Statistics, Shooting and the Myth of the Three Shot Group

By RocketmanOU

A number of years ago, I showed a shot grouping from one of my favorite rifles to my grandfather, a World War II veteran. “Looks like the sights are out of adjustment,” he remarked, “I can show you how to fix it, if you like.” Of course, he was correct – the group was high and right by several inches. In my mind at the time, it didn’t matter – sights can be adjusted, but look at how *consistent* the shots were! Five shots, all overlapping! As time has gone by, however, I have given the subject a lot more thought, and come to some conclusions, as well as some continued questions. These relate to the application of statistical analysis (from a mechanical perspective) to shooting, as well as some of the paradigms traditionally used to describe accuracy.

On shooting forums the world over, it is quite common to see three shot groups used to demonstrate the accuracy of a rifle+shooter combination (or, as is the case when shooting with numerous sandbags, just the rifle). Rifle shooting, of course, has innumerable variables involved, which is most of the attraction to me - it is the combination of (not limited to, but including) the ammunition, the firearm, and most importantly, me, that decides whether a shot will hit my intended target or become a devilishly aggravating puff of dust next to the steel. The general conclusion is that if my rifle can hold, let’s say, a 0.4 MOA group like the one in my awesome photo (I SWEAR it was at 300 yards!), then I can hit a target X size at Y distance, right? There are two misconceptions here: first, I am not necessarily measuring what I think I’m measuring and second, my three or five shot group is not necessarily statistically significant. The two misconceptions are directly linked, but we’ll start by addressing the first one.

Disclaimer: *what follows is a discussion of normal distributions and statistics. It is assumed that patterns such as vertical stringing from a heated barrel, lateral dispersion from jerking a trigger, etc. are negligible/non-existent. This is a discussion of <u>normally randomized</u> variations in shot placement, and as such, ignores the physical behavior of the mechanical system in question. There are lies, damned lies, and statistics.*

---

1. The Measurement Problem

Very often, we measure one thing that we correlate to another instead of measuring directly. One such example is a bulb thermometer – it does <u>**not**</u> measure temperature. It measures the expansion of mercury (or alcohol, as the case may be), which is a phenomenon we correlate (very closely) to temperature. To get an accurate measurement out of such a *transducer* (see what I did there?), we need to be sure that it is very strongly a function of the one variable we mean to measure and not strongly a function of other things that we don't want to measure. In the case of the thermometer, for example, if the column was not sealed, it would behave as a barometer **and** thermometer, reacting to both air pressure and temperature at the same time. This would effectively make it useless for measuring either individually.

Let's just skip forward a bit and say that my grandpa was right. I ask my students this question all the time: "What *really* matters?" Grandpa took one look at my target and saw instantly what he knew matters – my shots would have missed my target at an appreciably longer range than 100 yards –all 5 of them, with remarkably similar deviation. To determine how to measure something (i.e.: which statistic is important), It's sometimes useful to force ourselves to make a complete statement about what we're trying to do. For shooting, mine looks something like this: I want to be able to hit a small target at a long distance consistently. This means I need consistency (low levels of deviation) and good adjustment – a tight group means nothing if off target, and an average on-target group means nothing if the individual shots are all over the place. In short, I need both precision *and* accuracy. For most of us, this isn't news.

Interestingly enough, however, we often shoot a 3-shot or 5-shot group, and we're content with simply measuring the diameter of the group as a descriptor for our rifle+shooter system. But what does that actually measure? When we measure the diameter of our group, we're measuring the extremities of the variation in our shots. Unfortunately, this doesn't tell us <u>directly</u> whether our shots will hit the target, or how small of a target we may be able to hit at a given distance. What it tells us is the dispersion of that specific group. Statistics teaches us that we can use that dispersion to tell us about probabilities, but the only way to know whether I can hit a 2" diameter circle at 300 yards is to *actually shoot at* a 2" circle at 300 yards, and see how much of the time I hit it.

So the question remains – how can we use group size to talk about hit probabilities? Specifically, given a group dispersion *X* wide at *Y* distance, what's the probability that I will hit a 2" diameter circle at 300 yards? Spoiler: if you're shooting a 2" 3-shot group at 300 yards, it's not even nearly 100%. Since we're measuring one thing using a *separate, indirect* characteristic, we need to use some math to help us assess the uncertainty of our newfound transducer. This brings us to our statistics problem(s).

2. The Statistics Problem(s)

Let's assume, for a moment, that our shots are normally distributed about the mean of the overall group in a radial fashion. For reference, the probability density of the standard normal curve is shown below.

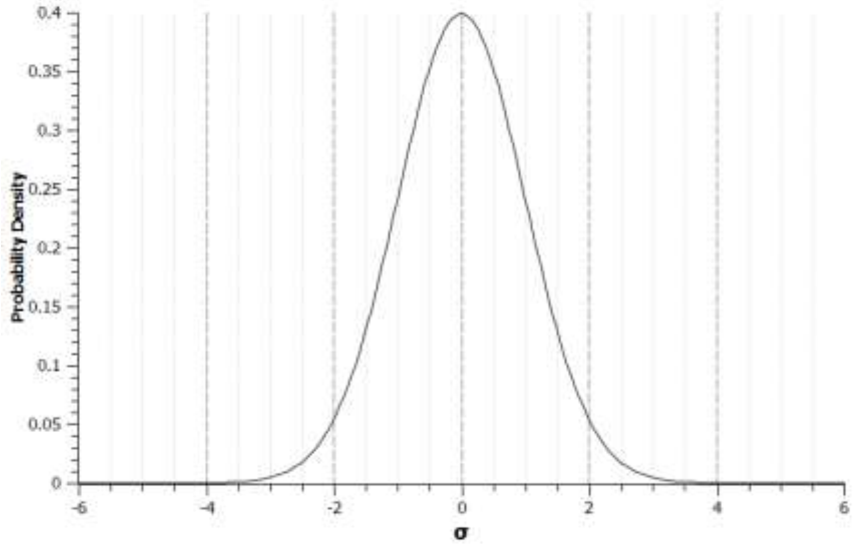

We are reminded that the probability of an event is equal to the area under the curve between two distances from the mean, e.g.: the probability that a random value x will be between -1σ and +1σ is equal to the area under the curve between those two points (68.3% in this case). Such a distribution applied in radial fashion would look like Figure 2:

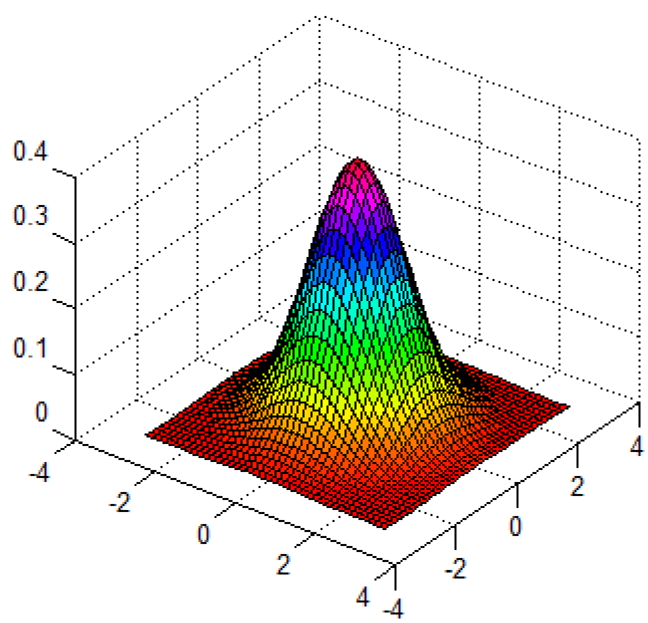

It's important to remember that the peak in the graph **does not** denote that the highest probability for shot placement is in the exact center of the target. Most shooters will agree – there are precious few times that we actually hit the x exactly in the middle. We can only express probabilities over an area, e.g.: "there is an X% chance the shot will land inside a Y-inch radius of the mean." This is calculated by taking the indefinite integral of the probability density function. A plot of this integral is shown on the next page.

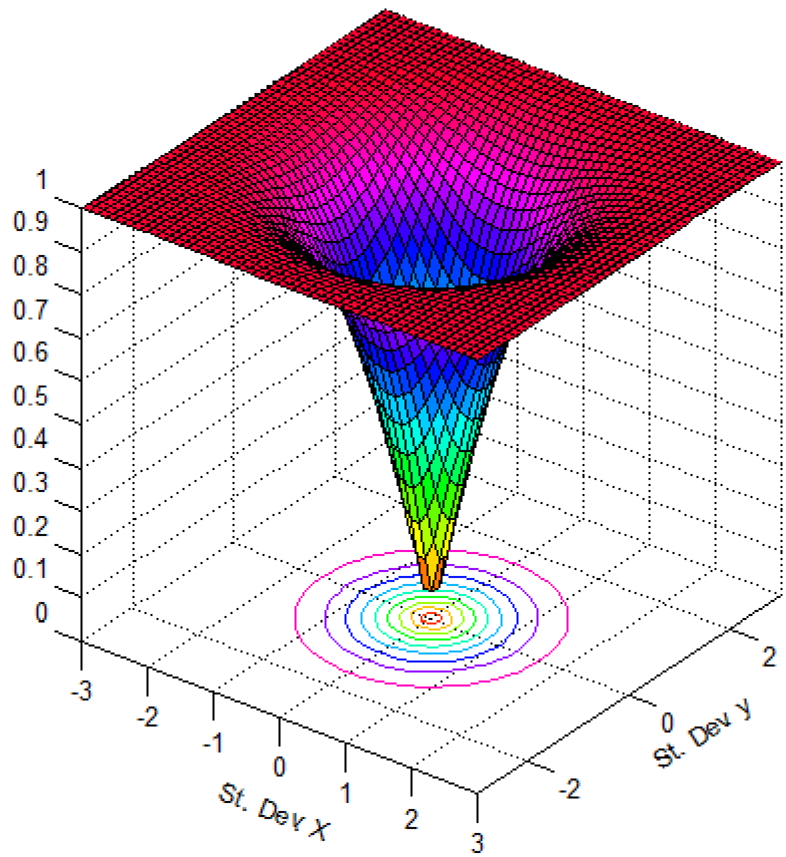

The above figure can be interpreted as follows: the probability that a shot will fall within x (or y) standard deviations (the flat plane with the circles on it is the x-y plane) is equal to the surface value in the z direction. This format is inconvenient for actually extracting data, but we can see the general relationship – probability is a minimum exactly at the mean, but increases as we increase radius. This makes sense physically, as it's harder to ensure all one's shots are inside the x-ring than it is to ensure they're all inside the 8-ring. A much more informative graphic is shown below: simply a 2-d, grayscale view of the contours shown above. Each successively larger ring denotes an additional 10% probability that a shot will land inside that ring, starting with 10% in the center.

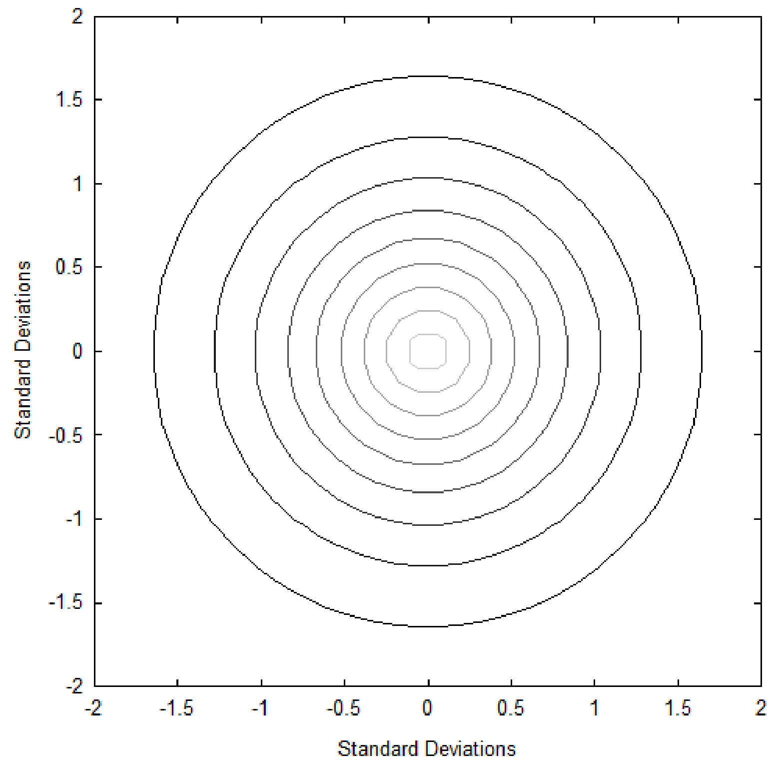

Now, if we shoot some groups and measure the standard deviation of the shots from the group mean, we can scale this distribution according to our rifle's performance. Let's take a look at an example. Four 3-shot groups taken at 100 yards are depicted below, all with similar dispersions. Per usual practice, we would call this a sub-moa setup, as the average group size is 0.85" at 100 yards, or about 0.81 MOA.

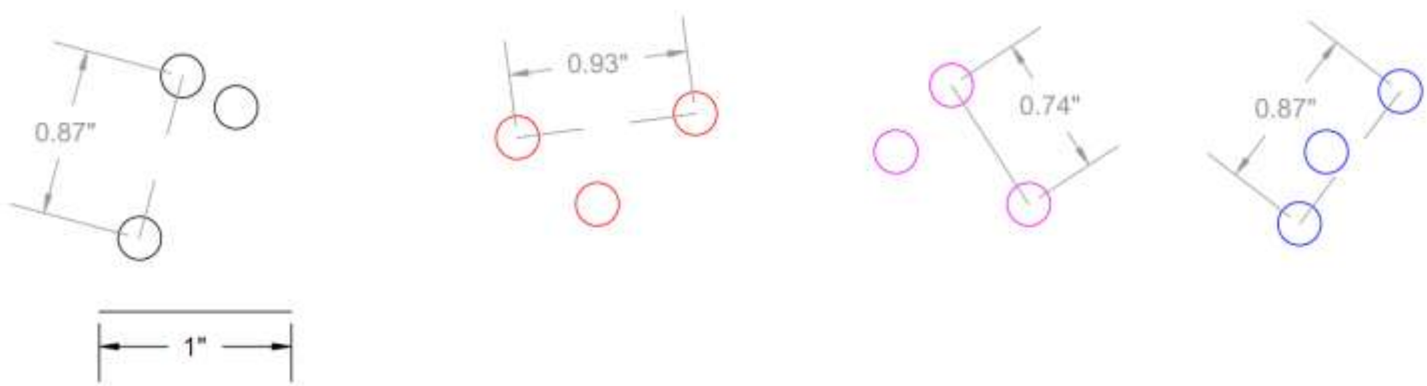

Next, we can load these shots into software (or draw them on graph paper), and get some interesting information about them. I use a script I wrote in FreeMat, an open-source MATLAB clone, to work the same way as On Target, but with quite a bit of additional functionality. Whatever way we choose, we can find the location of the mean of the group (the average shot location), as well as the resulting distance to each shot from that average. We can find the standard deviation, and apply the same process as above to find our probability rings. The figure below shows such a result from the first group in the previous figure. The green ring represents the 68.3% ring, while the red ring represents the 95% ring. In other words, roughly two-thirds of the shots in the group should be inside the green ring (true), and 95 out of 100 shots should be inside the red ring (also true).

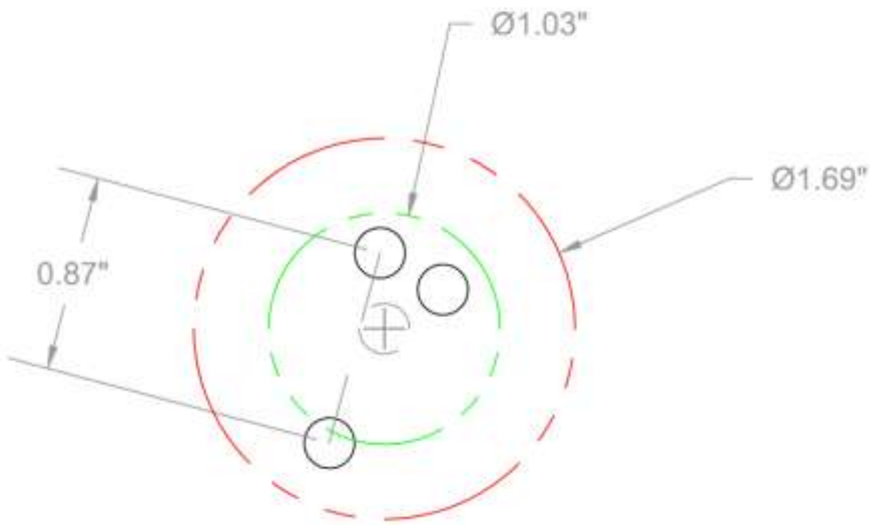

Right off the bat, we can make a rather important observation - our 0.83 MOA group, should we fire any more shots, will very likely *grow in size*. In fact, if we fire 10 shots, it's likely (based on our initial 3-shot group) that one will be outside of almost 1.7 MOA! There are several problems with this supposition, but one is very dominant – standard deviation is a *function of the number of points involved*. Remember what was said earlier about a transducer being a function of more than one variable? When the number of points is low, the effect of that variable on standard deviation is gigantic. For comparison, let's say we shot a 10-shot group that had the same maximum spread:

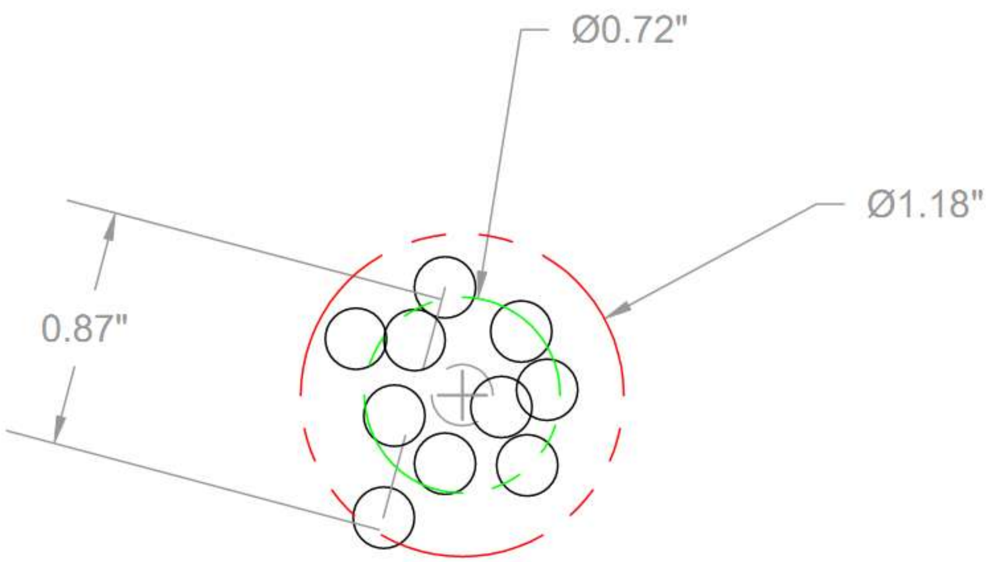

We see that the standard deviation decreases to the point that a reasonable number of shots fall inside the maximum spread diameter. In short, the more shots taken, the more accurate this picture becomes, with limiting returns. Just wait – *it gets better*.

When we take measurements, it is good practice to take a number of measurements and look at their overall behavior, as opposed to simply looking at one or two data points. Why? Here's a quick exercise – estimate the average height of all members of a given demographic by averaging the heights of a few people you know in said demographic. How accurate do you think your answer will be? It turns out that, for a normally-distributed dataset, we can calculate this – hooray for math! This is generally referred to as uncertainty in the mean value. Long story short (too late), often we use a *sample* to describe a *population*, which has errors because it's *not* the population. In other words, we're measuring one thing by measuring something else, then correlating the two (sound familiar?). When it comes to normal distributions, we can pretty accurately assess the uncertainty in this correlation.

Unfortunately, we don't have very good models for translating a sample's standard deviation into the population deviation. A brief look at the previous two graphics is a good example of why. What we *can* do, however, is estimate the population mean, which is very important here. The groups shown two pages back are fine and dandy, but they don't establish a zero. Hitting the target doesn't just mean shooting a tight group – it means shooting a tight group centered as well as can be on the target. Let's look at that first group again. The gray circle with the cross is the group's average.

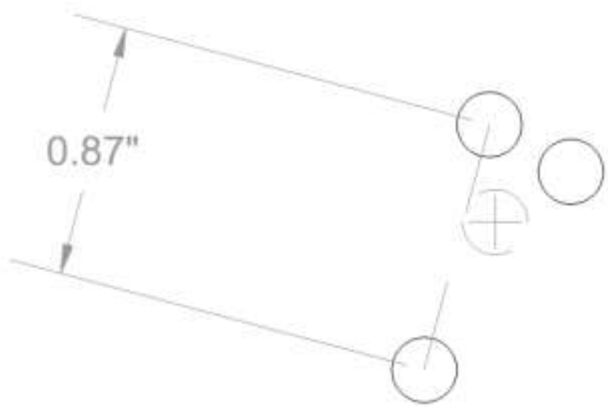

Now, thinking back to our question about estimating average height…how accurate do you think the gray circle is in describing the center of the overall population (the **actual** zero of the rifle)? Since we're using fewer than 30 data points, we consult the work of a certain Mr. Gossett (not Lou) who developed the student's t-table. We find quickly that for a 95% confidence (1 time out of 20, we'll probably be wrong), the uncertainty in the actual zero of the rifle based **on these three shots** is as shown:

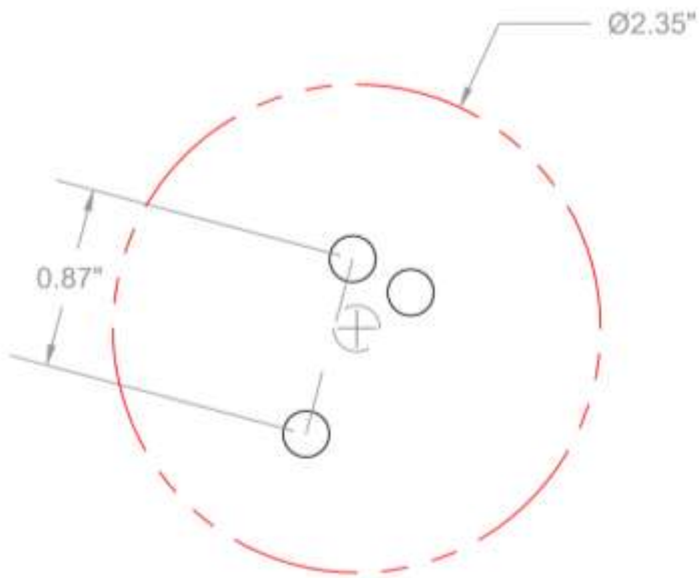

In case it's not readily apparent, this is a pretty terrible level of uncertainty. Quite realistically, what this means is that all those sub-moa groups would likely stack up something more like this:

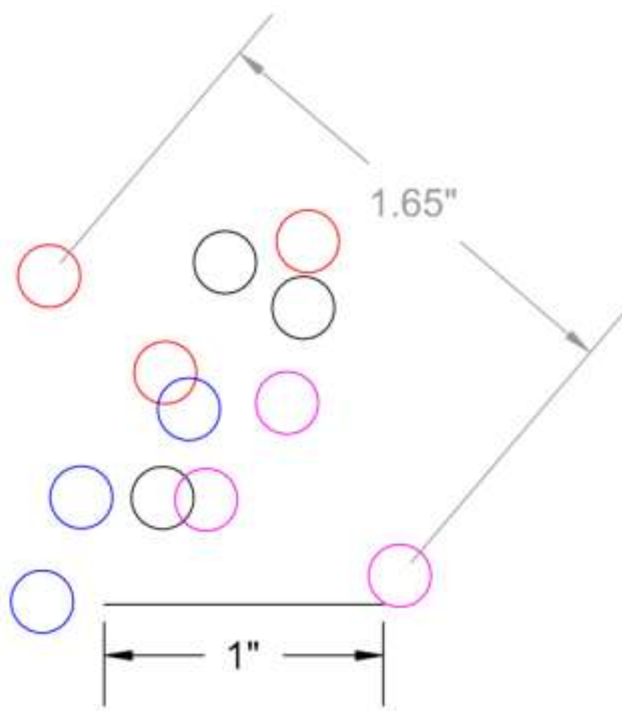

Again, we find that the actual distribution of shots is easily twice the size suggested by our 3-shot groups. When using a 10-shot group of the same extreme spread (like we did before), however, we find that the uncertainty in the mean value drops off very quickly. Compare the following figure to one at the top of this page:

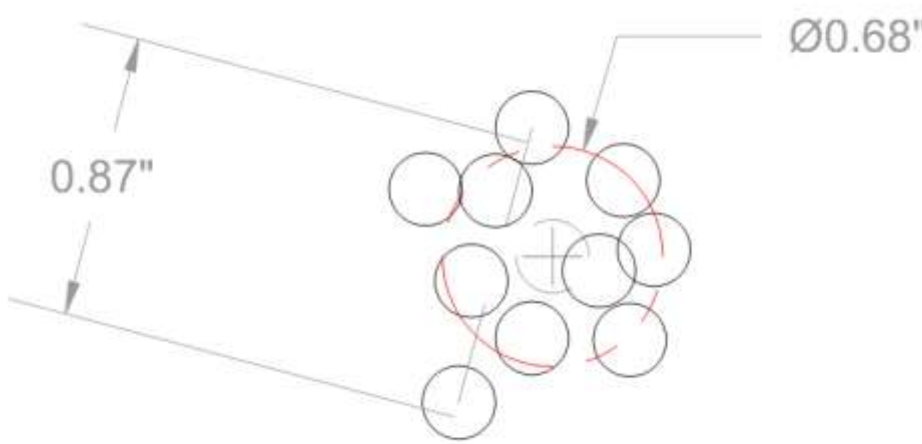

Our estimation of the location of the zero is *much* more accurate (95 chances out of 100 that we're within 3/8", as opposed to 1-1/8"). This translates to a much more accurate picture of where the rifle is placing shots, and how consistently it does so. There is even more to this conversation, but we'll leave it here so as not to kick a dead horse.

3. Conclusions/Continuing Questions

I'll do this as a bulleted list here for those of us who hate math/reading/attention span:

- 3-shot groups **cannot** precisely describe the actual distribution of shots. 10-shot groups **can** precisely describe the actual distribution of shots.
- 3-shot groups **cannot** precisely describe the zero of the rifle nearly as well as 10-shot groups.
- To what degree are shots normally distributed? Is a different probability density function more appropriate?
- What is the break-even point for number of shots/accuracy of estimation? How well does this correlate to longer distances?
- Final conclusion: Grandpa was right. What counts is hitting the target. Realistic target engagement **>** punching paper at 100 10 shots at a time **>** punching paper at 100 3 shots at a time **>** sitting in my office writing about statistics.