# Bayesian analysis of lunar laser ranging data

William H. Jefferys
Judit Györgyey Ries

ABSTRACT
In 1969, astronauts first placed a retroreflector on the moon for laser ranging of the moon, and since then the McDonald Observatory of the University of Texas has been ranging to these and other later-placed retroreflectors. By determining the round-trip time of a very short but powerful laser pulse, important and extremely precise information about lunar motion and earth rotation can be obtained. The problem is an interesting one from the point of view of signal-to-noise, for in unfavorable circumstances, nearly all of the detected photons are not laser returns but simply background photons. Other interesting features of this problem are the fact that the data are censored; and that it is necessary to take into account the Poisson nature of the data. Determining which photons are actual returns is critical to the initial data analysis. In this paper we describe how a Bayesian analysis of the return data can be used to improve the results.

## 1 Introduction

The only experiment of the Apollo lunar missions still in progress is the lunar laser ranging (LLR) experiment. The arrays of reflecting cornercubes that the Apollo astronauts left on the Moon, along with two other arrays delivered by Soviet spacecraft, do not require power, and their surfaces have not shown measurable degradation since they were deployed. Improvements in laser technology and timing devices have increased the accuracy of the range measurements in the intervening quarter century, but lunar laser ranging remains a technically and scientifically challenging measurement. Lunar laser ranging provides a wide range of scientific results as well as a three orders of magnitude improvement in the lunar ephemeris and lunar rotation variations over earlier techniques. These include, for example, measurements of the Earth's precession, of the moon's tidal deceleration, of the relativistic precession of the lunar orbit, and a test of the Strong Equivalence Principle [DBF$^+$94].

Lunar laser ranging is the measurement of the round-trip travel time of a photon emitted from an Earth-based laser. Changes in travel time, which indicate changes in the separation between the transmitter and the reflector, contain a great deal of information about the Earth-Moon system, which can be retrieved by estimating model parameters. An important signal in the difference between the observed and predicted range is the error in the predicted Earth rotation parameters. The motivation for better identification of photons returning from the moon is to better determine these Earth Orientation Parameters (EOP). The rotation of the Earth is far from constant at the millisecond level, and even after removing variations due to tidal and seasonal periodic variations, there are still signals which at this point are best described as random. Predicted Earth Orientation Parameters provided by the US Naval Observatory are important in navigation and in artificial satellite data acquisition. The more accurately the Earth orientation can be measured in near real time, the better the short-term predictions become. Artificial satellite observations provide very reliable information on the polar motion components of the EOP series. However, model deficiencies of the nodal precession of the satellite orbit (i.e., changes in the orientation of the orbital plane in space) cannot be separated from the rotational component of the EOP, a major problem. Although Very Long Baseline Interferometry can provide accurate measurement of all three components of Earth orientation using distant radio sources, the time-consuming data reduction gives results only about three weeks after the actual measurement. The precision of EOP determination from lunar laser ranging measurements is not as high as with the other two techniques, but it has the advantage of a quick turn-around time of about 12 hours. Combined with the satellite EOP determination, it can correct for the orientation of the satellite's orbital plane, and improve the predictions. However, the quality and quantity of lunar data depends strongly on atmospheric conditions and on the lunar phase. High humidity, atmospheric turbulence and high background noise can make the detection of lunar returns quite difficult.

## 2 Data acquisition

The basic elements of a lunar laser ranging station are a laser, timing equipment, a telescope through which the laser is beamed to the Moon, and which in case of the French and American stations also collects the returning photons, and a computer. The McDonald LLR has a mirror of 0.76 meter diameter, and a Nd-YAG laser which fires 10 pulses in a second, with 200 picosecond pulse width (full width at half maximum), delivering 120 mJoule of energy per pulse. The laser beam contains approximately $3 \times 10^{19}$ photons. There is a substantial loss of signal due to transmission through the optical elements and the atmosphere, beam divergence, and the distance to the retroreflector. When the laser beam reaches the Moon, it is spread over an area of about 7 km in diameter. The area covered by the retroreflector is $10^{-9}$ times smaller then the beam itself, and the cornercubes spread the beam further, producing a 20 km spot on the Earth.

This results in a factor of $10^{21}$ decrease in the signal; for every 30 laser firings, on average one photon arrives back at the detector, making the lunar laser ranging a single photon detection experiment.

An essential requirement in collecting LLR data is an adequate initial model of lunar dynamics, atmospheric refraction, station coordinates and Earth orientation. Based on this model, the telescope can be accurately pointed to the retroreflector on the Moon, which can be tracked during the observation. The interval timer is started as the laser fires, and it stops when the first appropriate photon reaches the detector. The precision of the epoch timing system is approximately 25 picoseconds. To eliminate most of the non-lunar photons, a filter centered on the wavelength of the laser is placed in the path of the returning beam. A temporal filter, a *range gate*, is also implemented. From the initial model, the predicted round trip travel time can be estimated for a given shot, and the detector allowed to observe only while the gate is open for a limited time interval centered on the predicted arrival time. The most commonly used range gate width at the McDonald Observatory LLR station is 400 nanoseconds. A typical lunar run is about 45 minutes long. During this time, depending on the lunar phase (i.e., on the illumination of the lunar disk), and on the atmospheric conditions, ten to a thousand photons reach the detector. When the Moon is in first and third quarter and the reflectors are in the dark, most of the detected photons are laser returns, but in the intermediate phases from first to last quarter, background photons can overpower the lunar returns because the location of the retroreflector is illuminated by the Sun. The individual returns are later analyzed and compressed into normal points of 1 cm precision. Data accuracy is at about the 2-3 cm level.

## 3 Present filtering method

In addition to the spectral and temporal filter the data are run through a software filtering program. The difference between the predicted and measured round-trip times, that is, the residuals, are calculated using the adopted model. For a perfect measurement and model the residuals should be zero. Different model errors introduce different signatures into the residuals, and the noise photons do not follow any pattern. We assume that the mathematical model we are using is correct except for UT0 (timing) errors in the predicted EOP. During the 45 minute run this error would cause the residuals of the returning photons from the laser shot to lie on a straight line of small but unknown slope. The background photons are still randomly distributed in the gate. Assuming that the background noise is uniformly distributed, to identify the lunar return the analyst looks for clumping, that is for significant deviation from a uniform distribution [RS92]. The residuals are binned (usually into 1 nanosecond bins) and the maximum number of photons expected in a bin is calculated from the total number of detections. The program looks at the bins in pairs, and looks for a significant deviation from this expected number (The slope and the width of the bins can be adjusted in this process). When it finds such pairs all photons in the two bins are identified as lunar returns and compressed into normal points. The EOP are recovered through another step using nightly corrections based on the normal points.

When the signal is strong, the laser returns can easily be identified even by the eye (Figure 1). This approach breaks down if the number of the total returns is small, or if the noise level is high (Figure 2); in such cases the program cannot decide based only on the maximum expected number of returns whether the detected photons are from the retroreflector or not. In this case no photons are so identified, and some data can be lost. The returning laser pulse is wider than the outgoing pulse. The width of the returning pulse, if we can estimate it, could provide additional information about the precision of the normal point, but at present is not taken into account, except that the bins are chosen to be wide enough to contain the smeared pulse. The EOP determination can be improved by increasing the amount of reliable data going into the calculations, and by recovering the parameters close to real time. The Bayesian approach could help in two ways: It could recover data deemed unusable before, and it eliminates the extra step of forming the normal point for EOP determination. In addition to determining EOP, the data are distributed to the LLR community for scientific analysis, and we hope that the Bayesian approach will help provide a better estimate of the uncertainty of the normal points.

## 4 Bayesian analysis

Loredo [Lor90] has given a nice introduction to Bayesian analysis for astronomers. Rather than try to duplicate his excellent discussion, we will summarize the results.

Bayesian analysis requires the scientist to provide two things: A prior distribution $p(\text{model})$ which is a function of all of the model parameters in the problem, and a likelihood function $L(\text{model}; \text{data})$ which describes in a probabilistic sense what sort of data we expect to see given any particular choice of model parameters. To give a simple example, we may be interested in estimating some quantity, such as the distance $D$ to the Moon. The prior distribution represents our knowledge and/or opinions about this parameter prior to taking some data set. It is a probability density, say, that tells us (on prior information) that we believe that it is more likely that the parameter lies in certain subsets of the possible values than in others. The prior distribution is a measure of our own ignorance: if we are very sure of the value of the parameter, it will be sharply peaked, and if we are less sure, it will be more broadly spread out.

The prior distribution can vary from individual to individual for various reasons, including the fact that different individuals usually have different prior information. Under many circumstances the result of a Bayesian analysis may not depend critically on the choice of prior within a fairly wide class of priors. In other cases, care must be taken and if the prior is uncertain in such a way that the results do depend sensitively on the prior, one should carefully investigate the way the results depend on the prior.

The likelihood function $L(\text{model}; \text{data})$ is any function that is proportional to the probability of obtaining certain data, given a model; but it is considered as a function of the model (i.e., the model parameters), since the data are considered fixed. Frequently the particular model is specified by a particular parameter set. In our simple example, each possible value of $D$ corresponds to a particular model, and the probability of obtaining a particular set of data—return timings, for example—depends upon the value of $D$. Data that are consistent with a high value of $D$ are more likely to be obtained if the value of $D$ is high than if it is low, and vice versa. Expressed in the language of conditional probability,

$$L(\text{model}; \text{data}) \propto p(\text{data}|\text{model}) \tag{1.1}$$

The Bayesian prescription tells us that, by Bayes' theorem, the posterior distribution of the parameter given the data, $p(\text{model}|\text{data})$, is proportional to the prior distribution times the likelihood, with a normalization factor that is just the reciprocal of this product integrated over all models (i.e., sets of parameters). Thus

$$p(\text{model}|\text{data}) \propto L(\text{model}; \text{data}) \times p(\text{model}) \tag{1.2}$$

with proportionality factor C given by

$$C^{-1} = \int_{\text{all models}} L(\text{model}; \text{data}) \times p(\text{model}) \tag{1.3}$$

A key characteristic of the Bayesian paradigm is that the results are *conditional* on the data that have been *actually observed*. In other words, the Bayesian analysis does not consider data that might have been observed but were not. This is displayed by the conditional nature of the posterior probability distribution.

All results of interest can be derived from the posterior distribution. For example, by integrating out (marginalizing with respect to) parameters that are not of interest, we can obtain a posterior distribution which is a function of just those parameters that we are interested in. Such calculations are often the most difficult part of a Bayesian analysis, since it is relatively straightforward in many cases to write down the likelihood function and even the prior, but the integration of the posterior may be difficult because it may be complex and not integrable in closed form. To handle this problem, a number of Monte Carlo methods have been developed in recent years that have proved rather effective [Tan93]. We will utilize this strategy in our discussion.

## 5 The Likelihood Function

The lunar laser ranging problem is an interesting one, not only because the signal is very weak but also because the observations are censored owing to the closing of the range gate whenever the first photon is detected. When the reflector is in the Sun, typically 95% of detections are in fact not genuine returns from the laser pulse, and most laser shots do not result in a detection.

Because we are counting discrete events which are for all practical purposes independent, the statistics governing this problem are Poisson. Often it is possible in large-signal situations to approximate Poisson statistics by an appropriate normal approximation, but that route is not available here. We must deal with the Poisson nature of the data at the outset, without fudging.

We approach the problem of writing down the likelihood function similarly to the way Loredo did in another small-signal problem involving a comparable number of neutrino detections from Supernova 1987A [Lor90]. Loredo observed that in a very short interval of time $\Delta t$, the probability that we detect no photons is

$$(r\Delta t)^0 \exp(-r\Delta t)/0! = \exp(-r\Delta t), \tag{1.4}$$

and the probability that we detect a single photon is

$$(r\Delta t)^1 \exp(-r\Delta t)/1! = r\Delta t \exp(-r\Delta t), \tag{1.5}$$

where $r$ is the expected rate per unit time of a detection. By making $\Delta t$ very small, we can ignore the probability of two detections in the interval. In our case, the rate varies with time, $r = r(t)$, because the gate is automatically opened and closed for each shot (closing can be at the end of the window or when a photon arrives), and also because the probability of the arrival of a photon is significantly enhanced during the very short interval that the range coincides with the actual light-time to the moon.

If the intervals $\Delta t$ are disjoint, then the detected events are independent and the likelihood function is just the product of (1.4) and (1.5) over all intervals:

$$L(\text{model}; \text{data}) \propto \exp(-\sum r(t)\Delta t) \prod_N r(t_i)(\Delta t)^N, \tag{1.6}$$

where the data $t_i$ are the times when a photon was detected and $N$ is the total number of detections. Since we only need the likelihood function up to a constant factor, we drop the last term $(\Delta t)^N$ then take the limit as $\Delta t \to 0$ to obtain

$$L(\text{model; data}) \propto \exp(-\int r(t)dt) \prod_N r(t_i), \tag{1.7}$$

where the integral is taken over the entire time that the range gate is open (or, equivalently, over the entire duration of the observation set, noting that the probability of detection $r = 0$ when the gate is closed).

The same result is obtained when we analyze the problem using the approach suggested in [Tan93, §2.1].

For our problem, we presume the following form for the rate $r(t)$:

$$r(t) = \begin{cases} 0 \text{ if the gate is closed} \\ r_{bg} \text{ if the gate is open but the return is not expected} \\ r_{bg} + r_s \text{ if the pulse return is expected} \end{cases}$$

Here, $r_{bg}$ is the background detection rate per unit time and $r_s$ is the signal detection rate per unit time.

What do we mean by saying that return is expected? We presume that the width of the returning laser pulse is very brief, a few hundred picoseconds. The pulse travels to the moon and is reflected back to the Earth, arriving some time later. We do not know the time of arrival, since that is what we have to determine. We have a model for return time that is dependent on certain parameters. We predict (on the basis of our model) that at some time $t$ the center of the pulse will arrive; that is the expected time for the return of the pulse.

By using the "box" function $\Pi(x)$, which is 1 when $-1/2 \leq x \leq 1/2$ and zero otherwise, we can express a simple model for $r(\mathrm{t})$ as follows:

$$r(t) = \sum_i \Pi((t - t_{gi})/a_i)(r_{bg} + r_s \Pi((t - t_{ri})/a_{\text{pw}})) \tag{1.8}$$

where in the i$th$ shot, $a_{\mathrm{i}}$ is the length of time the range gate is open (normally 400 nanoseconds, but shorter if a detection occurs), $t_{gi}$ is the mean of the opening and closing times of the range gate, $a_{pw}$ is the pulse width, and $t_{ri}$ is the predicted time of the pulse return. It is assumed that there is an unknown bias in the ephemeris of the Moon, so that there is an unknown offset in the Moon's distance; and that furthermore, the offset varies in time, linearly to first order, by an amount that is also unknown. Thus we can write

$$t_{\mathrm{ri}} = b + c(t - \bar{t}), \tag{1.9}$$

where $b$ is the expected pulse return time at time $t = \bar{t}$, $\bar{t}$ is the midpoint of the data take, and $c$ is the slope of the unknown pulse return time. The unknown parameters $b$ and $c$ of the model are to be determined.

The expression for the likelihood function can be simplified by carrying out the integration and product. We see immediately that

$$\int r(t)dt = Tr_{\mathrm{bg}} + (m - k/2)r_{\mathrm{s}}\Delta t, \tag{1.10}$$

where $T$ is the total time that the range gate was open, $\Delta t$ is the assumed width of the returning pulse, $k$ is the number of detections that occured within $\Delta t/2$ of the expected pulse arrival time $t_{ri}$, and $m$ is the number of times that the range gate was open when the pulse return was expected at time $t_{ri}$. The factor 1/2 in the last term takes into account the fact that the pulse is detected, on average, halfway through the pulse width, an approximation that is convenient but inessential. We adopt it for the purpose of this calculation.

The product can be written up to a constant factor as

$$r_{bg}^{N}\left(1+\frac{r_s}{r_{bg}}\right)^{k},$$

where as before $N$ is the total number of detections.

## 6 Prior distribution

The formal unknown parameters in this problem are the detection rates $r_{bg}$ and $r_s$, and the parameters $b$ and $c$ that describe the expected time of pulse arrival. Usually one would have a pretty good idea of the errors of the ephemeris, and can bound $b$ and $c$ reasonably well. For this investigation we usually adopted a simple prior that is uniform within a range typical of what might be expected for the parameter and zero outside that range. In some cases we adopted a normal prior. For some runs we made life as difficult as we could by assuming that the pulse could return at any time that the range gate was open and setting the width of the prior accordingly. In real life, that is far too pessimistic, but it allowed us to find out just how well we could pin down the actual return time from the data. We restricted the slope c so that over the entire run the expected time of arrival would not vary by more than 10 nanoseconds. The width of the priors on $b$ and $c$ were variable, i.e., we allowed ourselves to be very sure or quite ignorant, depending on the run.

## 7 Gibbs sampler

Once we have determined the prior and the likelihood, we can write down the posterior distribution (up to a constant factor). Now the fun begins. We consider some of the parameters, in particular $r_{bg}$ and $r_s$, to be "nuisance parameters." That is, we are not much interested in their actual values. We are most interested in the marginal distributions of $b$ and $c$, which provide the desired information about the Moon's orbital motion. These are obtained by integrating over all of the other parameters to obtain marginal posterior distributions $p(b|\text{data})$ and $p(c|\text{data})$ from the complete posterior $p(b, c, r_{bg}, r_s|\text{data})$. We might also be interested in the marginal distributions of $r_{bg}$ and $r_s$; during the initial runs, we assumed that we knew $r_{bg}$ and $r_s$; this information is actually pretty well known since the characteristics of the laser pulse and the reflection process on the Moon under ideal conditions are well understood after over 25 years of ranging, but can be influenced by weather and other conditions. In later runs we allowed these too to be uncertain.

Our integration method was the Metropolis subchain Gibbs sampler, as suggested in [Mül91] and described in [Tan93, §6.5.3]. The idea behind the Gibbs sampler is to generate a Markov chain using the posterior distribution to generate each next step in the chain. The transition probabilities at each step are prescribed by the posterior distribution, in such a way that one is more likely to make a transition from a region of lower posterior probability to one of higher posterior probability than vice versa. Thus, the Markov chain tends to spend more time in regions of high posterior probability than in regions of low posterior probability. The Markov chain is defined so that a step is taken first in $b$, then in $c$, then in $r_{bg}$, then in $r_s$, say, to constitute one iteration. Then the process is repeated indefinitely, always starting the new step where the old one left off. After each iteration, the current values of the parameters obtained are tallied separately, to obtain marginal distributions for each parameter. It can be shown that under reasonable conditions that are usually met in practice, the resulting Markov chain yields marginal distributions that approach the actual marginals in the limit.

The difficulty in carrying out this prescription is that the one-dimensional distributions of $b$, $c$, $r_{bg}$ and $r_s$ are themselves difficult to sample from, since in general (and in our case) they do not belong to the narrow class of distributions for which analytical sampling schemes exist. However, many probabilistic schemes exist for such sampling. We applied Müller's ideas to generate trial steps from the one-dimensional distributions at each iteration using a Metropolis-Hastings approach. The details are described in [Tan93], but the basic procedure is as follows. Suppose we are ready to generate the next step in a parameter, say $b$. From a symmetric, but otherwise *arbitrary* distribution $q(\Delta b)$, generate a delta step $\Delta b$. The new trial value of the parameter is $b^* = b + \Delta b$. We accept or reject this trial value probabilistically based on a simple function of the posterior distribution at the two points $b$ and $b^*$. In particular, with probability

$$\alpha(b^*, b) = \min\left\{\frac{p(b^*, c, r_{bg}, r_s|\text{data})}{p(b, c, r_{bg}, r_s|\text{data})}, 1\right\} \qquad (1.11)$$

accept $b^*$ as the new step; otherwise keep $b$. Then proceed in the same way with $c$, $r_{bg}$ and $r_s$ to complete the iteration.

A key advantage of this method is that it is unnecessary to work with the normalized posterior probability, since the normalization factor cancels out of the expression for $\alpha$.

We chose $q(\Delta b)$ to be *uniform* over an interval that was typically 10% of the allowed variation in the particular parameter in question (i.e., the range of the prior for that parameter). This is admittedly crude, and other choices are to be investigated, but it is remarkable how effective this choice turned out to be. Having done this, we then repeated the sample/accept/reject procedure for the the second parameter, then the third, and so on, until all parameters had been sampled once. One trip through this procedure for all of the parameters constitutes one iteration of the Gibbs sampler. At this point the current position of the point in parameter space is noted, and the process repeated. Typically we would iterate on the order of 10,000 times to allow the calculation to "burn in," and then start recording the data for another 10,000 iterates. Finally, histograms of the later iterates' marginal distributions were tallied and plotted.

## 8 Results

The data used were a set of simulated data from a simulator that has been used previously in the lunar laser ranging project ([RJ95]). The data consisted of 14,400 shots, of which 2313 generated detections. Most of those detections were of noise; only 8 were actual returns, or 0.34%. This particular set of data was intended to represent data of poor quality. Figure 3 shows the actual data as simulated.

The slope $c$ generated by the simulator was $+1.4$; we used various starting values to see if the Gibbs sampler would find the actual slope. Also, we sometimes started the procedure well away from the actual bias to see if we would converge on the actual bias (which for these data was $b = +2.35$ nanoseconds; we would typically start at positions like $\pm 10$ nanoseconds from the true value; the range gate width was 400 nanoseconds for these data).

Our biggest question was whether we would be able to pin down $b$ and $c$ sufficiently well to tell when the returns came. The likelihood function can be expected to have a very narrow peak near the return; would the Metropolis subchain Gibbs sampler find it? We tried a number of runs, first assuming that we knew the detection probabilities and later on computing their marginals as well. The results were really quite gratifying, as can be seen in the figures for one of the later runs.

In the run shown in Figures 4 and 5, we adopted a normal prior for b with mean 0 and standard deviation 15 nanoseconds, cut off at $\pm 30$ nanoseconds and started at $b = 0$ nanoseconds. This prior represents the typical state of knowledge for actual runs. The prior on c was still relatively pessimistic: uniform for $-10 \leq c \leq 10$ and 0 outside that range. The priors on the rates were uniform for rates $\geq 0$, and both the signal and background rates were estimated. The results were that the median of the smallest posterior distribution of $b$ was $+2.27$ nanoseconds (true value $+2.35$ nanoseconds) with the smallest 80% Bayesian confidence interval $(+2.25, 2.35)$ nanoseconds. The smallest 95% Bayesian confidence interval was $(+1.89, 2.44)$ nanoseconds. Thus, the tails of the distribution are quite heavy relative to a normal distribution, and the center is very strongly peaked near the true value. This run is actually quite typical, and it shows that despite the fact that the posterior probability is strongly peaked, the Metropolis subchain Gibbs sampler was able to handle it well and find the peak with no apparent difficulty.

The posterior distributions of the slope and of the rates are roughly normal, as expected, and not very interesting, so will not be discussed here.

Figure 5 shows the 8 actual lunar returns, together with the median line from the Gibbs sampler calculation. The vertical scale has been blown up substantially; the error bars in the vertical scale are $\pm 200$ picoseconds (0.2 nanoseconds). The fit is very satisfactory. Indeed, we have been delighted with the way the Gibbs sampler homes in on the input answer even when the signal is so small that one can't pick it out by eye. This is very promising for practical application.

## 9 Conclusions and Future Research

This paper describes a demonstration in principle that a Bayesian approach can be used to analyze the results of lunar laser ranging experiments. With poor data and somewhat pessimistic priors, one clearly and unambiguously picks out the return signal, even though over 99% of the photons detected were noise.

Much work remains to be done. We have merely scratched the surface of this interesting project. Extensive simulations need to be run, for all kinds of data. Real data have not yet been considered and must be analyzed. It will be particularly interesting to see if and how well this method can recognize signal where the older method cannot. We also need to investigate methods of improving the Gibbs sampler, e.g., convergence criteria and methods of deciding ideal step sizes for trial steps.

## 10 References

[DBF+94] J.O. Dickey, P.L. Bender, J.E. Faller, X.X. Newhall, R.L. Ricklefs, J.G. Ries, P.J. Shelus, A.L. Whipple, J.R. Wiant, J.G. Williams, and C.F. Yoder. Lunar laser ranging: A continuing legacy of the apollo program. *Science*, 265:482–490, 1994.

[Lor90] T.J. Loredo. From Laplace to Supernova 1987a: Bayesian inference in astrophysics. In P. Fogère, editor, *Maximum Entropy and Bayesian Methods*, pages 81–142. Kluwer Academic Publishers, Dordrecht, 1990.

[Mül91] P. Müller. A generic approach to posterior integration and bayesian sampling. Technical report 91-09, statistics department, Purdue University, 1991.

[RJ95] J.G. Ries and W.H. Jefferys. Application of bayesian statistics to lunar data analysis. *Bulletin of the American Astronomical Society*, 27(3):1200, 1995.

[RS92] R. L. Ricklefs and P.J. Shelus. Poisson filtering of laser ranging data. In *Proceeding of the Eighth International Workshop on Laser Ranging Instrumentation*. Annnapolis: NASA conference Publication 3214, 1992.

[Tan93] Martin A. Tanner. *Tools for Statistical Inference*. Springer-Verlag, New York, 1993.

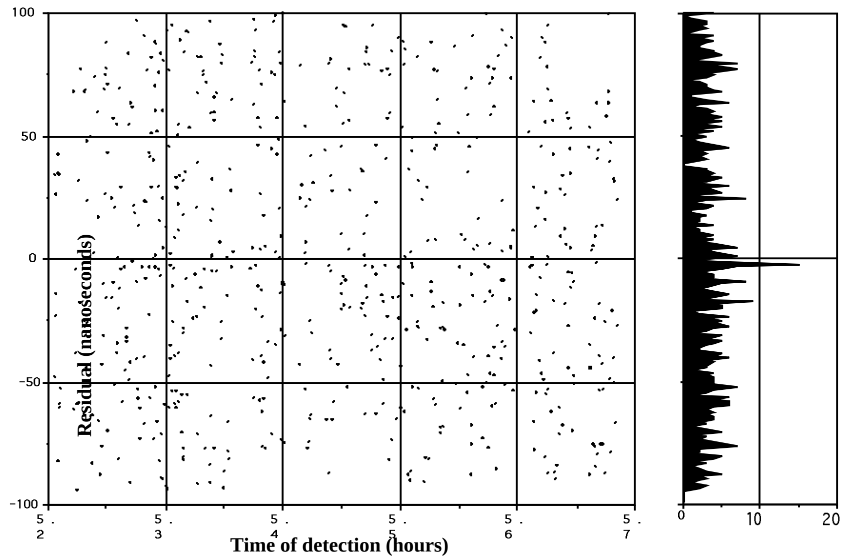


FIGURE 1. Fair-to-average data: The difference between the observed and predicted range is plotted as a function of time for an actual lunar run. The clumping is obvious to the eye on the histogram. The traditional filtering method can identify the laser returns.

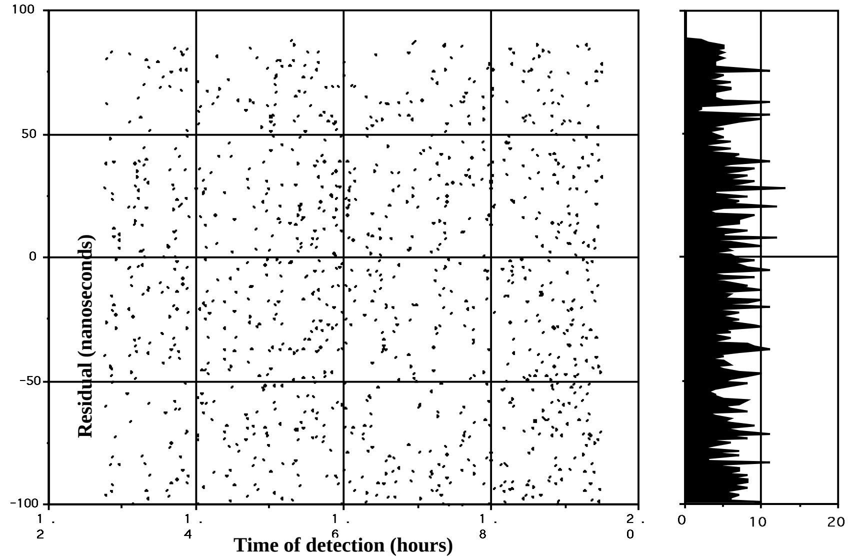


FIGURE 2. Poor data. The data are marginal, and the clumping is not obvious. The traditional filtering method breaks down. However, the ranging crew, with years of experience at lunar laser ranging, indicated that they think they obtained real laser returns in this run.

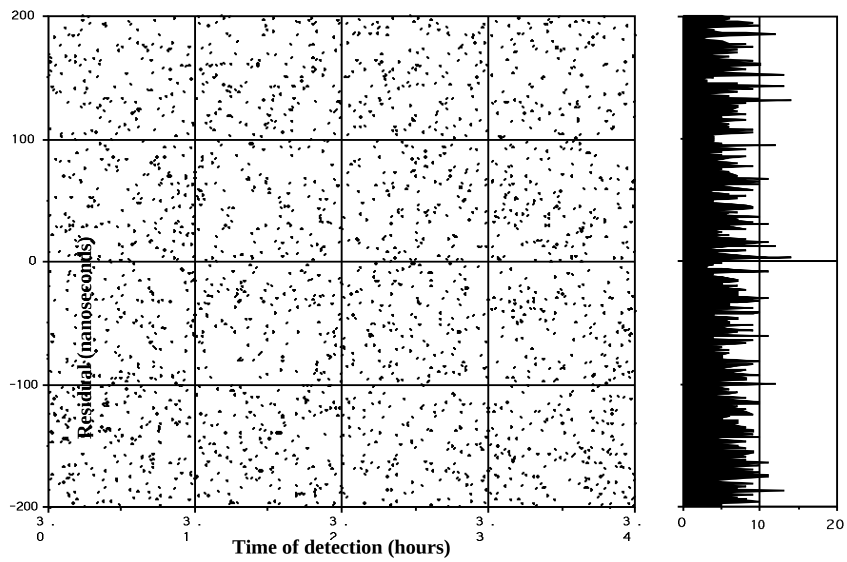


FIGURE 3. Simulated poor data: Simulated returns are plotted as a function of time. The 8 actual returns can barely be seen on the histogram and are difficult or impossible to see on the scatter diagram.

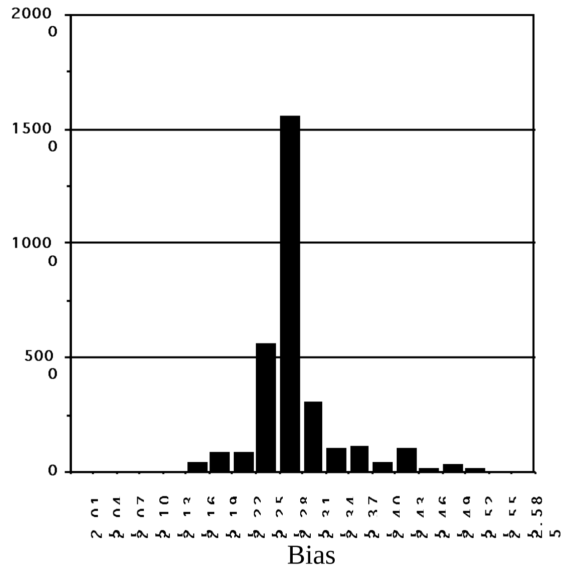


FIGURE 4. Posterior marginal distribution for the bias $b$ that represents the expected return of the photons from the laser at time $t = \bar{t}$. It is strongly peaked near the true value; 80% of the posterior probability is contained within an interval of 0.1 nanoseconds, and 95% within an interval of 0.55 nanoseconds.

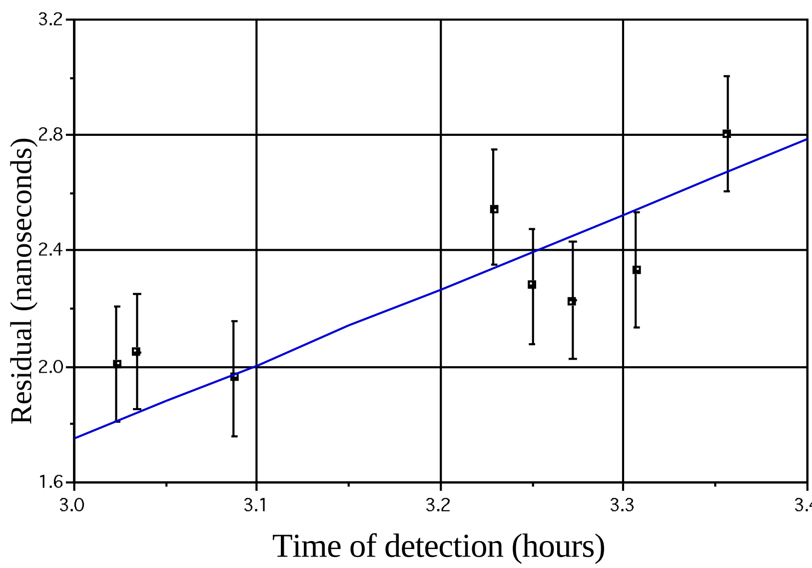


FIGURE 5. The eight actual returns are plotted together with the median line from the marginal distributions of $b$ and $c$. The error bars indicate the uncertainty in return time.