Decision Making Based on Information Gain

F B

August 10, 2012

1 Explanation

There is a $n \times n$ grid. Each cell in the grid can eitgher hold a real number in the range [0,1) or be empty. There also exist a sliding window of size $m \times m$ (m < n), which represents the operation range of the measuring device. figure 1 illustrates the idea. At each time instance t, the sliding window is placed somewhere over the grid and measurements are read for each of the cells in grid covered by window. Because of the nature of our settings and devices, each reading can return one of the following two:

- it might return Null, indicating the cell is empty (with 100% curtainty)
- it might return a real number. In this case our observation will be the true number contained in the cell plus a Gaussian noise. $z_{observed} = z_{real} + n$ where $n \in N(0, \sigma^2)$. This holds true for future measurments of this cell too, meaning that the reading will be within the gaussian noise range of the true value, so Kalman filter can be used to narrow down the readings.

It's worth mentioning that the Gaussian noise is caused by physical charectrisitics of the measuring device, it's zero mean and has fixed variance. Also please note that the sliding windonw can move freely, and its movement is not limited to its current adjacent cells.

Our objective is to find the best policy (best next move) for moving the sliding window, so that new measurement gives maximum amount of information.

2 Proposed Solution

We use relative/differential entroy as a measure of information. As a prior we assume each cell contains a real number with uniform distribution in the range [0,2]. The probability sums up to one-1 meaning there is definitely a number in that range;
$\int_0^2 p(x)dx = 1$ where $p(x) = 0.5$ in here.
Continuous entropy for that distribution will be:
$H_1 = \int_0^2 p(x)\log(p(x))dx = \int_0^2 0.5 \times (-1)dx = -1$.
We also assume that without any other knowledge, there is 50% chance that a cell is occupied (this agrees with our experiments). So after observation, if we find out the cell is empty (no real number between [0,1)) then we can say there is a uniform possibility that number is in imaginary range $[1, 2)$. Calculating relative entropy tells us that this particular observation has earned us:

Figure 1: Illustration of the grid and sliding window, cells that have not yet been observed are left blank

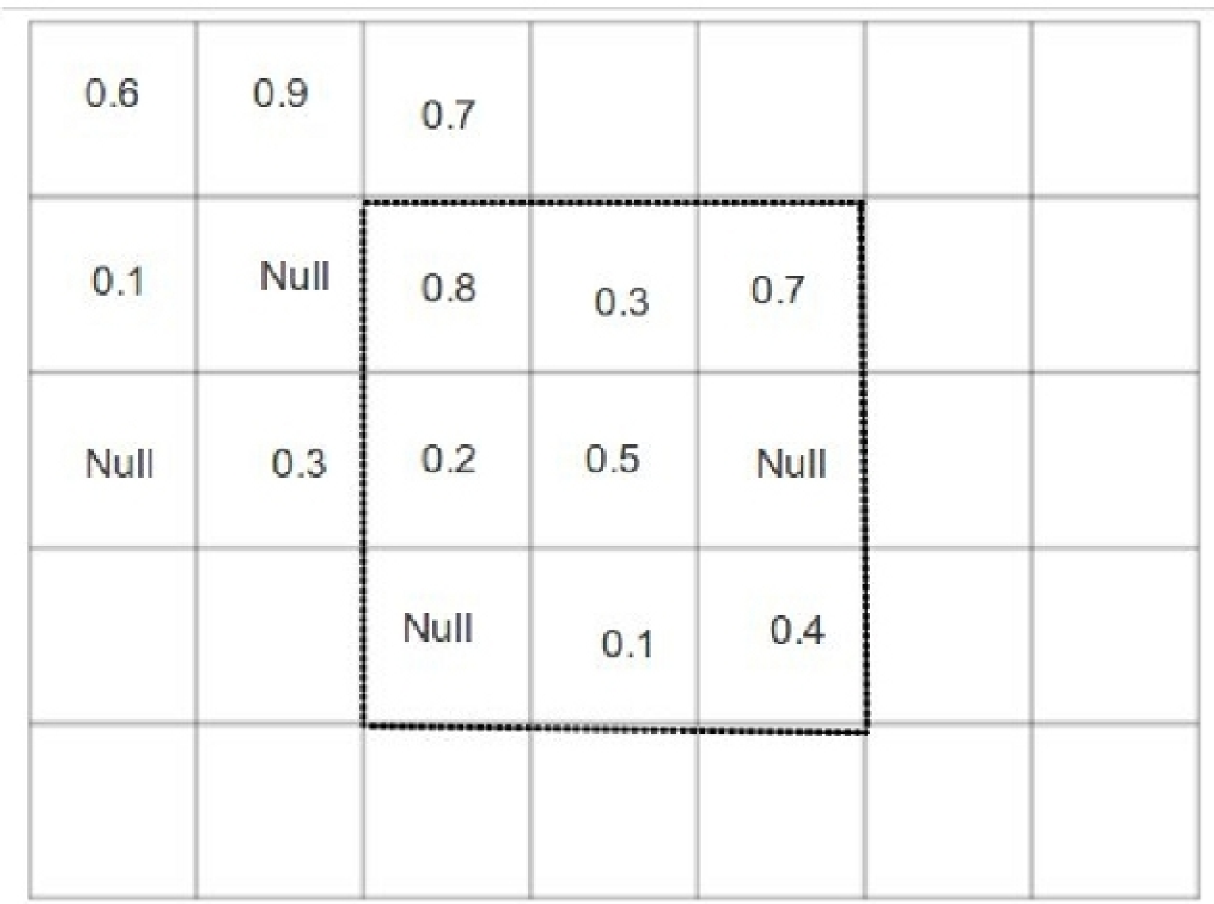

$H_2 - H_1 = \int_1^2 1 \times log(1)dx - \int_0^2 0.5 \times log(0.5)dx = 1$ bit of information.
In case reading returns a number, reletive entropy will be even higher, namely:
$\int_0^2 p(x) \times log(p(x))dx \approx \int_{-\infty}^{\infty} p(x) \times log(p(x))dx$ because: $sigma \ll 1$
$H_3 - H_1 = \frac{1}{2}log(2\pi e\sigma^2) - (-1)$

Second scenario is if a cell has already been observed. Now if we know cell is empty, new measurements won't help us and information gain will be zero. But if from previous readngs we have a gaussian distribution showing our belief in the reading, new reading will increase the curtainty and hence give some information. This extra information gain can be calculated as follows;
current variance $= \sigma_c^2$
variance of measurement $= \sigma_m^2$
variance or result after filtering $= \sigma_r^2 = \frac{\sigma_c^2 \times \sigma_m^2}{\sigma_c^2 + \sigma_m^2}$
information gained: $H_2 - H_1 = \frac{1}{2}log(2\pi e\sigma_r^2) - \frac{1}{2}log(2\pi e\sigma_c^2)$

Now that we know how to calculate expected information gain in every scenario, best-next-move for the sliding window will be decided by doing an exhaustive search over all possible moves, and by choosing the move that returns maximum information (sum of information gain over all cells in the window). Please note that there is no spatial correlation between cells, and belief in a cell doesn't tell us anything about neighboring or other cells.