

9.2: Searching Algorithms

Figure 9.2.1 shows a portion of the South Dakota Badlands, a barren landscape that includes many narrow ridges formed through erosion. Suppose you wish to climb to the highest point on this ridge. Because the shortest path to the summit is not obvious, you might adopt the following simple rule: look around you and take one step in the direction that has the greatest change in elevation, and then repeat until no further step is possible. The route you follow is the result of a systematic search that uses a searching algorithm. Of course there are as many possible routes as there are starting points, three examples of which are shown in Figure 9.2.1. Note that some routes do not reach the highest point—what we call the global optimum. Instead, many routes end at a local optimum from which further movement is impossible.




Figure 9.2.1. Finding the highest point on a ridge using a searching algorithm is one useful method for locating the optimum on a response surface. The path on the far right reaches the highest point, or the global optimum. The other two paths reach local optima. Searching algorithms have names: the one described here is the method of steepest ascent.

We can use a systematic searching algorithm to locate the optimum response. We begin by selecting an initial set of factor levels and measure the response. Next, we apply the rules of our searching algorithm to determine a new set of factor levels and measure its response, continuing this process until we reach an optimum response. Before we consider two common searching algorithms, let's consider how we evaluate a searching algorithm.

Effectiveness and Efficiency

A searching algorithm is characterized by its effectiveness and its efficiency. To be effective, a searching algorithm must find the response surface's global optimum, or at least reach a point near the global optimum. A searching algorithm may fail to find the global optimum for several reasons, including a poorly designed algorithm, uncertainty in measuring the response, and the presence of local optima. Let's consider each of these potential problems.

A poorly designed algorithm may prematurely end the search before it reaches the response surface's global optimum. As shown in Figure 9.2.2, when you climb a ridge that slopes up to the northeast, an algorithm is likely to fail if it limits your steps only to the north, south, east, or west. An algorithm that cannot respond to a change in the direction of steepest ascent is not an effective algorithm.




Figure 9.2.2. Example that shows how a poorly designed searching algorithm—limited to moving only north, south, east, or west—can fail to find a response surface’s global optimum.

All measurements contain uncertainty, or noise, that affects our ability to characterize the underlying signal. When the noise is greater than the local change in the signal, then a searching algorithm is likely to end before it reaches the global optimum. Figure 9.2.3, which provides a different view of Figure 9.2.1, shows us that the relatively flat terrain leading up to the ridge is heavily weathered and very uneven. Because the variation in local height (the noise) exceeds the slope (the signal), our searching algorithm ends the first time we step up onto a less weathered local surface that is higher than the immediately surrounding surfaces.




Figure 9.2.3. Another view of the ridge in Figure 9.2.1 that shows the weathered terrain leading up to the ridge. The yellow rod at the bottom of the figure, which marks the trail, is about 18 in high.

Finally, a response surface may contain several local optima, only one of which is the global optimum. If we begin the search near a local optimum, our searching algorithm may never reach the global optimum. The ridge in Figure 9.2.1, for example, has many peaks. Only those searches that begin at the far right will reach the highest point on the ridge. Ideally, a searching algorithm should reach the global optimum regardless of where it starts.

A searching algorithm always reaches an optimum. Our problem, of course, is that we do not know if it is the global optimum. One method for evaluating a searching algorithm’s effectiveness is to use several sets of initial factor levels, find the optimum response for each, and compare the results. If we arrive at or near the same optimum response after starting from very different locations on the response surface, then we are more confident that is it the global optimum.

Efficiency is a searching algorithm’s second desirable characteristic. An efficient algorithm moves from the initial set of factor levels to the optimum response in as few steps as possible. In seeking the highest point on the ridge in Figure 9.2.3, we can increase the rate at which we approach the optimum by taking larger steps. If the step size is too large, however, the difference between the experimental optimum and the true optimum may be unacceptably large. One solution is to adjust the step size during the search, using larger steps at the beginning and smaller steps as we approach the global optimum.

This page titled [9.2: Searching Algorithms](#) is shared under a [CC BY-NC-SA 4.0](#) license and was authored, remixed, and/or curated by [David Harvey](#).