

29 How to lower your putting score without improving

B. Hoadley
Middletown, New Jersey, USA

Abstract While playing golf, after you have mercifully reached the green, you have to decide where to target your putts. You might say that it is obvious - target the hole. But then the probability of being short is about 0.5. The conventional wisdom is to target beyond the hole, so that you won't be short. The pundits use expressions like, "never up never in." But how far beyond the hole? Here the advice gets diverse and qualitative. Dave Pelz (1989) says 17" for all putts. My pro says, "for long putts target a two-foot circle around the hole and for short putts target about one foot beyond the hole." The purpose of this research is to provide more specific quantitative advice. An approximate solution is to target a distance beyond the hole given by the formula:

[Two feet]^{*} [Probability of sinking the putt]. This solution applies to any putt. The dependence on putting ability and the difficulty of the putt is captured by the probability of sinking the putt. This tactic saves about 1 putt per round over conventional wisdom. The solution is derived from a Markov decision model of the putting process.

Keywords: Dynamic Programming, Markov Decision Model, Optimization, Probability, Putting, Simulation, Statistics, Tactics.

near 1. Another example is where one side of the hole is flat but the other side is very curved. Then you might target the flat area.

The next idea that I use in my model is the *footprint of success*. If the stopping position of a putt is in the footprint of success, then the putt sinks. It is an oblong area about 4 feet long on the far side of the hole and includes the hole. It is defined precisely in Section 3.

The sequence of events during putting is as follows: (i) select a target and putt, thereby generating a random stopping position; (ii) if the stopping position is in the footprint of success, then you are done putting, otherwise repeat steps (i) and (iii).

The first step in finding the optimal target is to note some limiting cases. If the probability of sinking the putt is near zero, then the optimal target distance beyond the hole is near zero. If you are not going to sink the putt, then you should minimize the expected distance for the next putt. To do this, you target the hole itself.

If the probability of sinking the putt is near one, then you should pick a target that maximizes the probability of sinking the putt. So the target should be in the middle of the footprint of success, which is about two feet beyond the hole.

So as the probability of sinking the putt goes from 0 to 1, the optimal target goes from 0' beyond the hole to 2' beyond the hole. Therefore, I assume that the optimal target beyond the hole is a monotonically increasing function of the probability of sinking the putt. This function starts at 0' and goes to 2'.

I then consider a two dimensional tactic space defined by the parameters x and y of a beta cumulative distribution function,

$B(p|x,y)$. This is a monotonic function of p , which goes from 0 to 1. p is the probability of sinking the putt, and the tactic is to target the putt for (Two feet)^{*} $B(p|x,y)$. I then show via simulation and sequential search that $x=1, y=1$ (i.e., $B(p|x,y)=p$) is nearly optimal.

1 Introduction

The problem addressed in this paper is introduced in the abstract. To attack the problem, I need a model of a single putt. Associated with a putt, is the position on the green that the ball would stop if there were no hole. I call this the *stopping position*. Of course, if the putt is sunk, then the ball never gets to the stopping position, because it falls in the hole. The stopping position is modeled as a random position on the green. The probability distribution of the stopping position depends on (i) the difficulty of the putt, (ii) the skill of the putter, and (iii) the target selected by the putter. The target selected by the putter is the expected value of the random stopping position. The question addressed in this paper is: what is the optimal target for each putt?

I assume in this paper that the green is flat. The advice derived from the flat green assumption is useful for most non-flat green scenarios. However, there are exceptions. For example, suppose the hole is at the edge of a cliff. Then you would aim short of the hole until the probability of sinking the putt was

2 Random model of putting

To model a single putt I use a polar coordinate system centered at the ball. Figure 1 shows the position of the ball and hole in polar coordinates. The coordinates of the ball are $(0,0)$. The first coordinate is the 0 angle and the second coordinate is the 0 radius. The coordinates of the hole are $(0,s)$. So the distance between the ball and the hole is s .

The stopping position of the ball is the random position (A,R) . A is the random angle of the putt and R is the random radius of the putt.

I define the target of the putt to be the expected value of (A,R) , namely $(E[A], E[R])$, where $E[\cdot]$ is the expectation operator.

I assume that A and R are statistically independent with Normal (sometimes called Gaussian) probability distributions. It is clear that the target angle should be 0, so $E[A]=0$. The standard deviation of A is denoted α . The mean and standard deviation of R are denoted μ and σ respectively.

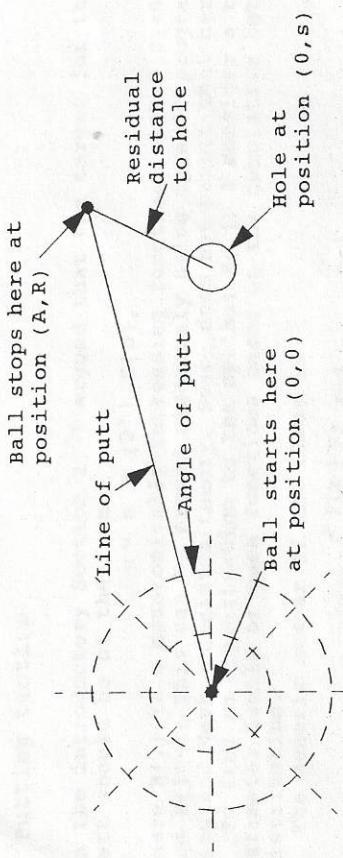


Fig.1. The ball and hole in polar coordinates.

It is clear that as μ get larger so does σ . A quick and dirty empirical study on a putting green suggested the relationship $\sigma = k\mu$, where k is a constant. The same study suggested that σ does not depend on μ .

In a putting situation, to select a target, you just select μ , which will depend on s . The parameters (α, μ) quantify the ability of the putter and the difficulty of the putt.

Note the assumption that the ability of the putter is given. I will show how to reduce the expected number of puts without improving the ability - hence, the title of the paper.

3 Probability of sinking a putt

Figure 2 shows the footprint of putting success (cross hatched area along with the hole). If the stopping position (A, R) is in the footprint, then the putt is sunk. The footprint is defined in terms of s =[distance to the hole], b =[footprint distance beyond the hole], and d =[radius of the hole]. From empirical studies I found that $b=4$. The radius of the hole is $d=2$.

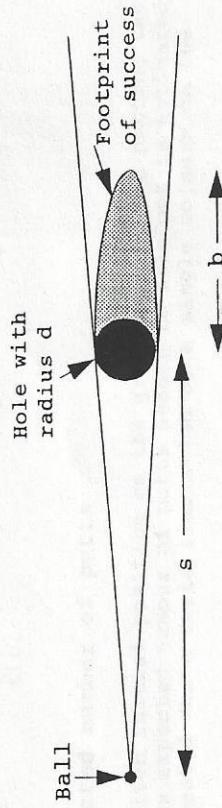


Fig.2. Footprint of putting success.

As an approximation, the conditions for putting success are

$$-\tan(d/s) < A < \tan(d/s)$$

$$s < R < s + \left[1 - \left| \frac{A}{\tan(d/s)} \right|^2 \right] b \quad (1)$$

The function \tan atan is sometimes called arc tan.

The right hand side of equation 2 is an approximate model of the curvature of the footprint. The left hand side of equation 2 ignores the front half of the hole, so it is also an approximation.

The probability of sinking the putt can be expressed as the integral of the joint probability density of (A, R) over the footprint of success. This integral is

$$\int_{-t}^t \int_s^{(s+g)} q(a, r) dr da, \quad (2)$$

where $t = \tan(d/s)$, $g = \left[1 - \left| \frac{A}{\tan(d/s)} \right|^2 \right] b$, and $q(a, r)$ is the Gaussian probability density of (A, R) .

4 Residual distance of a missed putt

Figure 3 shows the geometry needed to compute the residual distance to the hole if a putt is missed. In the figure, the residual is labeled e . The equations needed to solve for e are:

$$e^2 = w^2 + h^2 \quad (3)$$

$$\sin(A) = h/R \quad (4)$$

$$\cos(A) = (s+w)/R. \quad (5)$$

Simple algebra yields

$$e = \{ [R \cdot \cos(A) - s]^2 + [R \cdot \sin(A)]^2 \}^{1/2}. \quad (6)$$

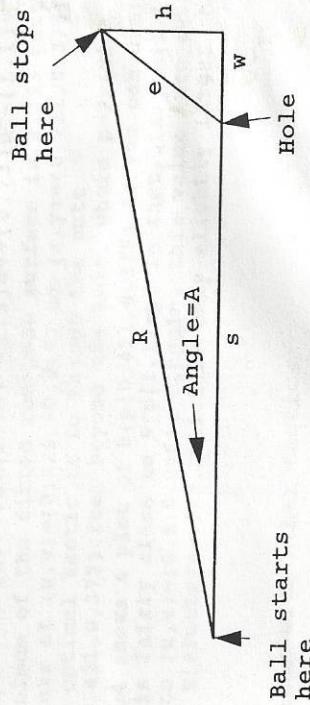


Fig.3. Residual distance of a putt.

5 Putting tactics

In the introductory Section 1, I argued that the target for the putt should be of the form

$$\mu = s + (2') \cdot B(p), \quad (7)$$

where $B(\cdot)$ is a monotonically increasing function, with $B(0)=0$ and $B(1)=1$. This can be proved rigorously using dynamic programming and Markov decision theory. Space does not permit that here.

To find an approximation to the optimal $B(\cdot)$, I consider a two parameter family of such functions based on the cumulative beta distribution.

The generic member of the class is

$$B(p|x,y) = \int_0^p \frac{\Gamma(x+y)}{\Gamma(x)\Gamma(y)} v^{x-1} (1-v)^{y-1} dv. \quad (8)$$

This is a fairly general class of functions with lots of different shapes - convex, concave, s-shaped, and backwards s-shaped. The parameters x and y are now the variables to be optimized.

6 Optimization methodology

The optimization methodology proceeds as follows:

Step 1. Pick an initial position on the green for which the probability of sinking the putt is near 0.

Step 2. Guess at the optimal values of x and y to define putting tactics.

Step 3. Estimate via simulation the expected number of putts (see Section 7) needed to hole out.

Step 4. Pick nearby values of x and y and repeat Step 3. This could be done using response surface methodology, but I did not.

Step 5. Continue to sequentially search for better and better values of x and y based on the sequence of values generated by Step 3.

Step 6. Stop when there is enough data to fit a parametric response surface. Then find the minimum of the response surface to get an approximate optimal putting tactic.

7 Expected number of putts

For a given initial position on the green and a given putting tactic, the expected number of putts needed to hole out is estimated via repeated Monte Carlo simulation of a single sojourn on the green.

The simulation of a single putting sojourn goes as follows: The initial position of the ball and the putting tactic (x,y) are given.

Step 1. Target the first putt according to equations 8 and 9.

Step 2. Compute the probability p of sinking the putt using equation 3 and numerical integration.

- Step 3. Pick a random number U between 0 and 1. If $U < p$, then the putt is sunk and you record the number of putts taken. Otherwise, continue.
- Step 4. Compute the residual distance to the hole using equation 7.
- Step 5. Redo Steps 2-4 until you hole out.

8 Simulation results

Consider an example where the distance (s) to the hole starts out at 60', the standard deviation (α) of the putting angle (A) is 1.44, and the standard deviation (σ) of the putting radius (R) is 0.22' (s); i.e., $k=0.22$. The parameters $(\alpha,k)=(1.44,0.22)$ were chosen by fitting my theoretical putting model to the lower curve on p. 38 of Pelz (1989). This curve gives the probability of sinking a putt as a function of distance. The lower curve is for the worst putters on the PGA tour. My fit to Pelz's empirical curve is very good, which lends credibility to my model. Space does not permit a demonstration of this.

Following the steps described in Section 7, I simulated many tactics of the form (x,y) . For every pair (x,y) , the simulation produced an estimated expected number of putts along with its standard deviation. Different numbers of iterations were used for the various runs. For example, for $(x,y)=(3,1)$, the number of iterations was 10,000, and the expected number of puts was 2.781 with standard deviation 0.0067.

An empirical plot of these results suggested that (x,y) be transformed before attempting to fit a simple quadratic response surface. The transformation selected was

$$u=x/(x+y) \quad (9)$$

$$v=[\ln(x+y)]/15. \quad (10)$$

The fitted response surface is

$$E[\#Putts] = 3.1 + (0.74)u + (1.1)v - (1.3)uv + (0.73)u^2 - (1.1)v^2. \quad (11)$$

The minimum of the fitted response surface is 2.753 putts, which occurs at $(u,v)=(0.529,-0.015)$ or $(x,y)=(0.421,0.375)$. So the optimal tactic is to target the putt $(2)*B(p|0.421,0.375)$ feet beyond the hole, where $p=pr[\text{Sink Putt}]$.

Figure 4 shows a plot of $B(p|0.421,0.375)$. As you can see, this

function is fairly close to $B(p|1,1)=p$. In fact, $(x,y)=(1,1)$ corresponds to $(u,v)=(0.5,0.046)$. Plugging this value into equation 2.753.

9 Comparison with other tactics

It is of interest to compare my optimal tactic with other proposed tactics. For example, Pelz (1989) on p. 127 says to target 17" beyond the hole for all putts. A simulation of Pelz's tactic with 5,000 iterations yielded an estimated expected number of puts of 2.81 with a standard deviation of 0.01. So my optimal tactic is 0.06 putts per hole better (for putting situations, which start at 60' from the hole). That is $(0.06)*(18)=1.08$ putts per round.

Of course, to get a more personalized result, you should average over your starting positions.

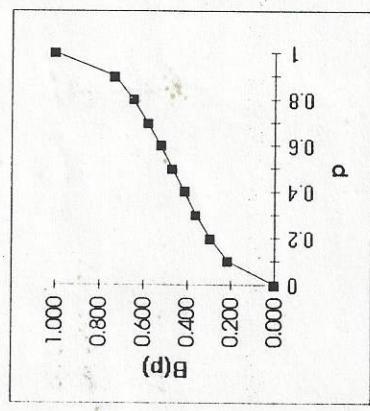


Fig. 4. The function $B(p|0.421, 0.375)$.

As a final benchmark, consider the tactic of always targeting the hole itself - not beyond the hole. A simulation of 10,000 iterations yielded a result of 3.059 putts per hole with a standard deviation of 0.0086. So targeting the hole itself costs you 0.3 putts per 60' hole.

10 References

Pelz, Dave with Mastroni, Nick (1989) *Putt Like the Pros*. Harper-Collins Publications, New York, NY.