = 0$). If possible, identify the temperature of the phase transition between ordered and disordered states. Otherwise, explain why it is not possible to find such a temperature.

3. **Mean field theory.** Computing the exact partition function in more than one dimension is very hard (Lars Onsager famously did this in the 1940s for the two-dimensional $H = 0$ Ising model). But there are various ways to approximate Q . One strategy for approximating Q is to use a computer to sample many possible microstates, the focus of next week's problem set. The advantage of a computational approach is that it can be made arbitrarily accurate, meaning you could systematically improve your estimate for Q by spending more computer time. In contrast, there is an uncontrolled approximation scheme called mean field theory that has the benefit of being tractable with pen and paper.

The idea behind mean field theory is simple. Spin i has four neighbors, each of which can either be $+1$ or -1 . Assume instead that these neighbors all equal their average value $\langle s \rangle$. Then the effect of the neighboring spins is replaced by their average effect (the mean field that they impose on spin i). As a consequence, we approximate spin i as being decoupled from the other spins but in the presence of an augmented external field of strength $h + 4J\langle s \rangle$.

- (i) Express $\langle s \rangle$ in terms of a derivative of $\ln Q$.
- (ii) Find the canonical partition function for N independent spins, each in a field of strength $h + 4J\langle s \rangle$.
- (iii) Combine your results to (i) and (ii) to get a self-consistent expression for $\langle s \rangle$. In other words, obtain an equation of the form

$$\langle s \rangle = f(\langle s \rangle) \quad (5)$$

for some function f which you will find.

- (iv) Numerically solve your equation in (iii) for the $h = 0$ case to generate a plot of $\langle s \rangle$ versus temperature. Does this plot display a phase transition? If so, at what temperature?

The final problems involve Monte Carlo sampling of the Ising model in one and two dimensions. If you really don't want to use python and would rather write your own code from scratch, you can skip problems 4 and 5. Otherwise, problems 4 and 5 should help walk you through the programming.

4. Monte Carlo simulation of the one-dimensional Ising model.

Open Ising1D.ipynb to find my code for simulating and visualizing a 1D Ising model at inverse temperature β . This code simulates $L = 40$ spins with periodic boundary conditions. Let's pass through the code block by block.

First, you import some packages. Nothing significant to say here.

```
1 # Import some packages that we will use
2 import matplotlib as mpl
3 import matplotlib.pyplot as plt
4 import numpy as np
5 plt.style.use('default')
```

(i) Next, you define lists of 40 spins to represent the configuration.

```
1 # Store three different types of configurations as a demonstration
2
3 # down_Config will be a configuration with every spin equaling -1
4 # up_Config will be a configuration with every spin equally +1
5 # random_Config will be a configuration with every spin randomly set equal to
6 # either +1 or -1
7
8 L = 40; # Length of the 1d grid
9
10 # Make the down_Config. Note that np.zeros(L) will return a vector of all 0's
11 # and Python is smart enough to know that subtracting the scalar 1 from the
12 # vector should really subtract 1 from each element of the zero vector.
13 down_Config = np.zeros(L) - 1;
14
15 # Make the up_Config based on the down_Config
16 up_Config = np.copy(down_Config) + 2;
17
18 # Make the random_Config. I start by copying the down_Config because it is a
19 # quick and easy way to make the random_Config vector be the right have the
20 # correct size
21 random_Config = np.copy(down_Config);
22 for i in range(L):
23     ### EXPLAIN THE NEXT LINE ###
24     random_Config[i] = 2 * np.random.randint(2) - 1
25
26 print('Here is a configuration of all down spins:')
27 print(down_Config)
28 print('Here is a configuration of all up spins:')
29 print(up_Config)
30 print('Here is a configuration of random spins:')
31 print(random_Config)
```

Explain what is happening in line 24. What is randint doing and why are we multiplying by 2 and subtracting one?

(ii) Next we define a function that graphically plots the spins to make them easier to visualize.

```
1 # Write a function to plot the spin configuration
2 def plotConfig(config, L):
3     ### EXPLAIN THE NEXT LINE ###
4     plt.imshow(np.resize((config + 1)/2, (1,L)), vmin=0, vmax=1, \
5                 cmap=plt.cm.Greys);
6     plt.axis('off');
7
```

```

8 # Use the plotting function
9 plotConfig(random_Config, L)

```

Explain what is happening in lines 4-5. Specifically, why do I use $(\text{random_Config} + 1)/2$ and not just random_Config ?

(iii) To compute the energy of a spin configuration, we will need to be able to sum over the nearest neighbors. Toward this end, it would be nice to have a function that returns the location of the neighboring spins. For example, I'd like to be able to know that spin number 5 is next to spins 4 and 6. It seems this function would be trivial, it should just return $i + 1$ and $i - 1$, but the following code is a little more complicated. Here is one way.

```

1 # Write a function to compute the index for spin i's neighbors
2 def neighbors1d(i, L):
3     ### EXPLAIN THE NEXT LINES ###
4     if i==0:
5         neighbor1 = L-1
6     else:
7         neighbor1 = i-1
8     if i==L-1:
9         neighbor2 = 0
10    else:
11        neighbor2 = i+1
12    return [neighbor1, neighbor2]

```

Explain the purpose of the if/else statements. Why did we not just set neighbor1 to $i - 1$ and set neighbor2 to $i + 1$?

(iv) Next we define two different energy functions.

```

1 # Compute the contribution of spin i to the total energy
2 def energy1spin(i, L, h, J, config):
3     energy = -h * config[i]
4     for neighbor in neighbors1d(i,L):
5         # Notice that we have J/2 rather than J since we will count up
6         # the i-j interaction for spin i *and* for spin j.
7         energy -= (J/2) * config[i]*config[neighbor]
8     return energy
9
10 # Compute the total energy of a configuration
11 def energy(L, h, J, config):
12     energy = 0
13     for i in range(L):
14         energy += energy1spin(i, L, h, J, config)
15     return energy

```

The first energy function computes the contribution from spin i . The second function sums over all of the possible i values to give the total energy. You will see why I separately defined `energy1spin` later. For now, we check that the energy function is working properly.

```

1 # Let's check that the energy function is working correctly
2
3 LCheck = 4
4 config1 = [1, 1, 1, 1]
5 config2 = [1, -1, 1, -1]
6 config3 = [1, -1, -1, -1]
7
8 print('When h = 0 and J = 1...')
9 print('The energy of ', config1, ' is ', energy(LCheck, 0, 1, config1))

```

```

10 print('The energy of ', config2, ' is ', energy(LCheck, 0, 1, config2))
11 print('The energy of ', config3, ' is ', energy(LCheck, 0, 1, config3))
12
13 print('\nWhen h = 1 and J = 0...')
14 print('The energy of ', config1, ' is ', energy(LCheck, 1, 0, config1))
15 print('The energy of ', config2, ' is ', energy(LCheck, 1, 0, config2))
16 print('The energy of ', config3, ' is ', energy(LCheck, 1, 0, config3))

```

Work out the energy of the 4-spin configurations config1, config2, and config3 by hand to compare with the output of the code.

(v) Finally we run the main Monte Carlo loop.

```

1 # This is the main Monte Carlo loop!
2
3 # Start with a random configuration
4 config = np.copy(random_Config);
5
6 # Set up the Hamiltonian you want to sample
7 J = 1;
8 h = 0;
9 beta = 1;
10
11 # Store the initial energy
12 old_energy = energy(L, h, J, config);
13
14 # Configurations will be added to this empty list so you can make a movie
15 # of the visited configurations
16 stored_configs = []
17
18 # Perform 100000 Monte Carlo steps
19 num_mc_steps = 100000;
20
21 # Create a vector to store the values of m
22 mag_values = np.zeros(num_mc_steps);
23
24 # Loop over the Monte Carlo steps
25 for mcstep in range(num_mc_steps):
26
27     # Pick a random spin to flip
28     index_for_spin_flip = np.random.randint(L);
29     # Flip that spin
30     config[index_for_spin_flip] *= -1;
31
32     # Compute the energy for the new configuration
33     new_energy = energy(L, h, J, config);
34
35     # Compute the change in energy due to the spin flip
36     delta_energy = new_energy - old_energy;
37
38     # Calculate the probability of accepting the move
39     acceptance_probability = min(1, np.exp(-beta * delta_energy))
40
41     # Draw a random number between 0 and 1 to determine whether or not to
42     # accept the move
43     if np.random.random() < acceptance_probability:
44         # we already flipped the spin, so we only need to update the energy
45         old_energy = new_energy;
46     else:

```

```

47     # Since the spin flip is rejected we retain the old energy but need to
48     # flip the spin back
49     config[index_for_spin_flip] *= -1;
50
51     # compute m for the configuration
52     mag_values[mcstep] = np.sum(config) / len(config);
53
54     # Store every 1000 configurations. You could adjust this if you feel like you
55     # have too many or too few frames
56     if (mcstep % 1000 == 0):
57         stored_configs.append(np.copy(config))

```

Discuss what line 43 is doing.

(vi) The remaining blocks of code plot the magnetization and generate a movie of configurations that are visited during the Monte Carlo sampling. My default parameters used $\beta = 1, J = 1, h = 0$. **Based on problem 2 (with the transfer matrices), do you expect the typical configurations to be ordered ($m = \pm 1$) or disordered ($m = 0$)?** Run the code to check.

5. **Monte Carlo simulation of the two-dimensional Ising model.** Open Ising2D.ipynb to find the basic framework for code that will simulate and visualize a 2D Ising model at inverse temperature β . This code simulates $L^2 = 20 \times 20$ spins with periodic boundary conditions. Let's pass through the code block by block. This time there will be some parts that you need to write to make it work!

We start with a very similar structure but now the list of spins must have L^2 elements. I store the spins in a single 1D list, which will require us to think a little bit when figuring out which spins are neighbors on the 2d grid. The following block of code would print out some of these lists for the configurations with 20×20 spins.

```

1  # Import some packages that we will use
2  import matplotlib as mpl
3  import matplotlib.pyplot as plt
4  import numpy as np
5  plt.style.use('default')
6
7  # Store three different types of configurations as a demonstration
8
9  # down_Config will be a configuration with every spin equalling -1
10 # up_Config will be a configuration with every spin equally +1
11 # random_Config will be a configuration with every spin randomly set equal to
12 # either +1 or -1
13
14 L = 20; # Length of the 2d grid
15
16 # Make the down_Config. Note that np.zeros(L*L) will return a vector of all 0's
17 # and Python is smart enough to know that subtracting the scalar 1 from the
18 # vector should really subtract 1 from each element of the zero vector.
19 down_Config = np.zeros(L*L) - 1;
20
21 # Make the up_Config based on the down_Config
22 up_Config = np.copy(down_Config) + 2;
23
24 # Make the random_Config. I start by copying the down_Config because it is a
25 # quick and easy way to make the random_Config vector be the right have the
26 # correct size
27 random_Config = np.copy(down_Config);
28 for i in range(L*L):
29     random_Config[i] = 2 * np.random.randint(2) - 1

```

```

30
31 print('Here is a configuration of all down spins:')
32 print(down_Config)
33 print('Here is a configuration of all up spins:')
34 print(up_Config)
35 print('Here is a configuration of random spins:')
36 print(random_Config)

```

I also updated the plotting function that will let you look at a 2D grid of spins.

```

1 # Write a function to plot the spin configuration
2 def plotConfig(config, L):
3     # The only thing that has changed from the 1D code is that (1, L)
4     # became (L, L)
5     plt.imshow(np.resize((config + 1)/2, (L,L)), vmin=0, vmax=1, \
6                 cmap=plt.cm.Greys);
7     plt.axis('off');
8
9 # Use the plotting function
10 plotConfig(random_Config, L)

```

(i) As I indicated, it will be a little harder to compute the neighbors in 2D. We can leverage the 1D function that we already wrote in the following manner.

```

1 # Write a function to compute the index for spin i's neighbors
2 def neighbors1d(i, L):
3     if i==0:
4         neighbor1 = L-1
5     else:
6         neighbor1 = i-1
7     if i==L-1:
8         neighbor2 = 0
9     else:
10        neighbor2 = i+1
11    return [neighbor1, neighbor2]
12
13 # Write a function to compute the index for spin i's neighbors
14 def neighbors2d(i, L):
15     row = int(np.floor(i / L))
16     column = i - (row * L)
17
18     [leftcol, rightcol] = neighbors1d(column, L)
19     [downrow, uprow] = neighbors1d(row, L)
20
21     upneighbor = (uprow * L) + column;
22     downneighbor = # FILL THIS IN;
23     leftneighbor = # FILL THIS IN;
24     rightneighbor = (row * L) + rightcol;
25     return [leftneighbor, rightneighbor, upneighbor, downneighbor]

```

Fill in the right hand side of lines 22 and 23 to complete the neighbors2d function. You may want to run the function with some test cases to confirm that it works properly.

(ii) The energy computation can be implemented in a manner that is *very* similar to 1D case.

```

1 # Compute the contribution of spin i to the total energy
2 def energy1spin(i, L, h, J, config):
3     energy = -h * config[i]
4     for neighbor in neighbors2d(i,L):
5         # Notice that we have J/2 rather than J since we will count up

```

```

6     # the i-j interaction for spin i *and* for spin j.
7     energy -= (J/2) * config[i]*config[neighbor]
8     return energy
9
10    # Compute the total energy of a configuration
11    def energy(L, h, J, config):
12        energy = 0
13        for i in range(L*L):
14            energy += energy1spin(i, L, h, J, config)
15        return energy

```

Identify and explain every difference between these two energy functions and the 1D versions of Problem 1(iv).

(iii) As in Problem 1(iv), we want to check that our energy is computed correctly. Since you're really using my energy function from part (ii), you can view this as a test of your nearest neighbor function from part (i).

```

1  # Let's check that the energy function is working correctly
2
3  config1 = np.zeros(9)+1
4  config2 = np.copy(config1)
5  for spin_index in range(9):
6      config2[spin_index] = np.random.randint(2)*2 - 1
7
8  print('When h = 0 and J = 1...')
9  print('The energy of ', config1, ' is ', energy(3, 0, 1, config1))
10 plotConfig(config1, 3)
11 plt.show();
12 print('The energy of ', config2, ' is ', energy(3, 0, 1, config2))
13 plotConfig(config2, 3);
14 plt.show();
15
16 print('\nWhen h = 1 and J = 0...')
17 print('The energy of ', config1, ' is ', energy(3, 1, 0, config1))
18 plotConfig(config1, 3);
19 plt.show();
20 print('The energy of ', config2, ' is ', energy(3, 1, 0, config2))
21 plotConfig(config2, 3);
22 plt.show();

```

Work out the energy of the 9-spin configurations config1 and config2 by hand to compare with the output of the code.

(iv) The Monte Carlo steps proceed in a manner that is very similar to the 1D code, but we will want to make one critical change. It will be wasteful to compute the full energy before and after a spin flip. The reason for this is that every computation of an energy requires you to add up the contribution from spin 1, from spin 2, from spin 3, etc., but in a single move you are only changing a single spin! We should be able to use our function energy1spinflip to more rapidly determine delta_energy.

```

1  # This is the main Monte Carlo loop!
2
3  # Start with a random configuration
4  config = np.copy(random_Config);
5
6  # Set up the Hamiltonian you want to sample
7  J = 1;
8  h = 0;

```



```

9 beta = 1;
10
11 # Store the initial energy
12 old_energy = energy(L, h, J, config);
13
14 # Configurations will be added to this empty list so you can make a movie
15 # of the visited configurations
16 stored_configs = []
17
18 # Perform 100000 Monte Carlo steps
19 num_mc_steps = 100000;
20
21 # Create a vector to store the values of m
22 mag_values = np.zeros(num_mc_steps);
23
24 # Loop over the Monte Carlo steps
25 for mcstep in range(num_mc_steps):
26
27     # Pick a random spin to flip
28     index_for_spin_flip = ### FILL THIS IN ###
29
30     # Compute the old contribution to the energy from spin i
31     old_energy = energy1spin(index_for_spin_flip, L, h, J, config)
32
33     # Flip that spin
34     ### FILL THIS IN ###
35
36     # Compute spin i's contribution to energy for the new configuration
37     new_energy = ### FILL THIS IN ###
38
39     # Compute the change in energy due to the spin flip
40     # Notice the factor of 2, which was needed because we have to count the way
41     # flipping i changes the (J/2) factor that energy1spin attributed to spin i
42     # but also the (J/2) factor that energy1spin attributed to spin j
43     delta_energy = 2 * (new_energy - old_energy);
44
45     # Calculate the probability of accepting the move
46     ### FILL THIS IN ###
47
48     # Draw a random number between 0 and 1 to determine whether or not to
49     # accept the move
50
51     ### FILL THIS IN (Several lines) ###
52
53     # compute m for the configuration
54     mag_values[mcstep] = np.sum(config) / len(config);
55
56     # Store every 1000 configurations. You could adjust this if you feel like you
57     # have too many or too few frames
58     if (mcstep % 1000 == 0):
59         stored_configs.append(np.copy(config))

```

Fill in the missing code to perform Metropolis Monte Carlo sampling for the 2D Ising model. To check your code, run the following block to see how the magnetization per site evolves as a function of the number of Monte Carlo steps.

```

1 # Plot the values of the magnetization, which have been stored throughout the
2 # Monte Carlo sampling in the vector mag_values
3 plt.plot(mag_values)

```

```

4 plt.xlabel(r'Monte Carlo Step', fontsize=18);
5 plt.ylabel(r'$m$', fontsize=18);
6 plt.title(r'Fluctuating values of the magnetization per site', fontsize=20);

```

Turn in your plot but not your code.

(v) In analogy with Problem 1(vi), you can run the final blocks of code in Ising2DFramework.ipynb to see the probability distribution for the magnetization per site and to generate a movie of typical configurations. If you look carefully, you will see that I only collect the samples after 40,000 Monte Carlo steps to make sure that the initial configuration does not strongly bias the sampling. **Discuss the major difference between the behavior of the 2D Ising model and the 1D behavior you observed (at the same temperature) in Problem 1(vi).**

6. Using your Ising simulation.

(i) In the absence of an external magnetic field (i.e., $h = 0$), the transition to a ferromagnetic state upon decreasing temperature is second order. In other words, the average magnetization per spin $\langle m \rangle$ changes continuously but not smoothly at $T = T_c$. Try to confirm this fact by computing $\langle m \rangle$ as a function of T from several Monte Carlo trajectories above and below T_c . Estimate the sampling error associated with each computed value of $\langle m \rangle$, and describe the procedure you use to do so. **Provide a plot of m as a function of T (showing error bars), along with snapshots of the system at a temperature below T_c , a temperature very close to T_c , and a temperature above T_c .**

(ii) Your results will differ from the exact behavior of this system (determined mathematically by Lars Onsager) in several respects. For example, the apparent transition in your simulated system should not occur precisely at $k_B T_c = 2.269J$. (Note that, in the code, the unit of temperature is J/k_B .) The transition should also appear smooth (rather than exhibiting a sharp kink). **Explain both of these deviations.**

(iii) [Optional] Singularities at the critical point arise from correlated fluctuations of spins separated by macroscopically large distances. Try to confirm the existence of long-ranged correlations at the critical point by computing $c(r_{ij}) = \langle s_i s_j \rangle - \langle s_i \rangle \langle s_j \rangle$ as a function of $r_{ij} = |\mathbf{r}_i - \mathbf{r}_j|$ (the distance between spins i and j) at several temperatures above and below the critical point. **Plot your results.** [Note: To accurately sample with $T \approx T_c$ you have to generate quite a few Monte Carlo samples, so it's a little harder to carry out this problem. I've made it optional, but I encourage those who expect to use computers in their research to embrace the challenge.]