Exercise sheet 5

Exercise 1: Monte-Carlo integration

1. Write a program to compute the volume of the unit ball (the set of vectors with norm ≤ 1) in d dimensions with a Monte-Carlo integration. Compare with the exact values $\frac{4\pi}{3}$, $\frac{8\pi^2}{15}$ and $\frac{\pi^5}{120}$ for $d = 3$, 5, 10.

2. *Importance sampling* in MC integration can be realized by a change of variables such that the new integrand is less strongly peaked. Suppose that the integrand is f, and that some function g has a similar shape and can be integrated analytically. Then one has

$$I = \int_a^b \mathrm{d}s \ \underbrace{f(s)}_{\text{peaked}} = \int_a^b \underbrace{\mathrm{d}s \, g(s)}_{\equiv \mathrm{d}\rho} \ \underbrace{\frac{f(s)}{g(s)}}_{\text{less peaked}} = \int_{\rho(b)}^{\rho(a)} \mathrm{d}\rho \, \frac{f(s(\rho))}{g(s(\rho))} \approx \frac{\rho(b) - \rho(a)}{N} \sum_{i=1}^N \frac{f(s(\rho_i))}{g(s(\rho_i))}$$

where the ρ_i are distributed uniformly on $[\rho(a), \rho(b)]$, and x and ρ are related by inverse transform sampling:

$$\rho = \int^{s(\rho)} g(s') \, \mathrm{d}s'.$$

We will apply this idea to the integrand

$$f(s) = \frac{\sin(\pi s/2)}{(s - m^2)^2 + m^2 \Gamma^2}$$

which is approximated by a *Lorentz curve* or *Breit-Wigner curve* or *Cauchy distribution*

$$g(s) = \frac{1}{(s - m^2)^2 + m^2 \Gamma^2}$$

with $\Gamma = 0.1$ and $m = 1$.

(a) Plot $f(s)$ and $g(s)$ on the interval $[0, 2]$.

(b) Calculate $I = \int_0^2 f(s) \, \mathrm{d}s$ using one of the deterministic methods of chapter 2. Aim for a high enough precision so that you can use the result as a reference value for what follows.

(c) Calculate I 30 times with the simple MC method of p. 187 of the lecture: 10 times each with $N = 10$, $N = 1000$ and $N = 100\,000$ points per integral. What do you obtain for the average numerical error for the different values of N?

(d) Redo this calculation using importance sampling, noting that

$$\int^s g(s') \, \mathrm{d}s' = \frac{1}{m\Gamma} \arctan\left(\frac{s - m^2}{m\Gamma}\right) + c$$

Exercise 2: Simulation of an ideal gas

Consider an ideal quantum gas in a box at temperature T. The gas consists of 1000 atoms, each of which is characterised by three quantum numbers n_x, n_y, $n_z = 1, 2, 3, 4 \ldots \infty$. The kinetic energy of a single atom is

$$E = \frac{\pi^2 \hbar^2}{2mL^2} \left(n_x^2 + n_y^2 + n_z^2\right)$$

where L is the box length and m the atom's mass. Since the atoms do not interact, their total energy is the sum of the kinetic energies. For now, we will use the approximation that the individual atoms are distinguishable.

1. Show that, if one of the atoms changes one of its quantum numbers by one unit (say $n_x \to n_x \pm 1$), then the energy changes by

$$\Delta E = \frac{\pi^2 \hbar^2}{2mL^2}(\pm 2n_x + 1)\,.$$

Use the Metropolis-Hastings algorithm to study this system at thermal equilibrium at temperature T. The probability that a given set of quantum numbers ϕ is realized is given by the Boltzmann distribution:

$$P(\phi) = \frac{1}{Z}e^{-\frac{E(\phi)}{k_B T}}\,.$$

The Markov chain is constructed as follows:

- At each step, randomly choose one of the atoms and one of its quantum numbers n_x, n_y, n_z.

- Suppose that n_x has been chosen for some atom. Propose a new configuration with the change $n_x \to n_x \pm 1$ (with the sign chosen randomly).

- If $n_x = 0$, reject this new configuration. Otherwise, compute the energy difference and accept or reject the new configuration according to the Metropolis-Hastings procedure.

It is convenient to use units where $\hbar = 1$, $L = 1$, $m = 1$, $k_B = 1$.

2. Plot the average energy per atom for $T = 10$ for a chain of $N = 200\,000$ points, starting with all atoms in the ground state, $n_x = n_y = n_z = 1$. How many points does it roughly take for the Markov chain to find the region of thermal equilibrium?

3. Plot the mean energy as a function of temperature for $T = 0, 2, 4, \ldots 20$ with samples of $N = 100\,000$ points and a number of burn-in steps which you determine from part 2.

4. *Bonus:* How would the method have to change for identical particles? Repeat exercises 2. et 3. for bosons. Then repeat them for a gas of 100 fermions for 10 values of T between 0 and 2000 with a suitable sample size.

Exercise 3: Simulated annealing

Write a program which determines the maximal number of dimers that can be placed on an $N \times N$ square tiling. Every dimer occupies two adjacent tiles. (The answer is of course $N^2/2$ or $(N^2 - 1)/2$ depending on the parity of N, but let's suppose we didn't know that.) Use the simulated annealing method.

Instructions: Start with an empty tiling. At each step, randomly select a tile and one of its neighbours. If both of them are empty, place a dimer there. If both are occupied by one single dimer, remove it with probability $e^{-1/T}$, where T is the "temperature". In all other cases change nothing. At each step, reduce the temperature by a factor $1 - \delta$. Experiment with several different values of δ and verify that the exact answer is approached as δ becomes small.
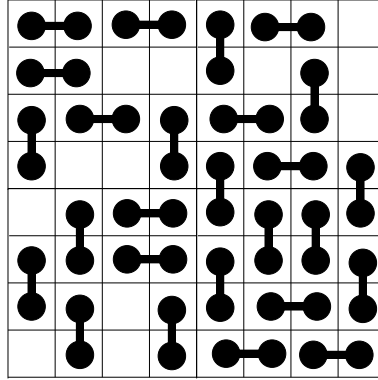
Example for dimers placed on a square tiling with $N = 8$.

Exercise 4: Kinetic Monte Carlo

Write a kinetic Monte-Carlo simulation of defect migration and annihilation in a simple cubic crystal lattice. Lattice sites are labelled by triples of integers representing their Cartesian coordinates. There are two types of defects, vacancies and interstitials, which can move to neighbouring sites at different rates. A vacancy will recombine instantly with an interstitial, removing both of them, if their euclidean distance becomes strictly smaller than two lattice spacings (or, equivalently, if all of their coordinates differ by at most one unit).

- Compute the migration rates $r_x = w_x \exp\left(-\frac{E_x}{k_B T}\right)$, where $w_i = 1$ fs^{-1}, $E_i/k_B T = 10$ (interstitials), $w_v = 10^{-3}$ fs^{-1} and $E_v/k_B T = 1$ (vacancies).

- Set up the system: Generate 100 coordinate triples for both interstitials and vacancies. The initial defect coordinates are normally distributed about the origin with standard deviations $\sigma_i = 20$ and $\sigma_v = 10$ respectively, rounded to the nearest integer. To sample from the Gaussian normal distribution with mean x0 and standard deviation sigma, use numpy.random.normal(x0, sigma). Check if any defect pairs can recombine immediately, and if they can, remove them from the system.

- For 100 000 simulation steps:

  - Calculate the total rate $R = N(r_i + r_v)$, where N is the number of defect pairs.
  - Select one of the defects at random, with a probability proportional to the respective migration rate.
  - Change one of its coordinates by ±1 unit at random, unless the new site is already occupied by a defect of the same type.
  - Check if the defect can recombine with any defect of the opposite type, and if it can, remove both of them from the system.
  - Draw a random number u from the uniform distribution on $[0, 1]$. Increase the time by $-\log(u)/R$.

- Plot the number of defect pairs as a function of time. Use a linear-logarithmic plot (matplotlib.pyplot.semilogx(x, y)). On a separate figure, produce a scatter plot of the defect coordinates projected onto the (x, y)-plane at the end of your simulation.