

THE POLYA URN PROCESS AND THE STOCHASTIC APPROXIMATION OF ITS BEHAVIOR

MICHAEL KANE

As a final project for STAT 637 (Deterministic and Stochastic Optimization) the Polya urn process was studied with special attention being paid to its relationship with stochastic approximation. This write-up is based on portions of Pemantle (2007), which provides a survey of random process but it also draws from other sources.

This paper begins by describing the Polya urn process. Next, distributions of ball counts and ball proportions are derived. These derivations should provide the reader with an understanding of how urn processes behave. After this intuition is developed, a stochastic approximation of the process is derived. This way, the relationship between the stochastic results and the previously derived results is clear. Finally, a case where the assumptions of stochastic approximation may not hold is introduced and an alternative solution is described.

1. THE POLYA URN

The Polya urn process is defined as:

```
Place  $r$  red balls and  $g$  green balls in the urn.  
repeat  
  Pick a ball  $l$  from the urn uniformly at random.  
  if  $l$  is a green ball then  
    Replace  $l$  and put another green ball into the urn.  
  end  
  else  
    Replace  $l$  and put another red ball into the urn.  
  end  
end ;
```

Algorithm 1: The Polya Urn Process

Clearly, the urn starts with $r + g$ balls. For the entire paper it will be assumed that r and g are greater than zero. At any time $t \geq 0$, the number

of balls in the urn is $t + r + g$. The next step is to understand the distribution of these balls at t .

2. COUNTING BALLS

At the first step of the process, there are $r + g$ balls in the urn and one ball is selected at random. This means that there is an $r/(r + g)$ chance of picking a red ball at the first step. Likewise, there is a $g/(r + g)$ chance of picking a green ball. Following Pollard (1997), let R_t be the event that a red ball is chosen at time t . Likewise, let G_t be the event that a green ball is chosen at time t . The probability of picking a red ball at time 2 can then be solved as:

$$\begin{aligned}\mathbb{P}R_2 &= \mathbb{P}R_1R_2 + \mathbb{P}G_1R_2 = \mathbb{P}R_1\mathbb{P}(R_2|R_1) + \mathbb{P}G_1\mathbb{P}(R_2|G_1) \\ &= \frac{r}{(r + g)} \frac{r + 1}{(r + g + 1)} + \frac{g}{(r + g)} \frac{r}{(r + g + 1)} \\ &= \frac{r}{r + g}.\end{aligned}$$

This shows $\mathbb{P}R_1 = \mathbb{P}R_2 = r/(r + g)$. These calculations work similarly for any $t \geq 1$, giving the result

$$\mathbb{P}R_t = \frac{r}{(r + g)}.$$

Now, let's examine the probabilities of a sequence of balls. For example, the probability of getting a green ball at time 1 and a red ball at time 2 is

$$\mathbb{P}G_1R_2 = \mathbb{P}G_1\mathbb{P}(R_2|G_1) = \frac{g}{(r + g)} \frac{r}{(r + g + 1)}.$$

The probability of getting a red ball at time 1 and a green ball at time 2 is

$$\mathbb{P}R_1G_2 = \mathbb{P}R_1\mathbb{P}(G_2|R_1) = \frac{r}{(r + g)} \frac{g}{(r + g + 1)}.$$

The probability of getting two green balls is

$$\mathbb{P}G_1G_2 = \mathbb{P}G_1\mathbb{P}(G_2|G_1) = \frac{g}{(r + g)} \frac{g + 1}{(r + g + 1)}.$$

Finally, the probability of getting two red balls is

$$\mathbb{P}R_1R_2 = \mathbb{P}R_1\mathbb{P}(R_2|R_1) = \frac{r}{(r + g)} \frac{r + 1}{(r + g + 1)}.$$

It may be noted that $\mathbb{P}R_1G_2 = \mathbb{P}G_1R_2$. In fact, the probability of any sequence of balls is only a function of the number of red balls and the number of green balls which appear in the sequence.

$$\mathbb{P}\{R_1 = \text{red}, G_2 = \text{green}, R_3 = \text{red}\} = \mathbb{P}\{R_1 + R_2 + R_3 = 2\}$$

In Panchenko (2007) it is shown that, by de Finetti Theorem, for an exchangeable sequence of Bernoulli random variables, there exists a distribution F on $[0, 1]$ such that

$$\mathbb{P}(R_1 + \dots + R_n = k) = \int_0^1 \binom{n}{k} x^k (1-x)^{n-k} dF(x).$$

To find the distribution of F , let's look at the moments.

We have already found that the first moment for R_k is $r/(r+g)$. The second moment was also found as

$$\frac{r(r+1)}{(r+g)(r+g+1)}.$$

Solving for the variance we get

$$\begin{aligned} \frac{r(r+1)}{(r+g)(r+g+1)} - \frac{r^2}{(r+g)^2} &= \frac{(r^2+r)(r+g) - r^2(r+g+1)}{(r+g)^2(r+g+1)} \\ &= \frac{rg}{(r+g)^2(r+g+1)} \end{aligned}$$

Which is the variance of the $\beta(r, g)$ distribution. The fact the F is a beta distribution can be confirmed by realizing that the moment generating function of the beta distribution are equal to the higher moments of this process.

Having k red balls after t steps can be expressed as

$$\begin{aligned} \mathbb{P}\{k \text{ red balls at time } t\} &= \int_0^1 \binom{t}{k} x^{k+r-1} (1-x)^{t-k+g-1} \frac{\Gamma(r+g)}{\Gamma(r)\Gamma(g)} dx \\ (2.1) \quad &= \binom{t}{k} \frac{\Gamma(r+g)}{\Gamma(g)\Gamma(r)} \frac{\Gamma(k+r)\Gamma(t-k+g)}{\Gamma(t+g+r)} \end{aligned}$$

3. THE DISTRIBUTION OF THE PROPORTION OF RED BALLS

Let X_{t+1} and Y_{t+1} be the number of red balls and green balls respectively at time $t+1$ and let W_{t+1} be the proportion of red balls at time $t+1$.

$$W_{t+1} = \frac{r + X_{t+1}}{X_{t+1} + Y_{t+1}} = \frac{r + X_{t+1}}{r + g + t + 1}$$

W_{t+1} may be written as a function of the random variable X_{t+1} and constants r, g , and t .

W_{t+1} is a martingale. Since W_{t+1} is a proportion we are guaranteed $\mathbb{P}|W_{t+1}| < \infty$. Also, if we let $\mathbb{P}_t W_{t+1}$ denote the probability of W_{t+1}

conditioned on all information up to and including t then,

$$\begin{aligned}\mathbb{P}_t W_{t+1} &= \frac{(X_t + 1)}{(X_t + Y_t + 1)} \frac{X_t}{X_t + Y_t} + \frac{X_t}{(X_t + Y_t)} \frac{Y_t}{(X_t + Y_t + 1)} \\ &= \frac{X_t(X_t + Y_t + 1)}{(X_t + Y_t)(X_t + Y_t + 1)} \\ &= \frac{X_t}{X_t + Y_t} \\ &= W_t.\end{aligned}$$

From Pollard (2001, page 151) we know that a submartingale $\{S_t : t \in \mathbb{N}_0\}$ with $\sup_t \mathbb{P} S_t^+ < \infty$ converges almost surely to an integrable limit. Since a martingale is also a submartingale, and $0 \leq W_{t+1} \leq 1$ for all $t \geq 0$, W_{t+1} must converge to an integrable limit.

To find this limit, first realize

$$W_t = \frac{r + X_t}{r + g + t} \rightarrow \frac{X_t}{t} \quad \text{as } t \rightarrow \infty.$$

From the last section, we know that the by de Finetti's theorem

$$\mathbb{P}\{X_t = k\} = \binom{t}{k} \int_0^1 x^k (1-x)^{t-k} d\beta(r, g)(x).$$

Again, using results from Panchenko (2007), for any function f with range $[0, 1]$,

$$\begin{aligned}\mathbb{P} f\left(\frac{X_t}{t}\right) &= \sum_{k=0}^t f\left(\frac{k}{t}\right) \binom{t}{k} \int_0^1 x^k (1-x)^{t-k} d\beta(r, g)(x) \\ &= \int_0^1 B_t(x) d\beta(r, g)(x).\end{aligned}$$

$B_t(x)$ is a *Bernstein Polynomial* which, in Wikipedia (2007), is shown to have the property

$$\lim_{t \rightarrow \infty} B_t(f)(x) = f(x).$$

Therefore,

$$\mathbb{P} f\left(\frac{X_t}{t}\right) \rightarrow \int_0^1 f(x) d\beta(r, g)(x) \quad \text{as } t \rightarrow \infty.$$

Any moment of X_t/t approaches the same moment of a $\beta(r, g)$ distribution for large enough t . The distribution of the proportion of red balls approaches a $\beta(r, g)$ almost surely as t approaches infinity.

4. EXTENSIONS TO OTHER URN MODELS

The previous sections study the *Polya urn* process in which, at time 0, there are r red balls and g green balls. At each time step, a ball is chosen at random from the urn and that ball, along with another ball of the same color, are put back. An extension to this process would be, for each iteration, to replace the picked ball along with c more of the same color ($c = 1$ in the Polya urn case).

In this case, the the probability of the number of red balls similar to Equation (2.1)

$$\mathbb{P}\{k \text{ red balls at time } t\} = \binom{t}{k} \frac{\Gamma\left(\frac{r+b}{c}\right)}{\Gamma\left(\frac{g}{c}\right)\Gamma\left(\frac{g}{c}\right)} \frac{\Gamma\left(k + \frac{r}{c}\right)\Gamma\left(t - k + \frac{g}{c}\right)}{\Gamma\left(t + \frac{g}{c} + \frac{r}{c}\right)}$$

The corresponding distribution of the proportion of red balls to the number of balls in the urn is $\beta(r/c, g/c)$.

The OK Corral process corresponds to the case where $c = -1$. The story motivating this process is a shootout between two gangs. Before the shootout starts there are r live members of the first gang and g live members of the second. At any time t , the probability that a member of first gang will be shot is equal to the proportion of size of the second gang to the total number of gunfighters, $Y_t/(X_t + Y_t)$. It is assumed that each shot is immediately fatal. The process is run until all members of one of the gangs is killed. The distributions of interest are the time when one gang kills all members of the other gang and, for the victorious side, the number of surviving gang members.

The OK Corral process falls under a more general set of process which are described as *sacrificial* in Flajolet, Dumas, and Puyhaubert (2006). For these urn processes $c < 1$. These models generally add the requirement that c divides both r and g so that on population is guaranteed to reach zero.

5. THE STOCHASTIC APPROXIMATION

Let's once again consider the Polya urn process, $c = 1$. The number of red balls at time $t + 1$ can be written recursively as

$$(5.1) \quad X_{t+1} = X_t + R_{t+1}.$$

Let $\Delta_{t+1} = X_{t+1} - X_t$, then we can rewrite Equation (5.1) as

$$\begin{aligned} \Delta_{t+1}X &= R_{t+1} - \mathbb{P}_t R_{t+1} + \mathbb{P}_t R_{t+1} \\ &= M_{t+1} + \mathbb{P}_t R_{t+1} \end{aligned}$$

M_{t+1} is a martingale with zero expectation. Furthermore,

$$\begin{aligned} M_{t+1} &= R_{t+1} - \mathbb{P}_t R_{t+1} \\ &= R_{t+1} - \frac{r + X_t}{r + g + t} \\ &= R_{t+1} - W_t. \end{aligned}$$

If we look at the total variation distance between R_{t+1} and W_t

$$\begin{aligned} \sup\{|\mathbb{P}R_{t+1} - \mathbb{P}W_t|\} &= \sup\{|\mathbb{P}\mathbb{P}_t R_{t+1} - \mathbb{P}W_t|\} \\ &= \sup\left\{\left|\mathbb{P}\left(\frac{r + X_{t+1}}{r + g + 1}\right) - \mathbb{P}W_t\right|\right\} \\ &= \sup\{|\mathbb{P}W_{t+1} - \mathbb{P}W_t|\} \\ &\rightarrow 0 \text{ as } t \rightarrow \infty. \end{aligned}$$

The last step comes from Section (3), where it was shown that the distribution of the proportion of red balls is converging to a constant distribution.

This result tells us that, for t big enough, the change in the number of red balls at time $t + 1$ behaves like the expectation of the number of balls added.

$$\Delta_{t+1}X \approx \mathbb{P}_t R_{t+1} = \frac{r + X_t}{r + g + t} = \frac{X_t}{r + g + t} + \frac{r}{r + g + t}.$$

The term $r/(r + g + t)$ approaches zero as t increases. Therefore, it will be ignored. The difference equation can be approximated by the differential equation

$$\frac{dX}{dt} = \frac{X}{r + g + t}$$

which can be rewritten as

$$\frac{dX}{X} = \frac{dt}{r + g + t}.$$

Taking the indefinite integral and the log of both sides yields the equation

$$(5.2) \quad X(t) = (r + g + t)e^m$$

for some constant m . Remembering that initially there are r red balls we can solve for e^m

$$e^m = \frac{r}{(r + g)}.$$

Substituting back into Equation (5.2)

$$X(t) = (r + g + t) \frac{r}{(r + g)} = \frac{r}{(r + g)}t + r.$$

The number of red balls increases linearly with time. This result is somewhat unsurprising. We already know that the distribution of the probability of getting a red ball at some time t is approximately $\beta(r, b)$ and so we can expect that the number of red balls will increase linearly in time. The derivation serves to confirm this suspicion.

This particular example, of creating a dynamic system from the expected behavior of a process, falls under the general category of *stochastic approximation*. These processes can be put into the form

$$Z_{t+1} - Z_t = \frac{1}{\eta_t} (F(Z_t) + \xi_{t+1} + D_t)$$

and satisfy the following requirements:

- F is a vector field on \mathbb{R}^n .
- ξ_{t+1} is a martingale which converges and satisfies $\mathbb{P}_t \xi_{t+1} = 0$.
- The remainder term D_t goes to zero as t increases
- $\sum_{t=1}^{\infty} D_t / \eta_t < \infty$.

When these requirements are satisfied, for a big enough t , the differential equation

$$\frac{dZ}{dt} = F(Z)$$

approximates the behavior of the process $\{Z_t\}$.

6. PATH COUNTING

Stochastic approximations rely on the convergence of a martingale. The question then arises: what happens if the process does not continue for an arbitrarily long time so that the martingale term does not have enough time to converge? This is potentially the case for the previously described OK Corral process. To understand these processes, Flajolet, Dumas, and Puyhaubert (2006) find urn configuration probabilities based on *path counting* which model the number of ways an urn can have a configuration, given the initial condition of r red balls and g green balls. This section will provide a sketch for the justification and describe some results reached in the paper.

As a motivating example, consider the Polya urn process with $r = g = 1$ going from time 0 and 1. During this iteration there is 1 red ball which may be picked. If it is picked, the ball count at time 2 is 2 red balls and 1 green balls. Similarly, at this iteration, there is 1 green ball to pick. If it is picked, then the ball count at time 2 is 1 red ball and 2 green balls. Compare this to the polynomial

$$xy(1+x)(1+y) = xy + x^2y + xy^2 + x^2y^2.$$

If we take the exponent of the x term as the red ball count and the exponent of the y term as the green ball count, then the terms x^2y and xy^2 correspond to these cases. The coefficient of the terms correspond to the number of different ways to reach a count of 3 balls. To consider the possible ball counts at time 2, consider the polynomial

$$\begin{aligned} x^2y(1+x)(1+y) + xy^2(1+x)(1+y) \\ = x^2y + x^3y + 2x^2y^2 + xy^3 + x^3y^2 + xy^2 + x^2y^3. \end{aligned}$$

At time 2, there are 4 balls in the urn corresponding to term x^3y , $2x^2y^2$, and xy^3 . The only way to have 3 red balls and 1 green ball in the urn at this time is if a red ball is picked twice. Likewise, the only way to have 3 green balls and 1 red ball in the urn at this time is if a green ball is picked twice. There can be 2 red balls and 2 green balls if a red ball is picked at the first iteration and a green is picked on the second iteration or a green ball is picked at the first iteration and a red is picked at the second. This corresponds to a coefficient of 2 for the term $2x^2y^2$.

Now that we can count the total number of *paths* to a given urn configuration for the Polya urn model, let's find the total number of configurations of the urn at a given time. Again, consider the urn model where the balls, along with the sampled ball, is returned to the urn at any time t . At $t = 1$, there are $r + g$ balls to choose from. At time $t = 2$, there are $r + g + 1$ balls to choose from. Let $s = r + g$, then the total number of different ways to choose balls through time t is

$$s(s+1)\dots(s+t-1) = t!c^t \binom{n+s-1}{t}.$$

This is referred to as the number of *histories* of the process.

From the difference equations describing an urn model, Flajolet, Dumas, and Puyhaubert (2006) give probabilities of urn configurations at given times based on path counts. These difference equations must have a linear recurrence representation, the sum of the rows of the state transition matrix must be equal, and if balls are being taken from the urn, the number being taken must divide the number of balls.

After finding a set of difference equations which describe an urn process, a polynomial generating function is derived. To get a configuration probability at time t , a derived differential operator is applied to the polynomial t times and the generating function variable is set to zero. The calculated number can then be divided by the the total number of paths to get a probability of a path.

As a final example, once again, consider the Polya urn process. The linear recurrence representation is

$$\begin{bmatrix} X_{t+1} \\ Y_{t+1} \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} R_{t+1} \\ G_{t+1} \end{bmatrix} + \begin{bmatrix} X_t \\ Y_t \end{bmatrix}.$$

Using the technique described in the paper, the differential system is

$$\begin{cases} \frac{dx}{dt} = x^2 \\ \frac{dy}{dt} = y^2 \end{cases}.$$

Solving this pair and remembering that initially there are r red balls and g green balls

$$\begin{cases} X_t = \frac{r}{1-tr} \\ Y_t = \frac{g}{1-tg} \end{cases}$$

which corresponds to a generating function $H(x, y, t)$ of

$$H(x, y, z) = \frac{x^r y^g}{(1 - tx)^r (1 - ty)^g}.$$

To get the probability that this urn will have a red balls and g green balls at time n , the differential operator is applied n times and x, y , and t is set to zero. Then, this number is divided by the total number of histories to get

$$\mathbb{P}\{k \text{ red ball at time } n\} = \binom{n}{k} \frac{\Gamma(r+g)}{\Gamma(g)\Gamma(r)} \frac{\Gamma(k+r)\Gamma(n-k+g)}{\Gamma(n+g+r)}$$

which agrees with the results of Section (2).

7. CONCLUSION AND FUTURE WORK

This paper describes several different approaches to solving the Polya urn process. Although it was only mentioned, each of these techniques lend themselves analyses of more complicated processes (some of which were described). Also, in many cases, extending the analysis to these processes is straight forward.

It may have been noted that the description of the history counting technique left some holes in its justification as well as the example. This was mostly due to time constraints. Since they were more relevant to the class, analyses from Sections 2, 3, and 5 were deemed more important.

Along with gaining a better understanding of the path counting technique, future work may consist of applying stochastic approximation to the graph coloring problem described in Achlioptas and Molloy (1997). This paper describes a greedy graph coloring algorithm which is applied to Erdős-Rényi graphs ($G(n, m/n)$). Consider the case where n is large. When the algorithm is run, trying to color the graph with k colors, there is

a threshold value γ such that if $m < \gamma$, the algorithm will succeed with high probability. If $m > \gamma$, then the algorithm will fail with high probability. The algorithm works by keeping track of possible colorings for each vertex. When a vertex is colored, adjacent vertices update their list of possible colorings. A stochastic approximation could be used to model these color list sizes.

REFERENCES

- Achlioptas, D. and M. Molloy (1997). The analysis of a list-coloring algorithm on a random graph (extended abstract). *Proceedings of the 38th IEEE Symposium on Foundations of Computer Science*, 204–212.
- Flajolet, P., P. Dumas, and V. Puyhaubert (2006). Some exactly solvable models of urn process theory. *Proceedings of Fourth Colloquium on Mathematics and Computer Science AG*, 59–118.
- Panchenko, D. (2007). Theory of probability. Lecture Notes from MIT Open Courseware.
- Pemantle, R. (2007). A survey of random processes with reinforcement. *Probability Surveys* 4, 1–79.
- Pollard, D. (1997). Unexpected symmetry. Lecture Notes for STAT 241/541.
- Pollard, D. (2001). *A User's Guide to Measure Theoretic Probability*. Cambridge University Press.
- Wikipedia (2007). Bernstein polynomial. [Online; accessed Dec-21-2007].