Lab 1: Photon Statistics.

Andy Friedman [1,2,3]
September 11, 2001

ABSTRACT

In this lab experiment, we study photon statistics using a Photo-multiplier Tube (PMT) and a computer interface. Photons are emitted from the PMT at a theoretically constant rate which we attempt to measure. We use our own software to vary the rate at which photons are sampled and to repeatedly count the number of photons in each sample. In this report, we discuss several small experiments which vary these parameters and we use the results to illustrate important statistical concepts. These include the mean, standard deviation, variance, Poisson and Gaussian Distributions, and the standard deviation of the mean. Basic introductions to the UNIX operating system, idl programming, and LaTeX document formatting were necessary to complete this lab report.

1. Introduction: Why We Care About Counting Photons

If this was a scientific paper or even a typical lab report, one would begin by discussing some of the basic physics of the Photo-multiplier Tube (PMT), the measuring apparatus which counts photons and displays the results via a computer interface, and if relevant, the potential applications to astronomy. But in this report, one may regard the PMT apparatus simply as a black box that emits photons at a rate that in theory is constant, but in practice fluctuates around the ideal value due to unspecified quantum mechanical processes which are inherently random.[4] The photons themselves can also be considered simply as discrete objects that can be counted. Their physics can for all practical purposes be ignored. Knowing that the details of the experimental apparatus are not the major point of this lab, this report therefore serves more as an introduction to the generally applicable tools of statistical analysis, programming, and scientific writing, which are relevant skills regardless of whether or not we continue on in astronomy. Photon counting simply provides a simple and interesting context within which to learn these skills.

---

[1] e-mail: friedman@ugastro.berkeley.edu

[2] Lab Group 3: Andy, John, Shane, Henry, Cheryl

[3] In addition to making the formatting of this report possible, LaTeXhas a really cool logo

[4] The count rate will be a different constant depending on the setting of the dial on the PMT apparatus. High count rates, for example, are illustrated via an audio signal composed of rapid discrete clicks, much like a Geiger counter.

2. Experiments:
Procedure, Results, and Data Analysis

In this section, we discuss the five small photon counting experiments we performed. For convenience, they will be titled "Histogram", "Mean and Standard Deviation", "Poisson Distribution", "Gaussian Approximation", and "Standard Deviation of the Mean". The experimental procedure, results, and data analysis for each of the experiments will be discussed individually, stressing the comparison between actual data and theory. Relevant idl plots with handwritten captions are referenced and can be found in the appendix section of this lab report. IDL programming specifics will not be discussed, but the source code for all relevant programs and generated data files can be found in my home directory in the folder /lab1 and its subdirectories SameRate, VariedRate, and SDOM.

2.1. Histogram

In this first experiment, we began to take data from the PMT by first adjusting the dial on the apparatus, listening for clicks to give a rough estimate of the count rate, and finally writing idl code to request 6 runs with sample sizes of 100 and 1000 respectively, each with count rates of 1000.[5] Each run gets sent to a data file with 100 or 1000 entries which measured the number of photons detected during each 1ms time period.

Once we have acquired data from the PMT, it becomes useful to display it in an easily interpreted fashion. One such method is to display the number of photons detected as a function of time. In such a plot, the data looks much like that from a polygraph test or an electrocardiogram (EKG) readout. The useful feature of such plots is that by eye alone, one can get a useful estimate of both the mean and the standard deviation from the sample. The mean here is simply the average number of counts per sample, while the standard deviation here is a measure of how far on average the counts deviate from the mean number of counts.

Mathematically, if we have a sample of N measurements: $x_i ... x_N$, the mean $\bar{x}$ is defined as:

$$\bar{x} = \frac{\Sigma x_i}{N} \tag{1}$$

Using this definition of the mean $\bar{x}$, the standard deviation σ_x is defined as:

$$\sigma_x = \sqrt{\frac{1}{N-1}\Sigma(x_i - \bar{x})^2} \tag{2}$$

These definitions will be implicitly referred to throughout the report whenever the words mean or standard deviation are mentioned.

Returning now to our actual data, in pages 1 and 2 of the appendix, plots of the Number of Counts vs. time for 6 independent runs can be found for sample sizes of 100 and 1000 respectively.

[5] A count rate of 1000 means that for each sample, photons are counted during a 1/1000 s or 1ms time period.

Both plots have count rates of 1000. From these plots, the mean is easily seen as the zero slope straight line that is your best estimate of the data, while the standard deviation is, roughly speaking, the amount of "jaggedness" in the plot. These qualitative measurements are nonetheless useful and often can give one good approximations for the measured mean and standard deviation, which are shown on both plots along with the relevant number of samples. Note that the plots with a sample rate of 1000 are have a lot more data and therefore have much darker lines. Notice also that the standard deviations for corresponding plots from the sample size 100 and sample size 1000 runs are not appreciably different. This tells us that the standard deviation is relatively independent of sample size, which makes sense because although taking extra data is helpful toward zeroing in on the true mean, it is not likely to cause points to deviate any more from the mean than they would do so for a smaller run. By analogy, increasing the size of a class is unlikely to change what the standard deviation is for the midterm. This assumes, of course, that the quality of the new students is comparable to the old students, since one could theoretically overload the class with geniuses and cause the mean to skyrocket.

Another useful way of displaying the data is to plot it in the form of a histogram. A histogram plots individual numbers of counts that were measured (also called bins) on the x-axis, while plotting the corresponding number of occurrences in each bin on the y-axis. For example, if there were 30 separate samples in which 300 counts were measured, the histogram would show bin number 300 on the x-axis with a corresponding value of 30 occurrences on the y-axis. In a graph which plots the number of counts vs. time, this type of data feature can be interpreted as the fact that the line y=300 intersects the jagged data curve 30 times. Simply put, histograms plot how often a particular value of counts occurs.

On pages 3 and 4 of the appendix, histograms are plotted for the same data that was plotted previously in the form of counts vs. time. Here the mean can be seen approximately as the bin in which there are a peak number of occurrences, while one standard deviation can be seen as the half width of a rectangle centered about the mean that contains roughly 68% of the area under the histogram.[6]

Notice that no matter how the data was displayed, for the both runs of 6 with the same rate and the same sample size, *the mean count rates were not the same!*. The reason the means are different has to do with the fact that the emission of photons by the PMT is not a deterministic process. The physics of photon emission relies on inherently random quantum mechanical rules and can therefore only be interpreted statistically. This tells us that we can never expect to perfectly measure the true mean. Assuming that the mean does indeed exist, as we have good reason to do so for the count rate in this experiment, we can only arrive at the true mean in the limit that our sample size N goes to ∞.

This fact is illustrated well by the concepts of the supermean, (or mean of means (MOM)) and the standard deviation of those mean values. In this case, we have a set of 6 means and we get the MOM by taking the mean of those numbers. In general, if we have N means $\bar{x}_i ... \bar{x}_N$, the MOM $\bar{\bar{x}}$

[6] Taylor, pg. 101

is given by:

$$\bar{\bar{x}} = \frac{\Sigma \bar{x}_i}{N} \tag{3}$$

Similarly, the standard deviation of the set of 6 means is a useful quantity. I will omit the explicit formula for the general case. It can be derived by replacing the x_i in Eq. (2) with $\bar{x}_i$ and $\bar{x}$ with $\bar{\bar{x}}$. [7]

As we increase the sample size, we find that the MOM comes closer to approaching the true mean. We can tell this because the standard deviation of our set of mean values (stdevm) decreases as we increase our sample size. This trend can be most easily seen in the table below, which lists the means for each set of 6 runs for both sample sizes, along with their respective MOM's and stdevm's. Note that the means for each sample size do not correlate perfectly. But what we can extract from the table is that the true mean is probably closer to 298 than to 301, because we have more confidence in our run with 1000 samples. In fact, we can interpret our stdevm as the uncertainty in our measurements, and when we choose the experiment with sample size 1000, we arrive at a best value for our mean count rate of $\bar{x}_{best} = 298.587 \pm 0.47623$.[8]

	Sample Size 100	Sample Size 1000
Mean 1	300.000	298.631
Mean 2	303.700	299.194
Mean 3	297.800	297.870
Mean 4	301.560	298.587
Mean 5	302.470	298.973
Mean 6	302.220	298.270
MOM	301.292	298.587
stdevm	2.09679	0.47623

Table 1: *Means for 6 Runs of Sample Sizes of 100 and 1000. Notice that the MOM's are different and that the stdevm is clearly smaller for the 1000 sample size set.*

[7] This refers to the standard deviation of a set of mean values from several data sets and should not be confused with the theoretical standard deviation of the mean (SDOM) for a single data set. The SDOM itself represents the uncertainty for a given set of measurements. There is an explicit formula for the theoretical SDOM, on pg. 102 of Taylor. It will be discussed later in this report.

[8] The theoretical SDOM is usually interpreted as the uncertainty in measurement for a single data set, but here, since we have 6 data sets, we can get an even lower uncertainty using stdevm.

2.2. Mean and Standard Deviation

In the first experiment, we varied the sample size while keeping the sample rate constant. In this second experiment, we do exactly the opposite, keeping the sample size fixed at 100 while varying the sample rate from 5000 down to 370, in increments dictated by the fact that the counter only accepts rates that are integer multiples of its clock rate of 10kHz. With 17 runs each with a different sample rate, this experiment is designed to illustrate the relation between the number of counts and the standard deviation. The first thing we notice is that as the sample rate decreases, the mean number of counts per sample increases. Why is this? Well, as we decrease the sample rate, we increase the number of counts in each sample because, although the PMT dial is fixed and the device is theoretically emitting photons at a constant rate, we increase the time period during which the instrument is measuring counts. For example, at a high sample rate of 5000, the counting for each sample occurs over a tiny period of 1/5000 s, whereas with the extremely low sample rate of 370, the counting occurs over a much larger time period of 1/370 s, yielding more counts. This is easily seen in the plot of mean vs. variance on page 5 of the appendix, as the mean rises linearly as the sample rate decreases from 5000 to 370.

What is variance, you might ask? Variance is simply another way to characterize the relationship between the mean and the standard deviation, because in Poisson statistics, the variance is simply the standard deviation squared. ($variance = \sigma^2$) As mentioned above, a plot of mean vs. variance can be found on page 5 of the appendix. Using this plot, we find that the standard deviation increases as the mean number of counts per sample increases. This makes sense because although the standard deviation may indeed remain a fixed percentage of the mean as we decrease the sample rate, as the mean number of counts increases, so too will the standard deviation. By analogy, if I have a 10% stake of someone's gambling winnings, my stake will certainly increase if they happen to win more money, simply because 10% of 1000 is certainly greater than 10% of 10, for example. Whereas in the previous experiment, the standard deviation remained relatively unaffected as we increased the sample size from 100 to 1000, here it clearly increases as the mean number of counts increases.

In Poisson statistics, the true mean μ is derived to be equal to the variance σ^2.[9] To test how close this Poisson statistics prediction comes to our actual data, in our plot, we overplot the relation $mean = variance$ on top of our actual data, which plots real means vs. their corresponding standard deviations squared. What we find is that the actual data does not perfectly fit the theoretical relation $\mu = \sigma^2$. In the plot, the actual data points (triangles) seem to fit the Poisson Statistics prediction (the solid line) much better for higher sample rates and thus lower numbers of mean counts. This tells us that the Poisson statistics prediction that $mean = variance$ is valid for smaller mean counts and becomes a worse approximation as the mean number of counts increases. Of course, this claim would be further strengthened by plotting mean vs. variance for several sets of data. Although only one plot is included here, we did try it for several data sets and in all cases,

[9] Taylor, pg. 247

the fit to the Poisson prediction was better for lower mean counts per sample.

2.3. Poisson Distribution

In the previous experiment, we touched upon the idea of Poisson Statistics, which predicts that $\mu = \sigma^2$. But where does this prediction come from? It actually comes from a theoretical probability distribution that depends only on the single parameter μ, which is none other than the Poisson Distribution. This distribution is a valid approximation for the results of experiments whose individual results are random, but who globally possess a fixed average rate. The Poisson Distribution plots the probability of measuring a certain number of counts x as a function of x and μ. It is given by the formula:

$$P_\mu(x) = e^{-\mu}\frac{\mu^x}{x!} \tag{4}$$

As derived in Taylor, $\mu = \bar{x}$, which leads to the previously discussed prediction $\mu = \sigma^2$. [10] Knowing this, we performed our third experiment where we plotted a histogram for a data set with sample size 1000 and sample rate 5000 (to ensure that we'll get only a few counts per sample), and overplotted the theoretical Poisson distribution curve, taking into account the prediction that $\mu = \sigma^2$. As can be seen from the plot on pg. 6 of the appendix, the Poisson distribution fits the data remarkably well.

In order to ensure that we were plotting a number of counts bigger than 1 on the y-axis, we had to multiply the Poisson distribution by an appropriate scale factor, since, as a probability, $P_\mu(x)$ is initially some positive number less than 1. As it turns out, this scale factor is simply the sample size (1000 in this case). This makes sense because since $P_\mu(x)$ is the probability of measuring some number of counts, in order to get the actual number of counts measured, one need simply multiply by the total number of samples. For example, if the probability for of measuring $x = 8$ counts for some sample is $P_\mu(8) = 0.2$ (i.e. 20%), and we actually measure 1000 such samples, then we should expect to measure 8 counts a total of 200 times since 200 is 20% of 1000. And as mentioned, when we multiply by the sample size, the theoretical Poisson Distribution does indeed match quite well with our data.

2.4. Gaussian Approximation

A question one might ask is whether the Poisson Distribution is always a good description of the data. To answer this, we performed a fourth small experiment where, this time, we aimed to have a rather large number of counts per sample, as opposed to the previous experiment, where we deliberately kept the number of counts low. We did this again with 1000 samples, but this time we lowered the sample rate to 100, ensuring that the counting period would be relatively long, but

[10] Taylor, pg. 245-249

still keeping the count rate well below 1MHz, as to not be responsible for the destruction of the PMT apparatus. This experiment increases the number of samples, raises the mean from roughly 7.4 counts to 369.6 counts, and causes the histogram to become more symmetric about the new mean. In fact, it changes the shape of the distribution enough so that the Poisson Distribution may no longer be the best approximation.

As it turns out, in the limit of large counts per sample, the Poisson Distribution is approximately equal to the Gaussian or Normal Distribution (also know more familiarly as the Bell Curve). This is remarkable because the distributions are so different. The Gaussian Distribution $G_{\mu,\sigma}(x)$ gives the probability of a continuous variable, depends on 2 parameters (μ and σ), and is always symmetric about its mean, whereas the Poisson Distribution $P_{\mu}(x)$ gives the probability of a discrete variable, depends on only one parameter μ, and in general, is not symmetric about its mean. So how can the curves ever be alike? Mathematically, the Gaussian Distribution is given by:

$$G_{\mu,\sigma}(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{\frac{-(x-\mu)^2}{2}} \tag{5}$$

And as the number of counts μ increases, this curve becomes a better and better approximation for the data. In fact, as μ approaches ∞, the Poisson Distribution becomes progressively more Bell shaped and approaches the Gaussian Distribution with the same mean and standard deviation.[11] When we plotted a histogram of the data for this experiment and overplotted a Gaussian Distribution, as seen on page 6 of the appendix, $G_{\mu,\sigma}(x)$ does in fact fit the data quite well. Evidently $\mu = 369.602$ is large enough for the Gaussian approximation to be valid.

2.5. Standard Deviation of the Mean

The final experiment we performed was designed to show us that we can measure the count rate more accurately with larger numbers of samples. For example, with ten sets of data each with a given sample size, due to statistical variations, the means of each of those ten sets will not be identical. One way then to characterize their variation is to take the supermean or mean of means (MOM) of those ten means. Additionally, one could take the standard deviation of those means (SDOM) to find out how much the means vary on average from the MOM. If larger sample sizes actually do produce better measurements, then ten sets of data with larger sample sizes should have a smaller SDOM's. In this experiment, we tested this hypothesis by varying the sample sizes from 2 up to 2048, doubling each successive number, and performing ten runs for each sample size. For each of these sets of ten, we then calculated the MOM and the SDOM. The results can be found on page 7 of the appendix as x-axis logarithmic plots of both the supermean and the SDOM as functions of sample size. The results are also summarized in the table at the top of page 8. What we find is that indeed, the SDOM decreases as the sample size increases, confirming our hypothesis. In addition, the MOM seems to be asymptotically approaching the true value of the mean, which

[11] Taylor, pg. 250-251

it should reach if the number of samples were to go to ∞

Sample Size	MOM (Actual)	SDOM (Actual)	SDOM (Theory)
2	74.1000	5.46606	6.08687
4	73.7750	5.63527	4.29462
8	74.0250	2.41509	3.04189
16	73.8625	1.57476	2.14858
32	74.7250	1.80847	1.52812
64	73.8156	1.01755	1.07395
128	73.3703	0.800511	0.757103
256	73.2602	0.560223	0.534951
512	72.6566	0.524985	0.376706
1024	72.0181	0.313905	0.265198
2048	72.0139	0.162648	0.187518

Table 2: *Lists of the MOM's, Actual and Theoretical SDOM's as a function of sample size. The MOM's seems to be asymptotically approaching a value of close to 72. Also notice that the SDOM clearly decreases as the sample size increases and that the Actual and Theoretical SDOM's correlate quite well.*

In addition, our knowledge of Poisson Statistics can be brought to bear to predict what the SDOM as a function of the measured mean counts per sample and the sample size N. In general, The equation for the standard deviation of the mean (SDOM) for a single data set of measurements $x_i ... x_N$ is given by:

$$\sigma_{\bar{x}} = \frac{\sigma_x}{\sqrt{N}} \tag{6}$$

In this case, we can use our knowledge of Poisson Statistics to substitute in the expression $\sigma_x = \sqrt{\mu}$, and in this case, our best value for μ, the average number of counts, is none other than the measured MOM. This then gives us values for the theoretical SDOM's which we can compare to the actual SDOM's. The actual SDOM's simply come from calculating the standard deviation of the measured means using Eq. (2). In the plot on page 7 of the appendix, and in the table above, we do indeed see a strong correlation between the actual SDOM's (triangles) and the theoretical SDOM curve extracted from Poisson statistics (solid line). In any case, both theory and experiment confirm our assumption that measurements of an average quantity increase in accuracy the more samples we take. And since the SDOM is the theoretical value for the uncertainty in our measurement, our best estimate of the average count rate for this data set is given by the run with 2048 samples, and can be written as $\bar{x}_{best} = 72.0139 \pm 0.162648$.

The fact that measurements can be improved by increasing the number of samples naturally leads to the quantitative question of how many samples one need in order to measure some quantity to with some arbitrary improvement in accuracy. For example, if we wanted to improve the accuracy of our measurement by a factor of two, by what factor would we need to increase the number of samples? Let's say we have N samples, then our uncertainty is given by the theoretical SDOM, $\sigma_{\bar{x}} = \frac{\sigma_x}{\sqrt{N}}$. So if we wanted $\sigma_{\bar{x}}$ to decrease by a factor of 2, to make our uncertainty twice as small, we would need the N under the square root in the denominator to go to $4N$, thereby allowing the denominator to change to $\sqrt{4N} = 2\sqrt{N}$. Thus to improve the accuracy of your measurement by a factor of 2, one must increase the sample size by a factor of 4. This tells us that although there are definite improvements that come from increasing the sample size, at some point, the improvements will not outweigh the experimental costs of continually increasing the sample size by a factor of 4 for every doubling in accuracy.

3. Systematic Errors

In this lab, a great deal of time was spent identifying and characterizing the various types of random errors that the experiments were subject to. These errors are not so terrible because they can, in principle, be dealt with in a statistical manner. On the other hand, systematic errors, which can not be handled statistically, have been largely ignored, so it is at least worth mentioning a few possible types of systematic error one might look out for in a photon counting experiment. First of all, systematic errors are errors that affect all of the data points in the same way, such as a ruler being miscalibrated. In this experiment, one such systematic error could be that there are always 10 extra counts added to every sample. This could be a bug in the software, or a glitch in the detector itself.

In any case, this effect would be hard to notice with high mean count rates of say a few hundred. But one way to test for the presence of such a systematic error[12] is to turn the PMT count rate way down, or even all the way down to zero! and see what the detector reads. If, in fact, the detector consistently lists on average 10 counts when we know the PMT has been turned off, then we can be reasonably sure that we have found our source of systematic error. As it happens, James performed this test and fortunately, only an occasional dark count of at most 1 or 2 showed up in each sample. This tells us that the effect of thermal dark counts is negligible in comparison to the actual counts and justifies our neglecting them as a source of systematic error, provided that the count rate yields, say, at least 5-10 counts per sample. Systematic errors will definitely be of more concern in future labs, but for this lab, we felt the topic deserved at least a token mentioning.

4. Conclusion

In this lab, we explored a variety of statistical concepts, and via the exciting context of a photon counting experiment, came up with some interesting general results that apply to all random systems that have a theoretical average rate we may try to measure. We will be very concise and

[12] sometimes referred to as dark current because the detector reads counts even when the light is turned off!

present the results in words in a list format.

1. The Standard Deviation is largely independent of sample size.

2. All measurements of the mean will be uncertain and the true mean can only be measured in the limit that the number of samples goes to ∞.

3. The standard deviation increases as the mean number of counts increases. y

4. The Poisson statistics prediction of $mean = variance$ is most accurate for smaller mean count rates.

5. The Poisson Distribution is a probability distribution that when scaled up by a factor of the sample size, approximates a histogram of the counts per sample very accurately for low numbers of counts.

6. The Poisson Distribution is approximately equal to the Gaussian Distribution in the limit that the number of counts x becomes large.

7. The SDOM decreases as the number of samples increases, increasing the accuracy of the best fit measurement, $x_{best} = \bar{x} \pm SDOM$

8. Random errors can be treated statistically, while systematic errors usually involve some calibration error in the measurement apparatus itself.

The bottom line is that statistics give you the ability to deal with randomness quantitatively. The universe is filled with messy natural processes which are either practically or theoretically non-deterministic, introducing irreducible uncertainties into our experiments. Nevertheless, we can almost always still extract some valuable information from the chaos. Understanding what these errors mean and how we can seriously treat them to legitimize our measurements is a skill all serious scientists must master. Scientist or not, if one can view the world through the selective eye of statistical analysis, nature often seems to make a lot more sense. To wrap one's mind around an inherently random quantum mechanical process such as photon emission, statistical thinking is a must, and this lab has served well towards putting me further into that mindset.

5. Acknowledgments/References

ACKNOWLEDGMENTS

I would like to thank in particular, James Graham, Nate McRady, and fellow Astro 122 students/prisoners Lee, Kirsten, Christina, Eric, Lindsey, John, and my lab group John, Shane, Henry, and Cheryl for being there to answer random esoteric questions several orders of magnitude faster than any online, unix, or paper based manuals. Our lab group really didn't spend much time working together since our schedules were all quite different. I ended up working primarily with

whoever happened to be in lab. I suggest that lab groups for the second lab are rearranged based on compatible schedules, now that we're all officially UNIX gurus, through no real fault of our own.

REFERENCES

1. Taylor, John R., *An Introduction to Error Analysis.* University Science Books, Sausalito, CA 1997, 2nd ed.

2. Graham, James R., *ERRORS & STATISTICS: Ay 122 Lab Notes*, 2001

3. Unix, LaTeX, and IDL Tutorials http://ugastro.berkeley.edu/irlab/tutorials.html