


Simulation on Spatial Grids

Mathematical Programming with Python

MATH 2604: Advanced Scientific Computing 4
Spring 2025
Monday/Wednesday/Friday, 1:00-1:50pm

https://people.sc.fsu.edu/~jburkardt/classes/python_2025/grid_simulation/grid_simulation.pdf



A Mexican jumping bean makes a random walk.

Simulations on a grid

- *Some simulations involve changes in position;*
- *Some simulations consider influences from neighboring points;*
- *A cellular automaton is an abstract “machine” whose state is a string of 0’s and 1’s, changing over time;*
- *Forest fires and diseases can be modeled as a spread to neighbors;*
- *Percolation is a model of a loosely packed material which might allow or block the flow of a liquid ;*

1 A cellular automaton

Steven Wolfram, the developer of *Mathematica* has studied a class of 256 cellular automata. We will think a cellular automaton as a rule for filling in values in an $m \times n$, given an initial row of 0’s and 1’s, and a specification for how to use the current row to fill in the next row.

The value to be assigned to cell (i,j) will only depend on three values in the previous row, in positions $(i-1,j-1)$, $(i-1,j)$ and $(i-1,j+1)$. We also need a rule for what to do with the first and last cells in each row: (assume an extra 0? or perhaps use wrap-around?)

Wolfram found that some rules quickly produced an entire field of 0’s or 1’s, or made uninteresting patterns. But one rule, number 30, seemed to produce a wide range of results depending on the initial data. In fact, he found a way to turn the output of this rule into a random number generator.

The rule is summarized by listing the three values in the preceding row, and their result:

index	neighbors	result
0	0,0,0	0
1	0,0,1	1
2	0,1,0	1
3	0,1,1	1
4	1,0,0	1
5	1,0,1	0
6	1,1,0	0
7	1,1,1	0

Note that 00011110 is the binary representation for 30, which is how the rule got its name.

We can experiment with this rule by filling in a sample array using rule 30. We will assume that the cells on the left and right have an extra “ghost” zero when we are looking for neighbors.

```
0: 0 0 0 1 0 0 0
1: - - - - -
2: - - - - -
3: - - - - -
4: - - - - -
5: - - - - -
6: - - - - -
7: - - - - -
8: - - - - -
9: - - - - -
```

Now we can see how to write a program that will fill in the grid associated with a cellular automaton if we have the rules, and the values in the first row. Assuming we initialized the grid to zero, then we only have to decide whether to reset a cell value to 1.

Input:

```
m, n: the number of rows and columns
start[n]: initial value for row 0
```

Initialize


```
c = np.zeros ( [ m, n ] )
c[0,:] = start[:]
```

Compute:

```
For 1 <= i < m
  for 0 <= j < n
    jm1 = ( j - 1 ) % n
    jp1 = ( j + 1 ) % n
    if ( c[i,jm1] == 0 and c[i,j] == 0 and c[i,jp1] = 1 ) or # 001 -> 1
        ( c[i,jm1] == 0 and c[i,j] == 1 and c[i,jp1] = 0 ) or # 010 -> 1
        ( c[i,jm1] == 0 and c[i,j] == 1 and c[i,jp1] = 1 ) or # 011 -> 1
        ( c[i,jm1] == 1 and c[i,j] == 0 and c[i,jp1] = 0 ) # 100 -> 1
      c[i,j] = 1
```

Return c

We can plot the cellular automaton easily. Here, plotting with i and j , I got an upside down plot. Defining $i2=step_num-1-i$ fixed the discrepancy between matrix indices and Cartesian coordinates:



2 Random walks

A random walk seems like a peculiar thing to study. However, it models a real physical situation, “Brownian motion”, noticed by Robert Brown in 1827, watching grains of pollen suspended in water, that randomly jiggled.

The puzzle wasn’t fully solved until Albert Einstein explained this by the collision of the pollen with individual water molecules whose velocities had a random variation.

From this, Einstein was able to show that a particle in Brownian motion would tend to drift away from its original position with a predictable variation.

3 Staggering Home

Bellefield Robotics has designed a food truck robot named ”BumbleBee” which offers food for sale at the corners of street blocks along Fifth Avenue. Its home station is at Bellefield Avenue. BumbleBee simply wanders up and down Fifth Avenue, from block to block, looking for customers. As a safety feature, it shuts down if it reaches a distance of 10 blocks from home base.




On Friday morning, BumbleBee starts at Bellefield, and chooses a random direction, east or west, to go to the next block. Today there are no customers, so after waiting at one block for five minutes, Bumblebee again randomly moves east or west to another block.

We may reasonably assume that at some point in the day, BumbleBee will randomly reach a distance of 10 blocks, and shut down. On average, when will this happen? Clearly, it takes a minimum of $10 * 5$ minutes to get that far from home base, and obviously, there is theoretically no maximum amount of time that could be involved. But realistically, would we have to wait two hours, two days, or what?

We can easily simulate one instance of this walk:

```
def walk_1d ( ):  
    directions = [ -1, +1 ]  
    location = 0  
    blocks = 0  
  
    while ( -10 < location and location < 10 ):  
        blocks = blocks + 1  
        location = location + np.random.choice ( directions )  
  
    return blocks
```

If we simulate BumbleBee's adventure five times, we see that walks of length 220, 42, 32, 118 or 48 blocks, to reach a destination that is only 10 blocks away. This is a little irregular and mystifying. The first thing we need to do is to generate lots of samples of this walk and find the average time it takes to reach a distance of 1 block, 2 blocks, ..., 10 blocks.




Surprisingly, the average data is very regular, and suggest that to reach d blocks in a random fashion, we actually typically have to walk d^2 blocks. Thus, on average, BumbleBee will shut down after $5 * 100$ minutes, or a bit more than 8 hours after starting.

4 Mexican Jumping Beans

Random walks occur in real life. Consider the case of the Mexican jumping bean. A moth lays eggs inside a bean pod, the bean pod later falls to the ground, the egg hatches as a larva, but is still in the bean pod, eating and developing. Some of these pods are in direct sunlight, and the larva will die if it overheats. The larva can try to move the bean pod by making sudden jerking motions. The bean pod flips around in response. If it lands in a shaded spot, the larva can sense the cooler temperature, and stop the jumping.

See the report *Mexican jumping beans use random walk strategy to find shade* available at <https://arstechnica.com/science/2023/02/taking-a-walk-on-the-random-side-helps-mexican-jumping-beans-find-shade/> which includes a link to a short video.



If you find this example interesting, you can try to simulate a very simple version of it. Suppose a jumping bean starts at the origin $(0,0)$, and randomly chooses steps north, south, east or west of one unit. On average, how many steps must be taken before the bean is 0, 1, 2, ..., 10 units away from the origin? Measure distance by $d = \sqrt{i^2 + j^2}$. Do about 100 simulations for each distance measurement. Make a plot of the average number of steps as a function of d , something like the plot we made for the food delivery robot.

If you can make a nice program and plot for this problem, you can turn these in instead of doing assignment #6.

5 The coughing passenger

Suppose you are a passenger on a plane, and you hear someone two rows ahead of you who is coughing repeatedly. You might begin to worry, because an airborne disease is most likely to spread to nearby neighbors of an infected person. If it was a really long flight, and the disease developed really rapidly, then it might spread first from the coughing passenger to the passenger in front of you, and then to you, taking two steps of infection.

In a simple model of disease propagation over time and space, we might have an $m \times n$ grid of people, who have agreed not to move during the extent of our experiment. We start with one infected person, at some random position (i, j) . According to our simplified disease model, the other people are “susceptible”, that is, they aren’t sick, but they could become so. However, a susceptible person can only catch the disease if they are sitting next to an infected person (left, right, in front or in back). A susceptible person who sits next to an infected person for an hour has a probability p of getting infected during that time.

Now given what we have said so far, it seems certain that if we wait long enough, everyone in the group must get infected; the one infected person we start with must eventually infect all their immediate neighbors, who must eventually infect all their neighbors, and so on.

This is not a very realistic model of disease. So let’s add the idea that an infection only lasts for h hours, after which the person is no longer sick, and in fact becomes immune. An immune person cannot get sick again.

So our model starts with three types of people: $m \times n - 1$ susceptibles (S), 1 infective (I), and 0 recovered (R). Over time, we expect the value of I to grow, but then it must eventually decrease to zero as the disease “burns out”. The value of interest, then, is the number of people who became sick versus those who never

got the disease. So we are interested in the value, at the end of the experiment, for the fractions $\frac{S}{m*n}$ and $\frac{R}{m*n}$. If everyone got sick, then these values will be 0 and 1, respectively, so the disease swept through the entire population. Higher values of $\frac{S}{m*n}$ indicate that the disease was less effective in spreading.

Using this model, and the value of $\frac{S}{m*n}$, we could also experiment with small changes to the problem, such as the location of the initial infected person, or the use of a vaccine to reduce the value of p , the transmissibility of the disease, or a treatment that reduces h , the duration of the disease, or some means of isolating infected persons so they can't transmit the disease to others.

Here is an outline of an algorithm for the SIR model, assuming we are given values for m, n, p, h . The current status is stored in the array P , where $P(i, j) =$

- 1, if person (i,j) is susceptible;
- -k, if person (i,j) is infected for k more hours;
- 0, if person (i,j) is recovered;

We continually updated the array P hour by hour until there are no infected persons:

```

P = ones ( m, n )
i, j = indices of a random location in P.
P(i,j) = -h # negative value; person will recover in h hours
hour = 0

while any P < 0

    hour = hour + 1

    for (i,j) in the grid

        Q = zeros ( m, n )
        if ( P(i,j) == 1 )
            for north, south, east, west infected neighbor, infect P(i,j) with probability p
                if infected, Q(i,j) = -h
            else if ( P(i,j) < 0 )
                Q(i,j) = P(i,j) + 1
                infected person has one less hour to be infected
            end
        end

    end


    replace P by value of Q

end

print hour, duration of epidemic
print sum ( P ) / m / n, proportion of population that did not get sick

```

Here is a plot of the variation in the values of S, I, R for a model in which at the end of the simulation, S = 0.14 and R = 0.86:



The behavior of S , I , and R for $m = n = 10, p = 0.2, h = 4$.

6 A Forest Fire

Let's consider a simple forest fire model. We assume an $m \times n$ grid. In each grid cell we have a measurement of the height of a single tree, which can be 0, 1, 2, 3 or 4 feet high. (Zero means no tree, of course.) Lightning strikes a random cell setting a serious fire there. The behavior of the fire is as follows:


- if a tree is on fire, then on the next time step, it is one foot shorter.
- if a tree is not on fire, it might catch fire on the next step. It has a probability of catching fire that is $1/2 * \text{number of burning neighbors}$.

A neighbor tree is one to the immediate north, south, east or west.

To simplify matters, we also use wrap around. So if `tree(i,j)` is in row $i=m-1$, its "south" neighbor is `tree(0,j)`, and so on.

Here is a plot suggesting the initial situation in the forest. The darkest green spots are the tallest trees, and the red spot is where the fire starts.

Forest Fire



Here is how we might implement the status of the forest over one step.

```
def forest_update ( forest ):
    import numpy as np
    m, n = forest.shape
    new_forest = forest.copy ( )
    for i in range ( 0, m ):
        for j in range ( 0, n ):
            #
            # Burning tree gets 1 step closer to zero:
            #
            if ( forest[i,j] < 0 ):
                new_forest[i,j] = forest[i,j] + 1
            #
            # Nonburning tree with burning neighbors might catch fire.
            #
            elif ( 0 < forest[i,j] ):
                im1 = ( i - 1 ) % m
                ip1 = ( i + 1 ) % m
                jm1 = ( j - 1 ) % n
                jp1 = ( j + 1 ) % n
                s = ( forest[im1,j] < 0 ) + ( forest[ip1,j] < 0 ) \
                    + ( forest[i,jm1] < 0 ) + ( forest[i,jp1] < 0 )
                ignite = np.random.random ( )
                if ( ignite < 0.50 * s ):
                    new_forest[i,j] = - new_forest[i,j]
```


```

forest = new_forest.copy ( )
return forest

```

If we keep following the updating rules, then the fire will eventually die out. We can ask questions such as how many trees are likely to survive, and what would happen if the direction of the wind influenced the fire spread.

Here is what the forest might look like after 10 steps of burning:



7 Percolation

Oil drillers are familiar with peculiar structures of underground rock. Often, they are drilling through a material that is a jumbled combination of rock and hollowed-out cavities which might contain oil, gas, or water. Depending on the size, density, and frequency of these cavities, a particular underground “lake” of oil might extend for a few hundred feet, or for miles.


As another example, we might consider an object composed of two metals, one conducting and one not. We suppose the object involves many separate regions of each metal, which are touching, but not melted together. We know the approximate size of a typical “cell” or clump of each metal, and the proportion of the two metals. We ask for the probability that an electrical signal can pass from the top to the bottom of the object. This will happen only if there is a some path, formed entirely of the conducting metal, that reaches from top to bottom.

John Hammersley proposed an approach for studying such problems. He imagined a porous stone such as pumice, which is made from volcanic activity. Suppose such a stone is submerged in a bucket of water. Would the center of the stone end up wet or dry? This depends on whether there is at least one microscopic channel between the surface and the center. If the water reaches the center, we say it has done so by *percolation*. This is the same word used to describe one method of making coffee, by having boiling water percolate through a layer of coffee grounds.

A simple 2d model of percolation can be constructed by setting up an $m \times n$ grid of square cells. We will allow some cells to be obstructions (filled, or nonconducting, or solid) and others to be channels (open, conducting, or hollow). We will be looking for the existence of a path from top to bottom. To be a realistic model, the number and size of these cells would have to approximate those of a physical structure. For our learning purposes, though, we will be satisfied with rather coarse grids. We will look at several factors to vary in our problems:

- the total number of cells $m \times n$;
- the aspect ratio $\frac{m}{n}$;
- the probability p that a given cell is conducting;

Once we have chosen m, n, p , we can randomly set the status of each grid cell. The hard part now is to figure out whether the resulting pattern of conducting cells forms a path from top to bottom. If we plot the grid, we can usually spot a path if it exists, at least for our small example grids. We have to come up with a way for the computer to do the same thing, without the benefit of eyes.



The raw occupation grid, and the clustering pattern. $m = n = 32, p = 0.60$.

In the illustration on the left, we can see that a path exists. On the right, each connected “cluster” of cells is given its own color, and we can observe that there is an enormous group of blue cells which do indeed include a path from top to bottom. The idea of clustering will be the key to helping the computer to identify such connecting paths.

An interesting feature of these mathematical percolation simulations is that, if we fix the values of m and n , then as we vary p we observe a fairly sudden transition in behavior. Below a critical value of p , we almost never get a connecting path, which above it, such a path almost always exists. By analogy with how ice changes to water, mathematicians describe this behavior as a *phase transition*. A variety of other numerical problems have this behavior, include something known as the *subset sum* problem, in which a set of integers must be divided into two sets of equal sum.

For our example with $m = n = 32$, the phase transition seems to happen at about $p = 0.60$. The example program `percolation.py` can be used to investigate this phenomenon further.