Numerical Solutions to the Ising Model using the Metropolis Algorithm

Danny Bennett
JS TP - 13323448

January 17, 2016

Abstract

Solutions to various versions of the Ising model were obtained using the Metropolis algorithm. The 2D square lattice was initially considered. After successfully using the Metropolis algorithm to update the system, the average energy per spin, average magnetisation per spin, specific heat capacity and magnetic susceptibility were plotted as functions of temperature in order to gain information about the system. The ground state energy was determined to be -2J, as expexted, and the Curie temperature was determined to be $T_C = 2.6 \pm 0.1\text{J/k}_\text{B}$ in this case, which compares well with the accepted value of $T_C = 2.269\text{J/k}_\text{B}$, and the ground state energy was determined to be -2J. The triangular lattice was then investigated; the ground state energy was determined to be -3J, and the Curie temperature was determined to be $T_C = 4.2 \pm 0.1\text{J/k}_\text{B}$. For the 1D system, a phase transition was initially observed, but this was due to a low value of J used in computations; when a larger value of J was used, there was no phase transition, which agrees with the theory. The 3D system was investigated and also determined to be ferromagnetic, with a larger Curie temperature of $T_C = 4.4 \pm 0.1\text{J/k}_\text{B}$, and a ground state energy of -3J. Finally, the methods developed in the previous parts of the project were used to investigate a simplified 2D model of NiO. The ground state was determined to be -36.75 ± 0.01meV, and the ordered state of the system was determined to be antiferromagnetic.

1 Introduction and Theory

1.1 The Ising Model

The Ising model is a model used in statistical mechanics, typically to simulate magnetic systems. It consists of a lattice with a discrete value σ_j ($\in \{-1, 1\}$, in the case of this project) assigned to each point, representing the spin at the site (down or up). The Hamiltonian of a system with a given configuration of spins is given by

$$H = -\sum_{\langle i,j \rangle} J_{ij}\sigma_i\sigma_j - B\sum_j \sigma_j, \tag{1}$$

where the first sum is over the interacting spin pairs and the second is over the entire lattice, J_{ij} is the exchange energy between the spins and B is an external magnetic field. If J_{ij} is positive, the neighbouring spins will tend to align; the material is ferromagnetic. If it is negative, they will tend to oppose each other; the material is antiferromagnetic. In the absence of an external magnetic field, $B = 0$, and so the Hamiltonian reduces to

$$H = -\sum_{\langle i,j \rangle} J_{ij}\sigma_i\sigma_j. \tag{2}$$

Ferromagnetism occurs in some magnetic materials below the Curie temperature, T_C. Electron spins of atoms in microscopic regions (domains) are aligned. The magnetic field is strong in the domain, but since in the material the domains will be randomly ordered with respect to one another, it will usually be unmagnetised. In the presence of an applied magnetic field, the domains may line up with each other, causing

the material to become magnetised. If the applied field is removed, the domains remain aligned, meaning ferromagnetic materials can exhibit spontaneous magnetisation (remanence): a net magnetic moment in the absence of an external magnetic field. Above the Curie temperature, the thermal motion is sufficient to disrupt the alignment, and the material becomes paramagnetic.

A number of simplifcations are made to the Ising model in this project. It is assumed that the exchange energy term is constant in all cases and is positive: $J_{ij} \equiv J > 0$. The 2-D Ising model is initially considered for a square lattice, which can be easily represented by an $N \times N$ matrix. It is also assumed that the nearest neighbour interactions are dominant and that longer range interactions are neglected in the Hamiltonian. These simplifications allow the Ising model to be easily implemented numerically, and the exact solution to this model was developed by Onsager in 1944[4], meaning the validity of the results can determined by comparing them to the exact solutions.

This simplified 2-D Ising model is implemented using the Metropolis algorithm: a modified Monte Carlo method. A standard Monte Carlo method would involve picking a number of states at random and making measurements, then weighting them by their Boltzmann factor, although this would require a large number of measurements to be significantly accurate and would be very expensive computation-wise. The Metropolis algorithm picks the states based on their Boltzmann factors and they are weighted equally.

1.2 Calculating Observables

Once the system has been simulated, the aim is to investigate its properties by calculating various observables and determine whether or not it undergoes a phase transition. This can be done by calculating several observables and investigating their properties; in particular, how they depend on the temperature of the system. The average energy per spin is determined by

$$\langle E \rangle = \left\langle \sum_{\langle i,j \rangle} H_{ij} \right\rangle = \frac{1}{2} \left\langle \sum_{i,j} H_{ij} \right\rangle, \tag{3}$$

the $\frac{1}{2}$ accounting for the fact that every pair is counted twice in the sum. If the system is ferromagnetic, then the expected value of $\langle E \rangle$ for the 2D square lattice would be -2J when the spins are all aligned. $\langle E \rangle$ should be a continuous function of temperature, with a point of inflection at T_C, and should approach 0 after this point, signifying that a phase change has ocurred and the material has become disordered and paramagnetic.

The average magnetisation per unit spin is given by

$$\langle M \rangle = \frac{1}{N^2} \sum_{(i,j)} \sigma_{ij}. \tag{4}$$

Calculating M can allow the time taken for the system to reach equilibrium to be determined for a given system size. By plotting the average magnetisation per spin after each implementation of the Metropolis algorithm against the number of iterations, the average magnetisation should be seen to approach a constant value signifying that equilibrium has been achieved, and the minimum number of iterations required to achieve is thus determined, which will help to increase the efficiency of the algorithm. The magnetisation can also be used to show that a phase transition has occurred in the material and approximate the Curie temperature, T_C. By plotting the average magnetisation after equilibrium has been reached against a series of values of temperature, a phase transition should be observed at T_C. The magnetisation should be 1 or -1 initially as it is ferromagnetic, but after T_C it should be zero, since it becomes paramagnetic.

The specific heat capacity, C_V, is given by

$$\begin{aligned} C_V &= \frac{\partial \langle E \rangle}{\partial T} \\ &= -\frac{\beta}{T}\frac{\partial \langle E \rangle}{\partial \beta} \\ &= \frac{\beta}{T}\frac{\partial^2 \ln(Z)}{\partial \beta^2} \\ &= \frac{\beta}{T}\left(\langle E^2 \rangle - \langle E \rangle^2\right), \end{aligned}$$

where $\beta = (k_B T)^{-1}$ is the inverse energy and Z is the partition function. Similarly, the magnetic susceptibility is given by

$$\begin{aligned} \chi &= \frac{\partial \langle M \rangle}{\partial H} \\ &= \beta\left(\langle M^2 \rangle - \langle M \rangle^2\right); \end{aligned}$$

$\langle E^2 \rangle$ and $\langle M^2 \rangle$ can be computed in a similar fashion to (3) and (4). Since C_V and χ are second derivatives of the free energy, when plotted against temperature they should exhibit a discontinuity at T_C, which would indicate that the phase transition that has occured is of second order.

1.3 Triangular Lattice

After a few adjustments, the Metropolis algorithm can be applied to a 2D triangular lattice. A triangular lattice is similar to a square lattice, the main difference being that every second row of points is horizontally offset, causing three neighbouring points to form an equilateral triangle. As a result, each point has six nearest neighbours, forming a regular hexagon. Since each lattice site has two extra nearest neighbours, the ground state energy is expected to be $\langle E \rangle = -3\text{J}$ using the same logic as in the case of the 2D square lattice.

1.4 Other Dimensions

There is no phase transition for the 1D Ising model with nearest neighbour interactions. This is a result of the fact that 1D systems with short range interactions cannot have a phase transition, which is known as van Hove's theorem[6]. One of the most famous arguments for this was made by Landau[3]: Consider a 1D lattice with N sites, each site either having spin "up" or spin "down", and that energy is minimised when all of the sites are in the same state. Assume that a boundary separates a region of "up" spins from "down" spins, and costs an energy ϵ. The energy for having n such boundaries is $E = n\epsilon$, and the entropy is $S = n\ln\left(\frac{n}{N} - 1\right)$, due to the number of ways to place the n boundaries among the N sites. The free energy is then given by

$$F = E - TS = n\epsilon - nT\ln\left(\frac{n}{N} - 1\right). \tag{5}$$

Therefore F is minimised in the limit $n \to L$, i.e. it is energetically more favourable to add another domain. So more and more boundaries will be added and there can therefore be no phase transition. This argument only works for a non-zero temperature, and thus numerically a phase change may be observed for low enough values of J, but should not be observed when using larger values.

The 3D Ising model may not be particularly enlightening, physically, but the comparision between it and the 2D Ising model could prove interesting, and it should prove to be an interesting exercise in programming.

2 Method

2.1 The Metropolis Algorithm

The Metropolis algorithm was implemented as follows:

- An $N \times N$ matrix is used to represent a 2-D square lattice; every entry of the matrix represents a lattice site and has entry ± 1, representing the spin of the particle at that site. A random matrix is initially used so that the initial configuration of spins is randomised.
- For this system we should have $N \to \infty$, meaning very large N, but this is very expensive computation-wise. One way to increase the accuracy of the results without increasing the size of the matrix is to use periodic boundary conditions. This is easily achieved by using modulo operator: taking the neighbouring sites to be $((i+1) \bmod (N), j)$, etc., instead of $((i+1), j)$, etc.
- A lattice site is selected at random and ΔE is calculated. If $\Delta E \leq 0$, the change is favourable and the spin is flipped. If $\Delta E > 0$, then the spin is flipped only if $\exp(-\Delta E / k_B T) > x$, where x is a random number on the interval $[0, 1]$ (i.e. the particle has enough energy to flip its spin). This process is repeated for every site in the lattice, selecting the sites at random.
- A new matrix is returned, which represents the updated state of the system.

A flow-chart for this algorithm is provided below.

The number of iterations of this process required to achieve equilibrium for a given system size was determined as follows: the process was repeated iteratively and the average magnetisation per spin was recorded in each case. A plot of the average magnetisation per spin vs the number of iterations was produced, and the magnetisation was seen to approach a constant value. From this a sufficient number of iterations required to reach equilibrium for a given system size was determined.

In terms of programming, the code was written as a module consisting of a class containing a series of simple functions so that it could be read and updated easily. The functions were used to: define a random matrix, calculate H_{ij} for a lattice site, calculate ΔE for a lattice site, flip the spin of a lattice site and sweep through the lattice, picking sites at random. The main part of the code is included in the appendix.

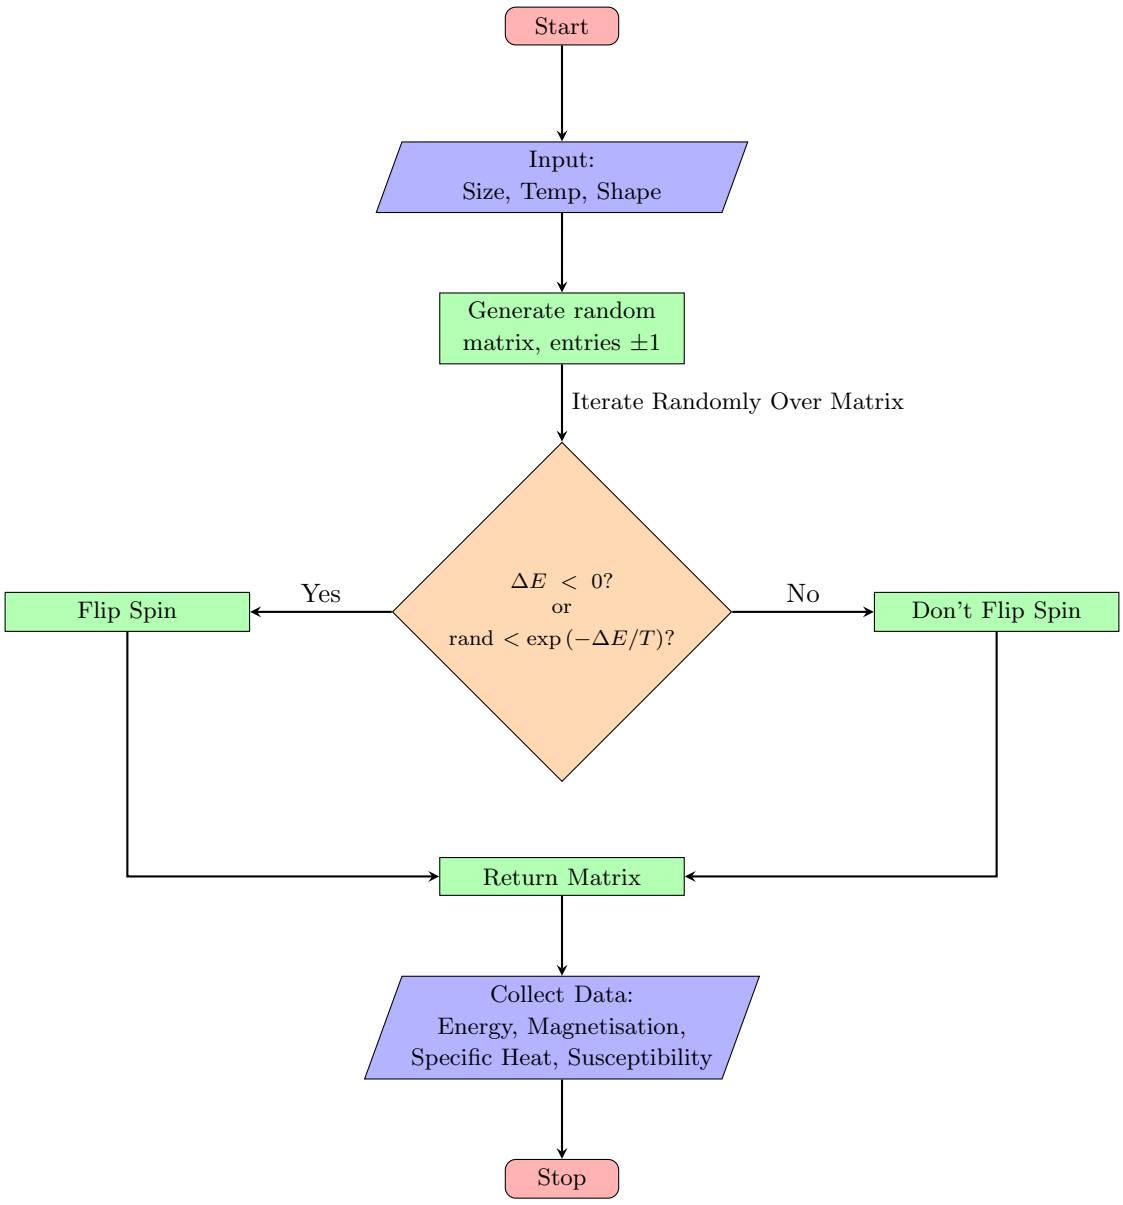

Figure 1: A flow-chart for one sweep of the Metropolis algorithm.

2.2 Calculating Observables

After the system is successfully updated, the various observables need to be calculated. Functions were defined as part of the main class to determine the average energy per spin, average magnetisation per spin, specific heat capacity and magnetic susceptibility of the system. The algorithm was used to bring the system to equilibrium over a temperature range, and these observables were recorded and plotted as a function of temperature.

Due to the random nature of the system, the plots contain large fluctuations after just one iteration, and an average of a large number of simulations was required to produce acceptable results. The process was repeated 2000 times and and an average was taken, however this method was very expensive, computation-wise. To reduce computation time, program was run simultaneously on multiple computers.

2.3 Triangular Lattice

In terms of programming, a triangular lattice would be difficult to work with, and it is therefore reasonable to look for a way to represent one by a square matrix. The following bijection can be used to relate a square and triangular lattice: shift every second row horizontally so that it is aligned with the rows above and below it, the nearest neighbours now being the points horizontally above and below the lattice site (as before), as well two of the closest diagonal points. Therefore the triangular lattice can be easily represented by adding two extra nearest neighbours to the existing represenation of the square lattice.

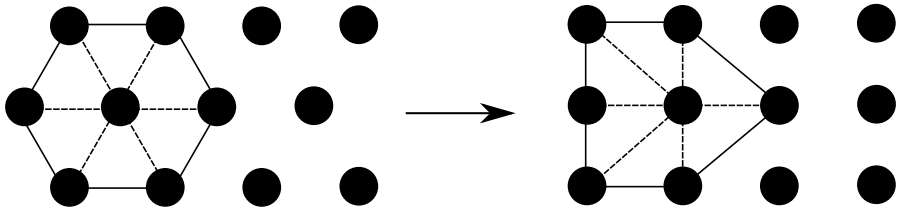

Figure 2: The triangular lattice (left), and the square matrix (right) which represents it in the code.

2.4 Other Dimensions

To preform an analysis of the Ising model in different dimensions, it was efficient to write one main module which worked for any spatial dimension. The main class was thus written for a 3D lattice, which could then be reduced to 2D or 1D lattices as required.

2.5 Example: A 2D model of NiO

The methods developed in the preceeding parts of the project were employed to simulate a simplified model of NiO. The material was described by a 2D square lattice having nearest and second nearest neighbour interactions with interaction energies of $J_1 = 2.3\text{meV}$ and $J_2 = -21\text{meV}$, respectively. As the oxygen atoms have no magnetic moment, only the nickel atoms were considered. The system was represented by a matrix with zeros at every second entry (the oxygen atoms) and ± 1 at every other entry (the nickel atoms with either spin "up" or "down"). Since the dominant interaction energy term is $J_2 < 0$, the system should in theory be antiferromagnetic. The expected ground state energy for the system is -42meV. When the system is completely misaligned, two of the nearest neighbours will have spin "up" and two will have spin "down", meaning their contributions to the energy will cancel. The contribution then comes from the four second nearest neighbours, giving -84meV, and dividing by two to account for double counting, the average

energy per spin should be −42meV. The main class from before was modified to account for second nearest neighbour interations with different interaction energies and used to investigate the simplified 2D NiO model.

3 Results and Analysis

The metropolis algorithm was successfully used to sweep through a matrix and update the spins appropriately. An example is shown below.

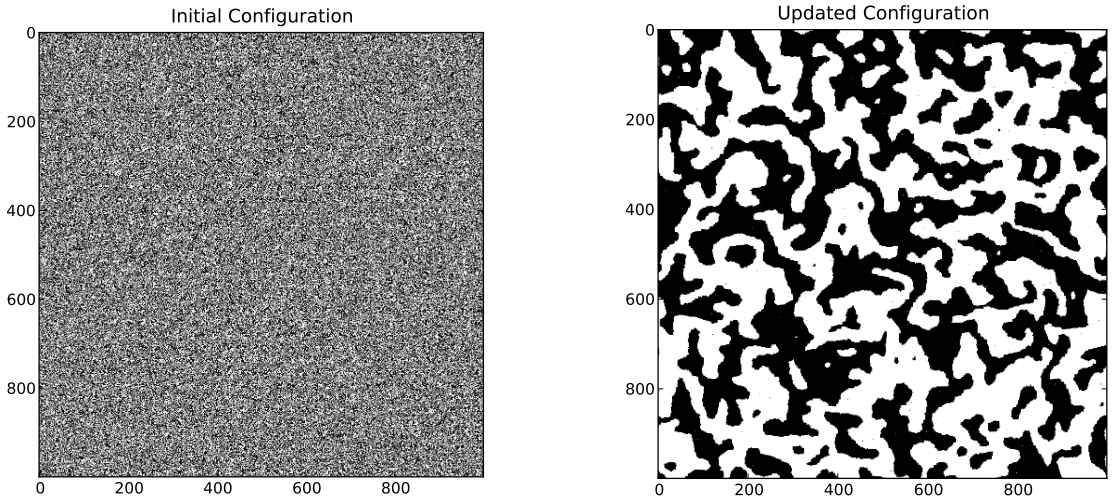

Figure 3: The initially randomised system (left) compared with the updated system (right). This system was represented by a 1000×1000 matrix at a temperature of $T = 1\mathrm{J/k_B}$, and the updated state shown is the result of 100 iterations of the Metropolis algorithm.

This system should be ferromagnetic (since $T = 1\mathrm{J/k_B} < T_C$), and this is clear from the updated system, as the spins are beginning to form domains. By inspection it is also evident that the periodic boundary conditions were correctly implemented. This result raises two issues: this system is not yet in equilibrium, so the required number of iterations of the Metropolis algorithm for a given system size needs to be determined. Also the matrix used in the above example was very large, and resulted in long computation times. In order to keep the computation times as short as possible, a matrix of smaller dimensions should be used. This will affect some aspects of the results (steepness and height of the curves), but the value of T_C and the numerical accuracy of the results should remain accuracte, to a degree. Also since the system appears to align at low temperatures, when taking measurements over a range of temperatures, an initially aligned matrix can be used when starting with a low temperature, which should slightly reduce the fluctuations in the resulting plots.

The Metropolis algorithm was applied to a random matrix a large number of times, and the average magnetisation per spin of the system was recorded for each step. A plot of magnetisation vs number of steps showed that after a certain number of steps, the magnetisation was approximately constant. An example is shown below.

From the plot it can be seen that a 100×100 matrix requires approximately 500 sweeps before it reaches equilibrium. In the proceeding parts of the project, a smaller matrix is used, as the results require a large number of simulations to be averaged, and a 100×100 matrix is not feasible for this; a 10×10 matrix only requires approximately 25 sweeps to reach equilibrium, and therefore much shorter computation times.

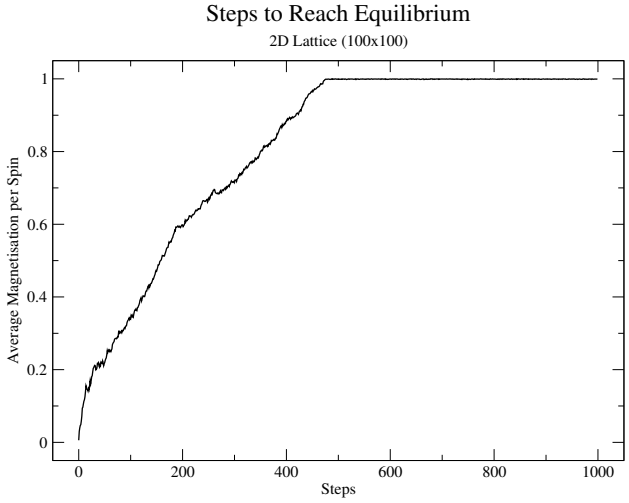

Figure 4: The magnetisation at each sweep of the Metropolis algorithm. From the plot it is clear that for a 100×100 matrix, approximately 500 sweeps are required to bring the system to equilibrium.

The observables for the 2D square lattice were then calculated. The program was run on multiple computers simultaneously using ssh to reduce the computation time. To prevent each computer from overwriting the data produced by the previous one, the time at which the process completed was included in the file name. The files were appropriately renamed using the rename feature in the terminal, so they could be averaged easily. Shown are the average magnetisation per spin, average energy per spin, specific heat capacity and magnetic susceptibility of the 2D square lattice, all plotted as a function of temperature.

As can be seen in the plot of magnetisation vs temperature, the magnetisation suddenly drops at $T_C \approx 2.6 \pm 0.1 \mathrm{J/k_B}$, which compares well with the theoretical value[2], and indicates that a phase transition has occured. The energy is a continuous function of temperature, as expected. As the temperature approaches 0, the energy approaches a ground state energy of $-2\mathrm{J}$, which agrees with theory. In both the specific heat and magnetic susceptibility plots, a discontinuity is seen to occur around T_C, indicating that the phase transition is indeed of second order. Thus it can be concluded that the 2D square lattice of collinear spins is indeed a ferromagnetic system.

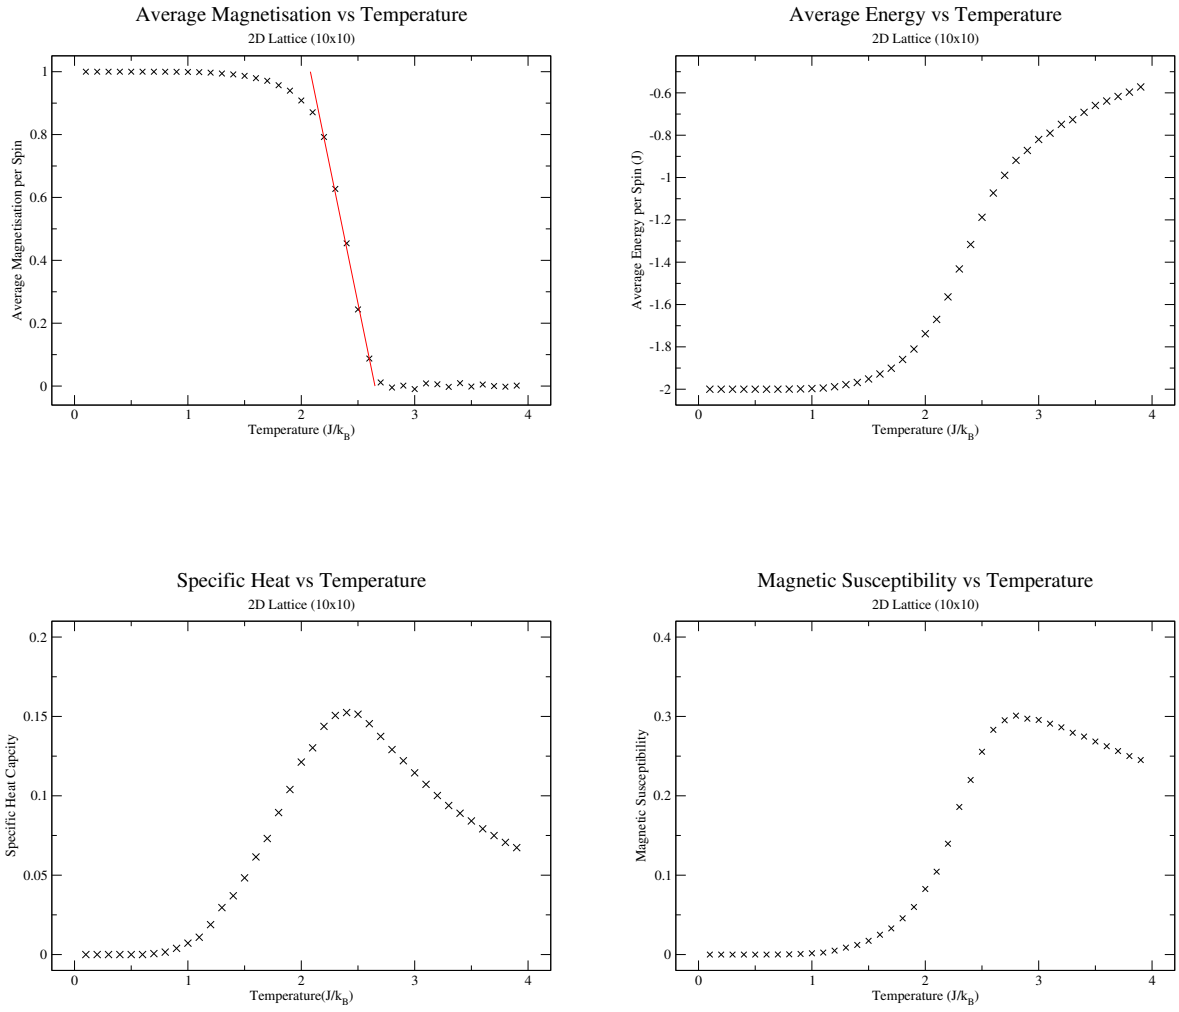

Figure 5: The measured observables for the 2D 10×10 square lattice. The Metropolis algorithm was implemented a sufficient amount of times to reach equilibrium, and the data was collected. This process was repeated 2000 times and average values were taken to reduce fluctuations in the plots. To save time, the program was run simultaneously on multiple machines and the process was repeated a smaller number of times, then the data was collected and averaged. The data was then formatted. The plots shown were formatted using Xmgrace.

The 2D triangular lattice system was then investigated. Since the only difference between square and triangular lattices is 2 extra nearest neighbours, it was easy to write the main class in such a way that the shape of the lattice could be specified when defining an object of the class. A plot of average energy per spin vs temperature for the triangular lattice is shown below, and the plots of the other observables vs temperature are included in the appendix. From the plot it can be seen that the ground state energy for the 2D triangular lattice is -3J, as expected. The Curie temperature was determined to be $4.2 \pm 0.1\mathrm{J/k_B}$, which is larger than the value of T_C found for the 2D square lattice. From the plots it is clear that the 2D triangular lattice is also a ferromagnetic system.

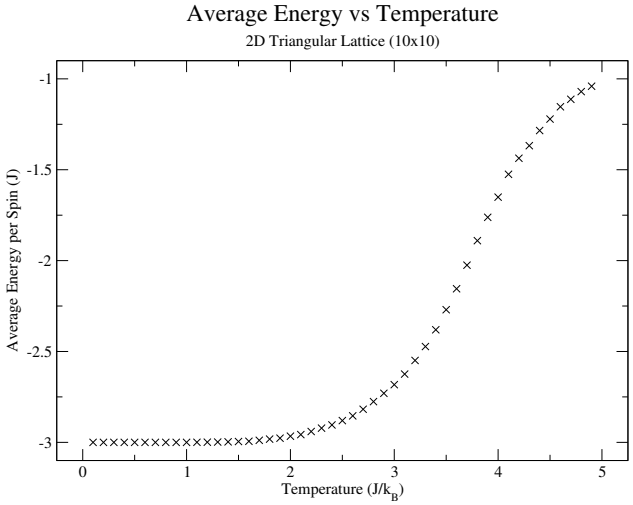

Figure 6: A plot of average energy per spin vs temperature for the 2D triangular lattice.

For the 1D lattice, a plot of average magnetisation per spin vs temperature appeared to exhibit a phase transition at a very low temperature, but this was in the case of $J = 1$. The plot was made at a larger value of J, and as expected, no phase transition was observed, agreeing with the theory that 1D systems with short range interactions cannot exhibit phase transitions.

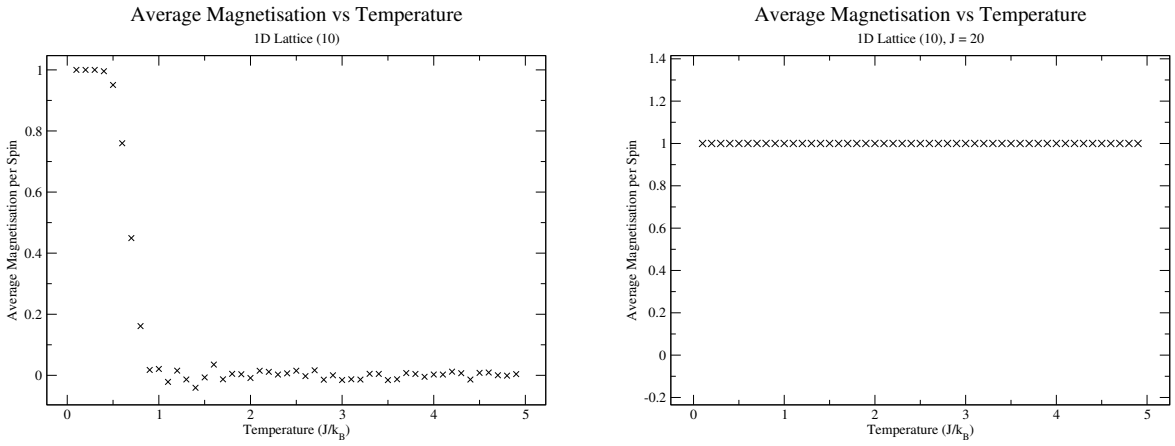

Figure 7: The plot on the left is of average magnetisation vs temperature for $J = 1$, and a phase transition appears to occur at a low temperature. The plot on the right is of average magnetisation vs temperature for $J = 20$ and the magnetisation remains constant, indicating that there is no phase transition, as expected.

"The model was suggested to Ising by his thesis adviser, Lenz. Ising solved the one-dimensional model, ... and on the basis of the fact that the one-dimensional model had no phase transition, he asserted there was no phase transition in any dimension. ... It is ironic that on the basis of an elementary and erroneous conclusion, Ising's name has become among the most commonly mentioned in the theoretical physics literature. But history had its revenge. Ising's name, which is correctly pronounced "E-zing", is almost universally mispronounced as "I-zing"!"
-Barry Simon [5]

The plots of the observables vs temperature for the 3D lattice are included in the appendix. The ground state energy as well as the value of $T_C = 4.6 \pm 0.1 \mathrm{J/k_B}$ are both larger than those of the 2D square lattice, which is expected, since the number of nearest neighbours is larger for the 3D lattice. From the plots it is clear that the 3D lattice is also a ferromagnetic system.

The average energy was plotted as a function of temperature for NiO, and from this the ground state was determined to be -36.75 ± 0.01meV. The plot of average magnetisation vs temperature shows that the average magnetisation per spin is approximately 0 for all temperatures, indicating that the system is not ferromagnetic. By examining the system at different temperatures, it is clear that it is disordered at higher temperatures, and the spins become misaligned for lower temperatures, indicating that the system transitions from a paramagnetic state to an ordered antiferromagnetic state.

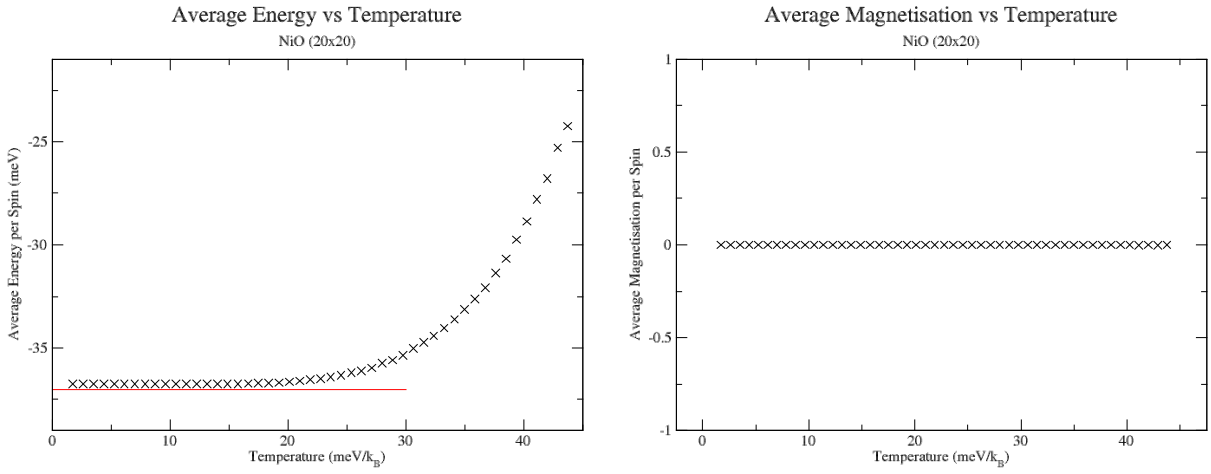

Figure 8: Plots of average energy (left) and average magnetisation (right) vs temperature for the 2D square lattice NiO model. The average energy appears to approach -36.75 ± 0.01meV as the temperature approaches 0, and the average magnetisation is approximately 0 for all temperature measurements.

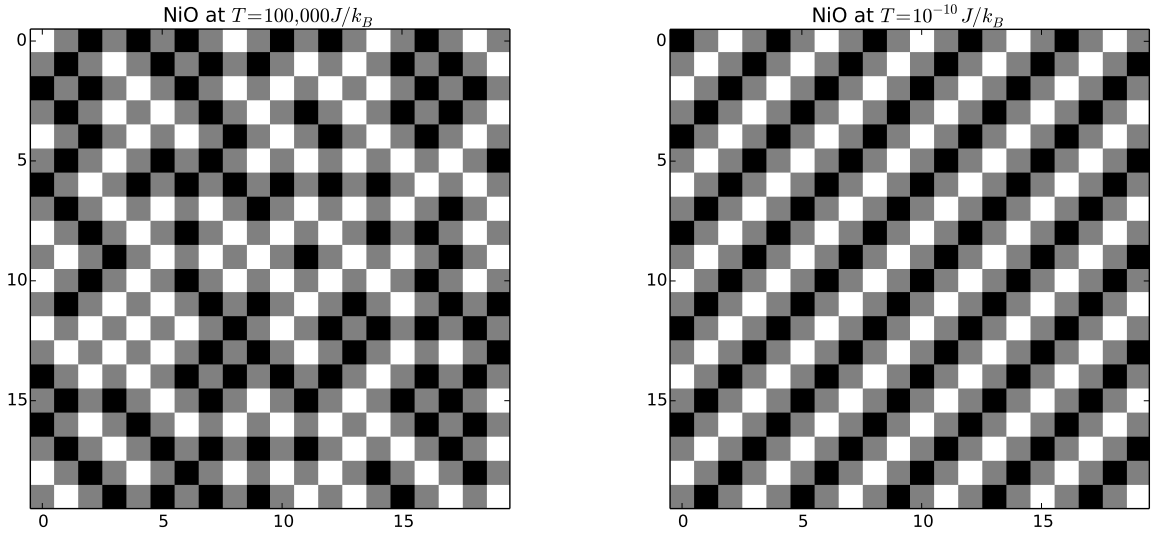

Figure 9: Plots of for high and low temperatures. The grey entries represent the oxygen atoms, and the black and white entries represent the spin "up" and "down" nickel atoms. At $T = 100,000\text{J/k}_\text{B}$ (left), the system is seen to be disordered (paramagnetic), and at At $T = 10^{-10}\text{J/k}_\text{B}$ (right), the spins are misaligned, indicating that the system is antiferromagnetic.

4 Discussion and Conclusions

The Metropolis algorithm was used to successfully update a system. The number of sweeps required to reach equilibrium for a given system size was then determined, and from this it was apparent that larger systems would be impractical to use when calculating observables due to the large number of sweeps required to bring them to equilibrium. Smaller systems were used for quicker computation times but at the cost of some numerical accuracy, although hopefully not too much due to periodic boundary conditions.

The 2D square lattice of collinear spins was determined to be a ferromagnetic system, as the plot of average magnetisation vs temperature showed a phase transition at $T_C = 2.6 \pm 1\text{J/k}_\text{B}$, which agrees well with the theory. The accuracy of the obtained value for T_C could be improved upon by using a larger system size, a smaller temperature step, or by averaging over a larger number of simulations, all of which would come at the cost of much longer computation times. The discontinuities in the plots of specific heat and magnetic susceptibility indicate that the phase transition is of second order. The ground state energy for this system was determined to be -2J. The 2D triangular lattice system was then investigated. The ground state energy was determined to be -3J, as expected, and the Curie temperature was determiend to be $4.2 \pm 0.1\text{J/k}_\text{B}$, which is larger than the value obtained for the 2D square lattice.

The 1D system appeared to exhibit a phase transition at a very low temperature, although this may have been due to the low value of J used during computation. A plot with a larger value ($J = 20$) shows that there is indeed no phase transition, which agrees with the theory. The 3D system was also determined to be ferromagnetic. The value of the Curie temperature for this system was determined to be $T_C = 4.6 \pm 0.1\text{J/k}_\text{B}$, which is larger than in the 2D cases. The ground state energy was determined to be -3J, as expected.

The results for the various systems are summarised in the table below.

System	Ground State (J)	T_C (J/k_B)
Square	-2	2.6 ± 0.1
Triangular	-3	4.2 ± 0.1
1D	-1	-
3D	-3	4.6 ± 0.1

A model of NiO was successfully developed and investigated. The ground state was determined to be -36.75 ± 0.01meV, which is close to the the expected value of -42meV. Once again, a larger system size, a smaller temperature step and a larger number of averaged simulations would have given a more accurate result. Unfortunately no useful information was obtained from the plot of average magnetisation vs temperature, other than the fact that it was approximately zero for all temperature measurements. Examinations of the system at different temperatures did show that the system undergoes a phase transition from a paramagnetic state to an antiferromagnetic state, as expected. However, the plots of specific heat and magnetic susceptibility unfortunately did not indicate what the Néel temperature was. Perhaps by using a larger temperature range and by determining the point of inflection on the average energy vs temperature plot, the Néel temperature could have been determined.

References

[1] Kerson Huang. Introduction to statistical physics. CRC Press, 2009.

[2] Hendrik A Kramers and Gregory H Wannier. Statistics of the two-dimensional ferromagnet. part i. Physical Review, 60(3):252, 1941.

[3] LD Landau and EM Lifshitz. Statistical physics, vol. 5. Course of theoretical physics, 30, 1980.

[4] Lars Onsager. Crystal statistics. i. a two-dimensional model with an order-disorder transition. Physical Review, 65(3-4):117, 1944.

[5] Barry Simon. The statistical mechanics of lattice gases, volume 1. Princeton University Press, 2014.

[6] L Van Hove. Sur l'intégrale de configuration pour les systèmes de particules à une dimension. Physica, 16(2):137–143, 1950.