

The Optimal Stopping Problem

Luce Skrabanek

24 September, 2024

1 Motivation

It is very useful in science to construct mathematical models of the systems that we are investigating. As the complexity of these systems increases, it can often be difficult or impractical to derive or locate in the literature exact mathematical solutions of our model systems. While one option is to incorporate simplifying assumptions (such as that of normality) that make analytical solutions tractable, an alternative approach is to more directly simulate real world phenomena, and use numerical techniques to obtain solutions and insights.

For example, we saw how a t-test (which uses an analytical solution based on the Student's t-distribution) can also be cast as a data shuffling problem, and solved by repeated numerical (random or exhaustive) sampling of such shufflings via the randomization test. Similarly, we've seen how simulations could be used to explore power calculations for two groups with different means and different standard deviations, or to illustrate how the Bonferroni and FDR multiple hypothesis corrections are controlling the error for our entire experiment.

In this lecture, we'll develop the tools that you can use to build and explore your own models.

2 The Optimal Stopping Problem

Algorithms to Live By, by Brian Christian and Tom Griffiths, explores how people solve all sorts of problems that are defined by a limited amount of time, space, or information or some combination of the above. The first chapter describes the so-called “secretary problem”, also called the “optimal stopping problem”. Although its origins are obscured by the mists of history, it was first described in print by Martin Gardner in his famous

Mathematical Games column in a 1960 issue of *Scientific American*. We'll frame it as a "dating problem".

2.1 Define the problem

Let's first define the problem: We want to maximize our chance of finding our soul mate in a limited pool of candidates. Each potential mate in the pool has a compatibility score, and our goal is to find (and then propose to) the person in the pool with the highest score for us. We can only ascertain a person's compatibility score by dating them, and once we move on to the next candidate, we can never go back to someone we've dated in the past. So we need to decide if we want to propose to someone without knowing for sure if there is a better match still to come.

We'll attack this problem using a so-called "look then leap" strategy, whereby we'll first date a predetermined number of people just to gather data about the range of compatibility scores. Once our data gathering phase is over, we'll then continue to date the remaining people from our pool until we find someone whose compatibility score is higher than any we've seen in the data gathering phase, and commit to that person. In this model, if we propose to somebody, they always accept, and if we reject them, we can never get them back.

The question now becomes: how many potential mates should be included in the data gathering phase, and how many should be left for the commitment phase?

In this model, there are two ways we can fail to meet our soul mate using this strategy: a) we meet our soul mate during the data gathering phase, and end up moving on to date others when we shouldn't have, or b) we can commit too early in the commitment phase, and never meet our soul mate. Note that in this model, success is binary and only occurs if we commit to our optimal match; committing to the second-best match is considered a failure, regardless of their numerical compatibility score. We'll call this the "black-and-white" approach. The optimal strategy requires finding the right balance between the two phases, or deriving the optimal proportion of the population that we should sample before committing.

2.2 A first model: "black-and-white"

It turns out that this problem does have a well-known analytical solution; the math (which in its exact form is a Riemann approximation to an integral) says that the optimal proportion of people we gather data on before switching to the commitment phase is $1/e$ (cf Thomas Ferguson, Who Solved the Secretary Problem? *Statistical Science*, 4(3):282-289 (1989)). Let's see if we can use simulations to verify this result.

We first develop the code for one scenario. Say we have 100 people in our potential candidate dating space, each with a score representing how good a match they are. The person with the best score is our real soul mate.

```
population <- 100
score <- rnorm(population)
optimal_score <- max(score)
```

We date the first *phase1_size* of the group (note that we are rounding to the nearest whole number) and note the maximum score in that group...

```
phase1_fraction <- 1/exp(1)
phase1_size <- round(population * phase1_fraction)
cutoff_score <- max(score[1:phase1_size])
```

Now select as your life partner the next date with a better score.

```
spouse_index <- phase1_size +
  which(score[(phase1_size+1):population] > cutoff_score)[1]
spouse_score <- score[spouse_index]
```

Note that in the above code snippet, the `which()` function returns the indices of the truncated array `score[(phase1_size+1):population]`. So, if your soul mate is found at position 1 in the truncated array, that will correspond to position `phase1_size + 1` in the whole array; adding `phase1_size` to the result of the `which()` function accounts for this shift.

We have successfully found our soul mate if the score of the match we found is the same as the optimal score.

```
spouse_score == optimal_score
```

If you run this a few times, you'll find that when our soul mate is encountered in the data gathering phase, we don't commit to anyone at all and the `spouse_score` is `NA`, which becomes an `NA` in our logical expression. We really want this to be recorded as `FALSE`, so let's set the score to the score of the final candidate if we haven't found anybody else by then.

```
if (is.na(spouse_index)) {
  spouse_index <- population
}
```

We probably want to run this multiple times, so that we can see how many times, on average, we find our soul mate. To do this, we'll write a for loop to run the code 1000 times.

```
num_iterations <- 1000
# vector to store result of each simulation
is_soulmate <- logical(length = num_iterations)
# Explicitly state the size of the total, and sampling phase, pools
population <- 100
phase1_fraction <- 1/exp(1)
phase1_size <- round(population * phase1_fraction)

for (case_idx in 1:num_iterations) {
  # scores of potential mates
  score <- rnorm(population)
  # score of soul mate
  optimal_score <- max(score)

  # maximum score in gathering phase
  cutoff_score <- max(score[1:phase1_size])

  # now select as your life partner the next date with a better score
  spouse_index <- phase1_size +
    which(score[(phase1_size+1):population] > cutoff_score)[1]

  # pick the last one if nobody better came along before then
  if (is.na(spouse_index)) {
    spouse_index <- population
  }

  is_soulmate[case_idx] <- (score[spouse_index] == optimal_score)
}
```

and calculate the mean of all the results (where the mean is just the sum of all the TRUEs, divided by the total number of trials):

```
mean(is_soulmate)
## [1] 0.381
```

We can see that the mean turns out to be approximately 0.37, or we find our soul mate 37% of the time, if we sample $1/e$ of the population before committing to somebody... which is coincidentally also the proportion of the sample we date before committing.

We probably want to explore this simulation a bit. What happens if we change the fraction of the population that we sample before switching to the commitment phase? Does the

size of the candidate pool make a difference? We can answer questions such as these by casting our code as a function, and then varying the parameters it is called with.

```
simulate_dating <- function(population = 1000,
                           phase1_fraction = 1 / exp(1),
                           num_iterations = 1) {
  is_soulmate <- logical(length = num_iterations)
  phase1_size <- round(population * phase1_fraction)

  for (case_idx in 1:num_iterations) {
    # scores of potential mates
    score <- rnorm(population)
    optimal_score <- max(score)

    # we date the first phase1_size people
    # and note the maximum score in that group
    cutoff_score <- max(score[1:phase1_size])

    # now select as your life partner the next date with a better score
    spouse_index <- phase1_size +
      which(score[(phase1_size+1):population] > cutoff_score)[1]

    # pick the last one if nobody better came along before then
    if (is.na(spouse_index)) {
      spouse_index <- population
    }

    is_soulmate[case_idx] <- (score[spouse_index] == optimal_score)
  }


  mean(is_soulmate)
}
```

Note here that we have assigned default values to all of the function's arguments. This way, when using the function, we only have to specify the arguments that differ.

Now we can call this function inside of another for loop.

```
phase1_fractions <- seq(0.05, 0.95, 0.02)
means <- numeric(length = length(phase1_fractions))
for (idx in 1:length(phase1_fractions)) {
  means[idx] <- simulate_dating(population = 100,
```


```
        phase1_fraction = phase1_fractions[idx],  
        num_iterations = 1000)  
}  
plot(means ~ phase1_fractions)
```



```
phase1_fractions[which.max(means)]  
## [1] 0.41
```

There is definitely a pattern, but the graph is a bit bumpy, so let's recompute this distribution with more iterations to get a smoother curve.

```
phase1_fractions <- seq(0.05, 0.95, 0.02)  
means <- numeric(length = length(phase1_fractions))  
for (idx in 1:length(phase1_fractions)) {  
  means[idx] <- simulate_dating(population = 100,  
                                phase1_fraction = phase1_fractions[idx],  
                                num_iterations = 10000)  
}  
plot(means ~ phase1_fractions)
```



```
phase1_fractions[which.max(means)]  
## [1] 0.33
```

The `phase1_fraction` that corresponds to the peak of the graph is around 0.30-0.40..., confirming the analytical solution of $1/e!$

The optimal strategy gives us a 37% chance of finding our soul mate. However, this also implies a 63% failure rate (given our current definition of success)!

2.3 Exploring an alternative model: “shades-of-gray”

Let’s consider an alternative model with a less stringent definition of “success”. Instead of defining success as finding our soul mate, and pairing up with anybody else as failure, we might consider the score of the person that we ultimately pair off with as a numerical measure of the success of our endeavor. In other words, committing to somebody with a relatively high score is still considered a pretty good outcome, even if it wasn’t the optimal one. So with this approach, the metric that we seek to optimize is not the probability of matching with one’s soulmate, but rather to maximize the average expected score of the person we match with. We’ll refer to this as the “shades-of-gray” approach.

While the underlying math for the analytical solution for this optimization is perhaps very different, modifying our simulation to reflect this different objective is straightforward.

```
# Numeric vector to store score of match  
spouse_scores <- numeric(length = num_iterations)  
# Record compatibility score of the matched person  
spouse_scores[case_idx] <- score[spouse_index]
```

Note that `spouse_scores` now holds the numeric compatibility value, not a TRUE/FALSE.

Here, our goal is to find the proportion of the population that needs to be sampled to maximize the expected score of the person we match with. As the simulation below demonstrates, the optimal fraction of the population that we should sample before we switch to a commitment phase is now much lower!

```
simulate_dating_2 <- function(population = 1000,  
                             phase1_fraction = 1 / exp(1),  
                             num_iterations = 1) {  
  spouse_scores <- numeric(length = num_iterations)  
  phase1_size <- round(population * phase1_fraction)
```

```
for (case_idx in 1:num_iterations) {
  # scores of potential mates
  score <- rnorm(population)

  # we date the first phase1.size people
  # and note the maximum score in that group
  cutoff_score <- max(score[1:phase1_size])
  optimal_score <- max(score)


  # now select as your life partner the next date with a better score
  spouse_index <- phase1_size +
    which(score[(phase1_size+1):population] > cutoff_score)[1]

  if (is.na(spouse_index)) {
    spouse_index <- population
  }

  spouse_scores[case_idx] <- score[spouse_index]
}

mean(spouse_scores)
}

phase1_fractions <- seq(0.05, 0.95, 0.02)
means <- numeric(length = length(phase1_fractions))
for (idx in 1:length(phase1_fractions)) {
  means[idx] <- simulate_dating_2(population = 100,
                                phase1_fraction = phase1_fractions[idx],
                                num_iterations = 10000)
}
plot(means ~ phase1_fractions)
```

```
phase1_fractions[which.max(means)]
```

```
## [1] 0.15
```

So, depending on your preferred criterion for “success”, the optimal proportion of the population that is in the sampling phase will differ.

There are other riffs and improvements on our model that we could imagine. For example, a previous assumption was that once you reject somebody, their feelings are irreversibly hurt and you can never go back to them. However, with a few more lines of code, you could simulate the situation where you would be able to select any of the last three people (for example) that you saw.

3 Group Project

Part I

One of the advantages of using a simulation to model your system is that you have access to all the information about the system.

Modify the “shades-of-gray” simulation to return the vector of the numerical compatibility scores of the match from each iteration. Prepare a histogram of the scores obtained when using the optimal `phase1_fraction` value for the “shades-of-gray” approach (~15%), and a second histogram of the scores obtained when using the optimal `phase1_fraction` value for the “black-and-white” approach (~37%). Plot the two histograms side-by-side, with the same scales and binning.

Can you identify the sources of each peak in the bi-modal distributions?

Do these histograms give you any insight as to why the “shades-of-gray” approach has a shorter optimal gathering phase than the “black-and-white” approach? Does this new information influence which metric of “success” you might adopt or recommend?

Part II

The simulation we described in class assumes that the candidates come from a normal distribution, with a standard deviation of 1. For each of the two approaches, does the optimal stopping point change if the standard deviation is different? What about if the distribution is different, e.g., a uniform distribution? Can you reason out why?

Part III

One of the assumptions in the classic statement of this problem is that the candidate will always accept when you propose to them. How would you change the model to include the possibility of rejection? Do you think that a constant probability of rejection is a good model? It may be that if you find somebody very compatible, chances are good that they like you a lot too, and vice versa. How might you model the case where the chance of rejection is inversely proportional to how much you like them?

Part IV: Extra credit

Extend the model in any way you think might be fun or interesting! For example, a group of 100 people are all trying to get paired up. What would be the objective function? Total happiness? Or would you try to maximize individual’s happiness?

Logistics

Answers are due by 11:59 PM on Wednesday, 9th October, 2024. E-mail your solution to kaa4023@med.cornell.edu, jbanfelder@rockefeller.edu, and las2017@med.cornell.edu

This Problem Set must be completed in groups of two or three. Please submit only one write-up per group.

Your submission should be generated using R Markdown. Include the code needed to reproduce your results and generate any figures in your writeup, as well as the compiled HTML file.