

PHYS 391 – Lab 4: Counting Statistics

Key Concepts

- Ionizing Radiation
- Counting Statistics
- Poisson Distribution
- Inverse Square Law

4.1 Introduction

This lab will explore the statistical properties of counting events in a random process. The specific process will be radioactive decays of nuclei observed by the presence of ionizing radiation in a Geiger counter. The basic concepts explored in this lab, however, are equally applicable to counting events in any random process.

Due to restricted access to campus, you will NOT be able to acquire the Geiger counter data yourselves. Instead, you will be provided with videos of the setup for each section of the lab and the corresponding data. The setup instructions are still provided below so that you know what is involved in the setup for each section. You will still need to think critically about your use of the data files provided to accomplish the physics goals in each section of the lab!

You will hand in your own ipython notebook with markdown boxes for discussion prompted in the following sections.

4.2 Note on Radioactive Sources

Ionizing radiation is produced naturally by many sources, and many common activities (like flying in an airplane or getting a single dental X-ray) will increase your exposure to radiation far more than the weak sources we will be using in the lab. Exercising some common sense, however, is in order. As we will see in this lab, the flux of radiation from a point source drops as $1/r^2$, also known as the inverse square law. It is always a good idea to reduce your exposure to any potentially dangerous situation, and with radioactive sources the best thing you can do is minimize the amount of time the source is very close to your body. In particular, you should not swallow the sources, and washing your hands after lab is always a good idea.

Because the University carefully controls any radiological material, the sources will need to be returned to the TAs at the end of each lab session. These will not be left out in room 17 over the weekend. As you will see, however, most of the radioactive items we are using are (or were) widely available household items.

4.3 Theory

4.3.1 Poisson Distribution

Any random process with two outcomes is described exactly by the Binomial distribution

$$P_{n,p}(\nu) = \frac{n!}{\nu!(n-\nu)!} p^\nu (1-p)^{n-\nu}$$

where n is the total number of trials, p is the probability of outcome A , and ν is the number of trials which resulted in outcome A .

While the Binomial distribution describes the probability of a given number of outcomes exactly, it is rather cumbersome to use in practice for any appreciably large number n . In the limit of n becoming large, but the product np being finite, the Binomial distribution can be re-written as the Poisson distribution

$$P_\mu(\nu) = \frac{\mu^\nu}{\nu!} e^{-\mu}$$

where $\mu = np$ is a real number and ν again is the integer number of times a certain outcome is observed. For the case of counting random events, outcome A can be thought of as the outcome where the event occurs,

while outcome B is where the event does not occur. The mean of the Poisson distribution is simply given by $\bar{\nu} = \mu$ meaning that on average you expect to see μ events in any given time interval, and the Poisson distribution gives the probability of observing an actual number of events ν . Figure 1 shows examples of the Poisson distribution for a range of values of μ .




Figure 1: Examples of the poisson distribution for values of μ (labeled as λ on the figure).

One key feature of the Poisson distribution, which is completely described by the single parameter μ , is the fact that the standard deviation of ν is given by $\sigma_\nu = \sqrt{\mu}$. If we then observe some number of events N , or best estimate for the average rate μ is given by N , and the uncertainty on this estimate is given by the expected standard deviation which would be $\sqrt{\mu}$, but is best estimated by \sqrt{N} . The relative uncertainty on any rate measurement, then, is given by $\delta R/R = \delta N/N = 1/\sqrt{N}$. In other words, the relative uncertainty on determining a rate improves with the number of events collected as $1/\sqrt{N}$.

4.3.2 Gaussian Approximation

As soon as the value of μ becomes appreciably large such that $\sigma_\nu/\bar{\nu} = 1/\sqrt{\mu} \lesssim 3$, the difference between a Poisson distribution and a Gaussian distribution with mean μ and width $\sigma = \sqrt{\mu}$ becomes negligible. This can be seen in Figure 1 where already for $\mu = 10$ the probability distribution looks very Gaussian.

4.4 Detecting Ionizing Radiation

4.4.1 Types of Ionizing Radiation

Ionizing radiation is defined as any process which can ionize (knock electrons out of) an atom. Visible light is an example of a type of radiation which is explicitly *not* ionizing radiation. Ionizing radiation generally comes in many forms depending on the particles involved, but the three most common are known by their historical designations as alpha, beta, and gamma radiation.

Alpha radiation, or alpha particles, are fully ionized helium nuclei (two protons and two neutrons) which are emitted in certain nuclear fission reactions like the decay of ^{238}U . These are relatively heavy particles which can cause tremendous amounts of biological damage, but fortunately can also be shielded quite effectively by even a small amount of material. A sheet of paper will readily stop alpha particles. The most common household use of alpha emitters is in smoke detectors, where the attenuation of Alpha radiation from ^{241}Am is used to detect the presence of trace amounts of particulate matter (smoke) in the air. Old Coleman lantern mantles used thorium dioxide which contains ^{232}Th , a naturally occurring alpha emitter. Alpha decays are generally only produced by very heavy nuclei.

Beta particles are electrons emitted in the weak decay of a neutron in the nucleus by the process $n \rightarrow pe^- \bar{\nu}_e$. In some cases, the inverse process $p \rightarrow ne^+ \bar{\nu}_e$ is also possible, and the positron emitted in this

process is also known as a beta particle. Beta decay is common in radioactive isotopes from lower down on the periodic table, such as ^{90}Sr . Beta particles are more penetrating through material than alphas, although a small amount of metal is usually enough to attenuate most commonly available beta emitters.

Gamma particles (often called gamma rays) are high energy photons which are often produced in association with either alphas or betas in a nuclear decay chain. For example, ^{60}Co emits both an electron and two photons as it decays to ^{60}Ni . Historically there has been a distinction between X-rays and gamma rays according to the energy of the photon, although the line between the two isn't very clear and there is no good physics reason to think of them differently. In fact, all the way down in energy to Ultraviolet light, photons are ionizing. The attenuation of photons by material is highly energy dependent, however, and gamma rays from nuclear decays typically will pass through centimeters of steel before interacting once. Dense materials provide the best shielding for gamma rays, and thick plates or bricks of lead are often used.

4.4.2 Geiger Counters

There are many techniques to detect ionizing radiation, but they all rely on observing the interaction of the ionizing radiation with some active detector material. In a Geiger counter the active material is an inert gas such as helium or neon with the addition of some halogens. This gas is contained in a volume with a high electric field. One typical geometry is shown in Figure 2 where a positive voltage of several hundred volts is applied between a wire (anode) in the center of a cylindrical tube (cathode) on the outside.

When ionizing radiation enters a Geiger tube, it knocks electrons out of the gas which are then accelerated in the electric field towards the anode. These electrons will collide with other atoms in the gas kicking out more electrons, and the end result is an avalanche of electrons falling onto the anode wire. Geiger counters are operated at a high enough voltage that even a single electron ionized from the gas will produce a large and easy to sense avalanche of charge falling on the central wire. The resulting "pulse" is passed through the coupling capacitor shown in Figure 2 and used to indicate the passage of a single ionizing particle.




Figure 2: Simple cylindrical geometry of a Geiger tube.

Unlike other types of detectors, the Geiger counter gives no information about the energy of the incident particle, but simply counts the number of times ionizing radiation passes through the Geiger tube. It has very high efficiency for detecting electrons, somewhat less for photons (due to their less frequent interactions) and can be effective for counting alphas only if there is a thin window to allow the alphas to enter into the gas volume.

4.5 Poisson Statistics

This part of the lab can be done without a source of any kind. We are going to look at the background count rate observed by the Geiger counters and compare the events seen in some time interval to a Poisson distribution. On one of the computers in Room 17, start the application LoggerPro. Get a Geiger counter from the TA and plug it into the LabPro green box using the special "Digital Radiation Monitor Cable". Placing the Geiger counter on a book or wooden block so that it is several centimeters off the table will likely improve your results in this lab.

To set up LoggerPro to see these monitors, click the small green icon in the upper left corner of the LoggerPro window (or select the menu item Experiment:Set Up Interface:LabPro). This will bring up a graphical depiction of the LabPro box. On the channel where the Geiger counter is attached, select “Radiation” (not “Student Radiation Monitor”) and close the box.

You now need to set up LoggerPro to count the number of events in some reasonable time interval. Click on the watch icon to bring up the DataCollection dialog box (or choose the menu item Experiment:Data Collection). Choose the mode to be Time Based, set the collection time to be some reasonably long time (like 150 seconds) and set the sampling rate to be something like 5 seconds/sample. LoggerPro will then count the number of events seen by the Geiger counter in a 5 second interval, and return this event count as a single value. Make sure the “Sample at time Zero” and “Oversample” options are not selected, and close.

Now, if you start the acquisition by pushing the big green button at the top of LoggerPro you should (after 5 seconds) start to see counts per 5 second interval displayed. The default time history isn’t so useful, so you should add a histogram graph to see a nice Poisson distribution taking shape.

Ideally we want a mean number of counts per time interval of order 1. If your mean appears to be significantly above or below 1, adjust the sampling rate such that you get on average about one count per time interval. What you are observing is some combination of natural background radiation, cosmic rays, and spontaneous breakdowns in the Geiger tube gas. Occasionally you will see an event where many counts happen at once which would seem to be statistically very unlikely. Not all of these background processes are completely random, for example cosmic rays often result in several particles at once passing through a small area.

Once you have a table of counts per time interval, copy this data to python and find the mean and standard deviation for this data (you can also do this on the spot in LoggerPro). How well does the relation $\sigma = \sqrt{\mu}$ hold? Try making a histogram of your data in python. Note that the `hist` function isn’t very good at dealing with integers. You will need to specify the exact binning, and to avoid ‘edge’ problems, it is best to specify the bin centers (not the edges) at integer values using the numpy `arange` function. For example, something like `hist(dat, np.arange(-0.5,5.5,1.0))` should work to make a histogram of integers in 6 bins from 0 to 5.

Using the measured average count rate μ , make a figure with a Poisson distribution with that value of μ overlaid on top of a histogram of your data. Remember, a Poisson distribution is normalized so you should multiply the probabilities by the total number of time intervals observed to make the vertical scale work out properly. I have shown some examples of similar plots in class, or just check out some of the `matplotlib` examples. You should include this figure in your report. The `fig.savefig('myflie.png')` function can be used to save a figure to a variety of file formats. How would you (qualitatively) assess the agreement between the Poisson distribution and your data?

What is your measured background rate in counts per second with uncertainty? Note that you only need the total number of events seen in the entire data collection time (and the total time) to determine the average rate. Does this agree with your Poisson mean?

4.6 Gaussian Limit

Now, take one of the sources and put it close enough to the Geiger counter to get an appreciable count rate. If you are using one of the Fiestaware salt shakers, 10 cm should be about right. Measure the distance from the counter to the source, as we will use this information later. Repeat your measurement of counts per time interval. You should see an average count rate of over 10 counts per time interval. Anything over 5 is probably OK, but if it is below that, try moving the source closer.

Again, import the data into python (or use the built-in analysis tools in LoggerPro) and find the mean and standard deviation of the number of counts per time interval. Again, does the relation that $\sigma = \sqrt{\mu}$ appear to be valid? Save your histogram (either from python or LoggerPro) so that you can include this in your report and comment on how close to a Gaussian this appears to be.

4.7 Inverse Square Law

Here we will quantitatively verify the inverse square law. For any point source, the flux of particles through a fixed area will fall as $1/r^2$ where r is the distance from the source to the fixed area. In our case, the source

isn't really a point source which may complicate things. Also, there may be some attenuation of the particles in the air (this could in fact be highly significant for alpha particles) so we may expect to see deviations from a perfect inverse square law.

Using the results from the previous section as a first data point, take at least six more measurements, moving the source further away from the detector each time, and calculate the average event rate (with uncertainties) for each distance. Make sure at least 4 measurements are taken at distances greater than 20 cm from the detector. You can either calculate the rate per second, or just count a total number of events and divide by the elapsed time. Don't forget to subtract the background if relevant. Plan your distances carefully so that you actually see a reasonable number of counts at each position you measure. Also note that you may need to count for a longer time as the rate goes down to get a reasonable uncertainty on your rate. Be sure and explain in your notes what distance you are quoting so that you can correctly include this in your report: a figure is probably helpful here. Moving the source closer than about 10 cm will likely cause trouble due to the finite size of the source and detector.

The rate vs. distance data, for a point source, should follow a functional form given by $R(r) = C_0/r^2$ where C_0 is some constant and r is the distance. To test this, we would like to fit to a function like $R = C_0/r^n$ and experimentally measure the parameter n and verify whether this is consistent with 2. It is probably easiest to "linearize" the problem so we can fit a straight line to the data. There are several ways this could be done, but taking the log of the equation above is probably the most straightforward. Write down explicitly what two quantities you need to plot against each other, and how the parameters of a linear fit correspond to the parameters C_0 and n .

Enter your data into python (by hand is fine) and perform a linear fit to this data using one of the methods we have discussed before. If you are really feeling ambitious, you can perform a weighted fit by specifying a vector of weights for each point, but this isn't necessary. If you use the errors in the fit, be extremely careful to specify the errors (or weights) correctly!

Make a plot of rate vs. distance including error bars (the `errorbar` command is useful here) and overlay this data with the results of the fit. Include this plot in your report. Be sure to label the axes and specify appropriate units. You can either plot rate vs. distance directly (which should not be linear) or you can plot the actual quantities which you performed the linear fit on. Alternately, `fig.set_yscale("log")` function can be used to directly change the y-axis from linear to log scaling.

What is your measured value of n (including uncertainty), and how does this compare with what you expect? If this is not consistent with the expectation, is there is any trend in your data which might explain the discrepancy? For instance, is there evidence of attenuation of particles in the air, or is there just one point which seems to be pulling the fit away from the expected value? A residual plot (data value - fit value) might be useful if the trend isn't clear from your plot above. Would the non-point nature of the sources produce a coefficient greater, or less than $n = 2$, and why? The flux of particles from a source follows exactly the same logic as the electric field strength from a charge distribution under Gauss' Law. Discuss your results and include any plots which can support your arguments.

4.8 Attenuation in Matter

Finally we want to verify the attenuation of ionizing radiation in matter, and measure the attenuation length. Place the source at a distance where you get a good counting rate (this can be quite close), and take measurements with varying amounts of material placed between the source and the detector without changing the distance. Sheets of aluminum foil work well, although notebook paper will also work. Increase the thickness of the material in a geometric progression (1, 2, 4, 8, ...) until you see at least a factor of 4 attenuation in the observed rate. To save time, it is best to figure out how much material you need to reduce the rate by at least a factor of 4 first, then work backwards in reasonably spaced steps so that you can cover the full range of rates with somewhere between 5-10 data points. Again, you probably want to count for a longer time when the rate is low so that the uncertainty on the rate is roughly constant for each point.

We expect the observed rate to vary as $R(l) = R_0 e^{-l/\lambda}$ where again R_0 is the unattenuated rate, l is the thickness of material attenuating the particles, and λ is the attenuation length. Again, figure out how to linearize the problem and perform a linear fit to extract the attenuation length λ . It is fine to do this fit in units of sheets of material (in other words l is just an integer number with some funny units). Include plots and a final measurement of λ in your log book. To really get a proper uncertainty on λ , you should use the

individual rate errors on each point in the fit. Be very careful about your error propagation, however, since you will need the absolute uncertainty on $\log R$, since that should be the quantity you are fitting to in order to get a linear dependence.

To try to make this result a little more meaningful, we need to convert λ into suitable units. First, estimate somehow the thickness of your absorber material and convert λ into units of centimeters. If you used aluminum foil, you should just measure the thickness of a stack of sheets using a set of calipers. Second, estimate the density of your material somehow (you can just look it up if you wish) and calculate the product of the attenuation length times the density $\lambda\rho$ which has units of g/cm^2 . This product is the typical unit used for quoting attenuation length (or sometimes the inverse, known as the mass attenuation coefficient μ/ρ where $\mu = 1/\lambda$). Write down your best value for $\lambda\rho$ including an uncertainty based on your linear fit above. We have discussed various ways to find an uncertainty on a fitted slope in this class. Pick one of them.

Electrons (beta particles) in aluminum tend to have a mass attenuation length of around $0.1 \text{ g}/\text{cm}^2$ although it depends a bit on the exact electron energy. For alpha particles the attenuation length is many orders of magnitude shorter, while for photons it typically is much longer (again depending on energy, but around $10 \text{ g}/\text{cm}^2$ for 1 MeV photons). If you used paper, the attenuation length won't be significantly different when written in these density-corrected units. From your observations, does it seem more likely that we are detecting alpha, beta or gamma particles? Justify your answer.

4.9 Final Thoughts

Most radioactive processes don't just produce one type of ionizing radiation, although usually one type is more apparent than others. Often, alpha or beta particles can be shielded quite readily with a small amount of material, but then there is a small component of gamma radiation which doesn't go away so quickly. On the other hand, just because a particular type of radiation is produced doesn't mean a given detector is sensitive to it. Understanding the response of your detectors to different radiation types is critical for making accurate measurements of ionizing radiation. So don't be surprised if your results in Section 4.8 don't look completely like you might expect. Historically, a considerable amount was learned about the nature of radiation in the early days with detectors even more rudimentary than the ones you are using (which were invented around 1928).