Image Segmentation and Labeling Using the Polya Urn Model

A. Banerjee† P. Burlina† F. Alajaji‡

†Department of Electrical Engineering
University of Maryland at College Park
College Park, MD 20742-3275
and
‡Department of Mathematics and Engineering
Queen's University
Kingston, ON K7L 3N6, Canada

Abstract

We propose a segmentation method based on Polya's model for the spread of contagion. An initial segmentation is obtained using a Maximum Likelihood (ML) estimate or the Nearest Mean Classifier (NMC). The resulting clusters are then subjected to a morphological process operating like the development of an infection to yield segmentation of the image into homogeneous regions. This process is implemented using contagion urn processes and generalizes Polya's scheme by allowing spatial interactions. The urn mixture model provides a fuzzy representation of the pixel label. The composition of the urns is iteratively updated by assuming a Markovian relationship between neighboring pixel labels. The asymptotic behavior of this process is examined. Examples of the application of this scheme for the segmentation of Synthetic Aperture Radar (SAR) images and Magnetic Resonance Images (MRI) are provided.

1 Introduction

We describe a segmentation method using contagion urn schemes that rely on a modified version of the Polya-Eggenberger sampling process [8]. It consists of a biologically inspired sampling procedure originally designed to model the development of contagious phenomena.

Many approaches have been studied for segmentation. Unsupervised segmentation approaches include the Nearest Mean Classification (NMC) and the branch and bound procedure [3]. Supervised methods generally proceed by formulating statistical model assumptions for the image formation and the region generative processes. Maximum likelihood (ML) or maximum *a posteriori* estimation (MAP) are then used for segmentation. Examples of such approaches abound in the literature [4; 6; 11]. Techniques modeling images as Markov random fields (MRF) have been extensively investigated [4]. MRF's attempt to represent spatial dependencies and the MRF-Gibbs equivalence allows for the computation of the maximum *a posteriori* (MAP) estimate of the original image [4].

This paper models images using urn processes. The motivation for employing urn schemes is twofold: first, urn processes can generate Markov chains as well as MRF's [5]. Second, urn schemes are of particular interest because they provide a natural representation for fuzzy image labeling. Therefore, they constitute an attractive generative process for the underlying image regions which exhibit strong spatial dependencies. Our work is related to the Gibbs sampling procedure [4] by preserving key features of the Gibbs sampler but using instead a contagion sampling scheme. The spatial dependencies of the pixel labels are captured by the contagious behavior which promotes smoothing of the image into contiguous regions. The urn process for segmentation is related to relaxation labeling algorithms, except that the urn process is not deterministic [10].

We begin by applying either ML or NMC segmentation technique to the image. The contagion process is then applied to the image labels. In this scheme, each pixel is represented by an urn with a mixture of balls of different colors, one color for each class label. A neighborhood system is also defined on each pixel. The balls of the urns of the neighborhood system are then combined to determine the next state of the urns. The iterative nature of the algorithm incorporates temporal memory, while the inclusion of the neighboring urns in the update promotes spatial contagion. Moreover, the neighborhood system is modified, pending the existence of an edge element in the neighborhood. This is done to preserve edges by containing the propagation of similar class labels within closed boundaries.

This paper is organized as follows: the initial NMC and ML segmentations are presented in Section 2. The contagion-based smoothing process is then described in Section 3. In Section 4, the stochastic properties of the resulting image process are discussed. Finally, experimental results on SAR and MR images are shown in Section 5.

2 Initial Segmentation

When no a priori information on the image statistics is available, general clustering algorithms such as NMC usually are applied. In the NMC method, an initial arbitrary labeling is used from which centroids of the feature vectors of each class are computed. Next, all samples are reclassified to the cluster corresponding to the nearest mean, and the centroids are recomputed. This process is iterated until a stopping criterion is met [3].

In contrast, when a stochastic model for the image can be justified, it is possible to apply ML segmentation. The conditional distribution of the image; i.e., the form of

$$p(X_s/C_s = l, C_r; r \in N_s^k) \qquad (1)$$

is assumed. Here, C_s is the label for pixel s, C_r represents the pixel labels of N_s^k, the k^{th} order neighborhood of pixel s, and X_s is the given image data [4].

For segmentation purposes, we estimate the pixel labels by assuming that the conditional probability of each class label, i.e. $p(X_s/C_s = l)$, is governed by a multivariate Gaussian distribution on the second order neighborhood N_s^2.

After obtaining the parameters of the different classes, the ML test determines the label for each pixel in the image. The ML decision rule is described as

$$\hat{l} = \mathit{argmax}_l \; p(X_s/C_s = l, C_r; r \in N_s^2). \qquad (2)$$

The above schemes do not capture the statistics and connectedness of local regions. Since the ML test assumes that each pixel label is equally likely throughout the image, it produces noisy segmentation. This assumption is incorrect, for in a local region dominated by one class, the dominant class has a higher prior probability than the other classes. Such contextual information is not taken into account in either the ML or NMC estimate of the pixel labels.

This drawback is usually addressed within the framework of MAP segmentation. The MAP estimate of the class label $\hat{l}$ for a pixel given the observed image X_s is

$$\hat{l} = \mathit{argmax}_l \; p(C_s = l/C_r; r \in N_s^2, X_s) \qquad (3)$$

Indeed, it can be shown that maximizing $p(C_s = l/C_r; r \in N_s^2, X_s)$ is equivalent to maximizing $p(C_s = l/C_r; r \in N_s^2)p(X_s/C_s = l)$.

If segmentation of the image into homogeneous regions is desired, it is intuitively appealing to model the prior distribution $p(C_s = l/C_r; r \in N_s^2)$ using an MRF, as the MRF model relates the label of a pixel with the labels of its neighboring pixels [6].

If the prior is modeled as an MRF, the Gibbs-MRF equivalence can be exploited by techniques such as simulated annealing (SA) or other stochastic relaxation methods to derive the MAP estimate [6].

Unfortunately, techniques such as simulated annealing have high computational costs. Indeed convergence to the MAP estimate is possible only when impractically slow annealing schedules are followed. Instead, we propose to replace the annealing step by an urn contagion process to model the spatial dependencies between neighboring pixels.

3 Image Sampling with Contagion

The labeled image is described by an urn process. Each pixel in the image is represented by an urn containing a mixture of balls of different colors representing the classes. The proportion of each class in the urn indicates the similarity of the pixel to the class. The urn representation is therefore a fuzzy representation of the segmented image. At each iteration the current urn is modified by re-sampling with contagion, a process that is inspired by the original urn sampling process introduced by Polya and Eggenberger in [8].

3.1 Temporal Contagion

The work reported in [8] introduced the following urn scheme as a model for the spread of a contagious disease through a population. An urn originally contains T balls, of which W are white and B are black ($T = W + B$). Successive draws from the urn are made; after each draw, $1 + \Delta$ ($\Delta > 0$) balls of the same color as was just drawn are returned to the urn. Let $\rho = W/T$ and $\delta = \Delta/T$. Define the binary process $\{Z_n\}_{n=0}^{\infty}$ as follows:

$$Z_n = \begin{cases} 0, & \text{if the } n^{th} \text{ ball drawn is white;} \\ 1, & \text{if the } n^{th} \text{ ball drawn is black.} \end{cases}$$

It can be shown that the process $\{Z_n\}$ is stationary and non-ergodic [9; 7]. The urn scheme has infinite memory, in the sense that each previously drawn ball has an equal effect on the outcome of the current draw.

3.2 Temporal and Spatial Contagion

The urn sampling scheme proposed in this work incorporates both temporal and spatial contagion. Instead of representing an image by a finite lattice of

pixels, we consider an image as a finite lattice of urns. In the "one-dimensional" urn sampling described above, the effect of each sample propagates through time. For the "two-dimensional" case, the sampled ball at each iteration must depend not only on the composition of the pixel's urn, but also on the composition of the neighboring urns to encourage contagious behavior. Thus, we need to allow for spatial interactions at each time instant by *associating* the urns of the neighboring pixels in the determination of the newly sampled ball.

3.3 A Fuzzy Image Labeling Representation

The following presentation considers, without loss of generality, a binary labeling problem. Let $I_n = [p_n^{(i,j)}]$ be a binary label image of size $K \times L$, where $p_n^{(i,j)} \in \{0,1\}$ is the label of pixel (i,j) at iteration n, $n = 0,1,\ldots$, $(i,j) \in \mathcal{I}$ where

$$\mathcal{I} : \{(i,j) : i = 0,\ldots,K-1; j = 0,\ldots,L-1\}.$$

To each pixel we associate an urn $u_n^{(i,j)} : (B_n^{(i,j)}, W_n^{(i,j)})$ with each pixel (i,j) at time n, where $B_n^{(i,j)}$ and $W_n^{(i,j)}$ are respectively the number of black and white balls in the urn. With this representation we define a membership function for each pixel as:

$$m_F^B(p_n^{(i,j)}) = \frac{B_n^{(i,j)}}{(B_n^{(i,j)} + W_n^{(i,j)})}$$

3.4 An Algorithm for Segmentation with Spatial Contagion

The general class of algorithms for the contagion-based smoothing process can be described as follows:

• Initialization

Let I_0 be an initial segmentation (at time index $n = 0$). For each pixel (i,j), the initial urn composition $u_0^{(i,j)} = (B_0^{(i,j)}, W_0^{(i,j)})$ is obtained by computing the relative frequencies of white and black pixels in a spatial neighborhood centered on (i,j). For this work, the second order (3x3) neighborhood system for each pixel is adopted.

• Iterative Image Sampling

For $n > 0$, the urn composition of each pixel (i,j) is updated by sampling from a combination of the participating urns $\mathcal{V}_{n-1}^{(i,j)}$ with $\mathcal{V}_{n-1}^{(i,j)} : \{u_{n-1}^{(r,s)} : (r,s) \in N_s^q\}$, where N_q^k is the neighborhood system defined as in [4]:

$$N_q^k : \{q = (r,s) \in \mathcal{I} : (i-r)^2 + (j-s)^2 \leq k\}.$$

A simple, yet effective, sampling procedure is as follows: the urn $u_n^{(i,j)}$ for pixel $p_n^{(i,j)}$, is updated by first combining the balls of $u_{n-1}^{(i,j)}$ and the N neighboring urns:

$$\mathcal{C}_{n-1}^{(i,j)} = ASSOCIATE(\mathcal{V}_{n-1}^{(i,j)}). \quad (4)$$

The $ASSOCIATE$ function forms a collection of balls, $\mathcal{C}_{n-1}^{(i,j)}$, from the urns of the neighborhood. Examples of the $ASSOCIATE$ function include grouping the urns of $\mathcal{V}_{n-1}^{(i,j)}$ into a "super" urn or sampling one ball from each urn to form the collection. Furthermore, the neighborhood may be modified if an edge element exists in that neighborhood; if so, those neighboring urns which lie on the other side of the edge are excluded. This is necessary to preserve edges and limit contagion to local areas.

Next, an operation on the new collection of balls, $\mathcal{C}_{n-1}^{(i,j)}$, is performed i.e.,

$$Z_n^{(i,j)} = SELECT(\mathcal{C}_{n-1}^{(i,j)}). \quad (5)$$

The select function may determine the next state of the urns by sampling one ball from $\mathcal{C}_{n-1}^{(i,j)}$ or by taking the majority class of $\mathcal{C}_{n-1}^{(i,j)}$.

We denote $Z_n^{(i,j)}$ as the outcome of the $SELECT$ function:

$$Z_n^{(i,j)} = \begin{cases} 0, & \text{if the } n^{th} \text{ ball drawn is white;} \\ 1, & \text{if the } n^{th} \text{ ball drawn is black.} \end{cases}$$

If $Z_n^{(i,j)} = 0$, add Δ white balls to urn $u_n^{(i,j)}$; if $Z_n^{(i,j)} = 1$, add Δ black balls to urn $u_n^{(i,j)}$.

This yields a new urn composition for each pixel as

$$u_n^{(i,j)} : \begin{cases} W_n^{(i,j)} = W_{n-1}^{(i,j)} + (1 - Z_n^{(i,j)}) * \Delta, \\ B_n^{(i,j)} = B_{n-1}^{(i,j)} + (Z_n^{(i,j)}) * \Delta. \end{cases}$$

The above procedure is iterated until $n = N$. At time N, the final composition of each individual urn $u_N^{(i,j)}$, $(i,j) \in \mathcal{I}$ determines the final labeling of the image. As described above, each urn represents a fuzzy measure of the pixel label.

4 Statistical Properties

4.1 Temporal Contagion

The resulting sequence of generated images exhibits both spatial and temporal dependencies represented by a Markovian relationship in terms of the urns $u_n^{(r,s)}$; more specifically:

$$Pr\{u_n^{(i,j)} | U_{n-1}, U_{n-2}, \ldots, U_0\} = Pr\{u_n^{(i,j)} | \mathcal{V}_{n-1}^{(i,j)}\},$$

where $U_n : [u_n^{(i,j)}]$ is the urn matrix associated with I_n, and $\mathcal{V}_{n-1}^{(i,j)}$ is the set of participating urns defined in the previous section.

Consider the original Polya sampling scheme. The asymptotic properties of the joint distribution can be characterized in the "one-dimensional" case, i.e., when all spatial interactions are inhibited at each sampling step. In this case, it can be shown [9] that the proportion of white balls in each urn after the n'th trial $\rho_n^{(i,j)}$, where

$$\rho_n^{(i,j)} = \frac{\rho + \left(Z_1^{(i,j)} + Z_2^{(i,j)} + \cdots + Z_n^{(i,j)}\right)\delta}{1 + n\delta},$$

is a martingale [2] and admits a limit Z as the number of draws increases indefinitely. Indeed, $\rho_n^{(i,j)}$ (or equivalently the sample average $\frac{1}{n}\sum_{k=1}^{n} Z_k^{(i,j)}$) converges with probability one to Z [2]. This limiting proportion Z is a continuous random variable with support the interval $(0,1)$ and Beta probability density function with parameters $(\rho/\delta, (1-\rho)/\delta)$:

$$f_Z(z) = \begin{cases} \frac{\Gamma(1/\delta)}{\Gamma(\rho/\delta)\Gamma((1-\rho)/\delta)} z^{\frac{\rho}{\delta}-1}(1-z)^{\frac{1-\rho}{\delta}-1}, \\ if\ 0 < z < 1; \\ \\ 0, \text{otherwise.} \end{cases}$$

$\Gamma(\cdot)$ is the gamma function described by

$$\Gamma(x) = \int_0^\infty t^{x-1}e^{-t}dt \quad \text{for} \quad x > 0.$$

The behavior of this pdf can be interpreted as follows: assuming $\delta = 1$ for simplicity, if the original proportion of white balls in the urn is close to one, then the limiting distribution of $W_n^{(i,j)}$, will be skewed towards 1. A similar behavior is obtained for the case when ρ is close to zero. Therefore, the limiting pattern will reflect the underlying probability

$$Pr\left(p_1^{(i,j)} = x\right) = \rho^x(1-\rho)^{(1-x)}.$$

For the M-ary labeling case, the above observations generalize with convergence to the Dirichlet distribution [5].

4.2 Temporal and Spatial Contagion

We examine the asymptotic behavior of two examples of a general urn sampling scheme for segmenation.

• Method1

Consider sampling from the "super" urn. Restating the problem, suppose there are N urns in the neighborhood of pixel X_s, each *initially* with b_i black balls and w_i white balls, and $b_i + w_i = T$ for all i, $i = 1, 2, \ldots, N$. We put the contents of all N urns into a "super" urn, sample one ball, and add Δ balls of the same color into the urn of the pixel X_s. The following properties are easily derived.

The probability of sampling exactly k black balls from n iterations of the "super" urn is

$$Pr(X = k) = \binom{n}{k} \frac{B(\alpha + k, \beta + n - k)}{B(\alpha, \beta)}, \quad (6)$$

where $\alpha = \sum_i \frac{b_i}{\Delta}$, $\beta = \sum_i \frac{w_i}{\Delta}$, and the beta function $B(\alpha,\beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$.

The above process can be regarded as being generated by a sequence of independent Bernoulli trials with parameter Z, where Z is random with beta distribution. In fact, it is identical with different parameters to the Polya-Eggenberger distribution in the "one-dimensional" case given above.

The average number of black balls in the "super" urn at any given time is expressed as

$$E[B_n] = \sum_j b_j \frac{NT + n\Delta}{NT}. \quad (7)$$

Therefore, the average proportion of black balls in the "super" urn is

$$E\left[\frac{B_n}{(NT + n\Delta)}\right] = \sum_j \frac{b_j}{NT} \quad (8)$$

Remarkably, the average proportion of black balls in the "super" urn at any time instant equals the original proportion of black balls. The above results show that the composition of the urn is highly dependent on the original proportion of the balls. Eventually, the majority class of the urns in a given neighborhood will spread and dominate the population of balls in that neighborhood. Therefore, we conclude that this urn sampling scheme will reinforce the majority class in a local spatial neighborhood; it constitutes a positive-feedback system that yields limiting patterns of the self-reinforcing type [1]. The contagion effectively models the Markovian dependencies of the pixel labels.

• Method2

This second example is described as follows: We sample one ball from each of the urns in pixel $X_{(i,j)}$'s neighborhood, $\mathcal{V}_{n-1}^{(i,j)}$. From this collection of balls, we compute the majority class, denoted $Z_n^{(i,j)}$. We update urn $u_n^{(i,j)}$ in the same manner described in the previous section, i.e.

$$u_n^{(i,j)} : \begin{cases} W_n^{(i,j)} = W_{n-1}^{(i,j)} + (1 - Z_n^{(i,j)}) * \Delta, \\ B_n^{(i,j)} = B_{n-1}^{(i,j)} + (Z_n^{(i,j)}) * \Delta. \end{cases}.$$

The urn composition of urn $u_n^{(i,j)}$ is governed by the Polya-Eggenberger distribution as explained above. Eventually, the initial majority class of each urn in the neighborhood will dominate its composition.

It is difficult to find a general closed-form expression for $P(Z_n^{(i,j)} = k)$, the probability that class k is the majority of the individual samples. The difficulty arises because we are trying to find the majority of a set of samples of a non-i.i.d. process. Hence, we resort to heuristic arguments. Results of this method are given in Figure 1.

5 Experimental Results

For segmentation of SAR imagery, we start with ML segmentation. As shown in Figure 1(a), the resulting labeling is spotty, a characteristic of the ML segmentation technique. Application of simulated annealing generates a contiguous segmentation of the image (Figure 1(b)). Likewise, Figure 1(c) shows that ten iterations of the urn sampling scheme operating on the ML segmentation yields an image labeled into locally homogeneous regions. The SAR image used in this example is Lincoln Laboratory's ADTS SAR data, which is fully polarimetric with 1 foot resolution.

Whereas simulated annealing achieves segmentation by optimizing a function (the MAP pixel label estimate), modified urn schemes smooth the segmented image by morphologically processing the pixel labels. Since the Polya urn schemes model contagious behavior in a population, modified urn schemes allow dominant pixel labels to propagate within local regions, analogous to diffusion methods for segmentation. The advantage in using the urn scheme lies in the reduction of the computational complexity of the segmentation algorithm. We avoid the time and computational costs of simulated annealing by employing a simpler algorithm.

To segment the MR images, we obtain an initial segmentation by the NMC. The inherent noise of this image modality leads to the speckled segmentation. The contagion urn process then operates on the pixel labels to produce a smoother segmentation. The output after one and ten iterations are shown in Figures 2(b) and 2(c), respectively. Note that the edges are preserved by limiting contagion to local areas. An edge map, computed by the Canny edge detector, is employed to modify each pixel's sampling neighborhood to prohibit sampling over different region types.

In both cases, urn sampling method 2 is implemented; one ball is sampled from the urns of the neighborhood and a majority rule is applied to determine the next state of the urns. For this work, each urn is initialized with ten total balls, and Δ, the number of balls added at each iteration, is two.

6 Conclusion

In this paper, we have illustrated how modified Polya urn sampling schemes can be implemented for image segmentation. Given an initial speckled segmentation, the contagion process obtains a smoother segmentation into homogeneous regions by its Markovian properties. Two general properties incorporate temporal and spatial contagion. First, iterative updating is required for temporal contagion. Second, sampling from neighboring urns, similar to the Gibbs sampler, is necessary to encourage spatial contagion.

Further lines of research include the evaluation of the optimal values for the parameter δ, the ratio of Δ to the initial number of balls in an urn. For instance, if δ is too high, then the segmentation is over-smoothed; if it is too low, then the algorithm may not converge to the appropriate segmentation. As mentioned above, the initial composition of the urns determines to a great extent the outcome of the contagion process. Therefore, finding an appropriate method to initialize the urn composition is critical to accurately segment the image.

References

[1] W. B. Arthur, Y. M. Ermoliev, and Y. M. Kaniovski, "Path-Dependent Processes and the Emergence of Macro-Structures," *European Journal of Operational Research*, pp. 294–303, 1987.

[2] W. Feller, *An Introduction to Probability Theory and its Applications*, vol. 2, John Wiley & Sons Inc., 1971.

[3] K. Fukunaga, *Statistical Pattern Recognition*, Academic Press, 1990.

[4] S. Geman and D. Geman, "Stochastic Relaxation, Gibbs Distribution, and Bayesian Restoration of Images," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 6, November 1984.

[5] N. Johnson and S. Kotz, *Urn Models and Their Application*, John Wiley & Sons, Inc., 1977.

[6] P. A. Kelly, H. Derin, and K. D. Hartt, "Adaptive Segmentation of Speckeled Images Using a Hierarchical Random Field Model," *IEEE Transactions on Acoustics, Speech and Signal Processing*, 1988.

[7] G. Polya, "Sur Quelques Points de la Théorie des Probabilités," *Ann. Inst. H. Poincarré*, Vol. 1, pp. 117–161, 1931.

[8] G. Polya and F. Eggenberger, "Uber die Statistik Verketteter Vorgänge," *Z. Angew. Math. Mech.*, pp. 279–289, 1923.

[9] G. Polya and F. Eggenberger, "'Sur l'Interpretation de Certaines Courbes de Fréquences," *Comptes Rendus C. R.*, pp. 870–872, 1928.

[10] S. Ranade and A. Rosenfeld, "Point pattern matching by relaxation," *Pattern Recognition*, Vol. 12, pp. 269–275, 1980.

[11] E. Rignot and R. Chellappa, "Segmentation of Polarimetric Synthetic Apperture Radar Data," *IEEE Trans. on Image Processing*, Vol. 1, pp. 281–300, 1992.

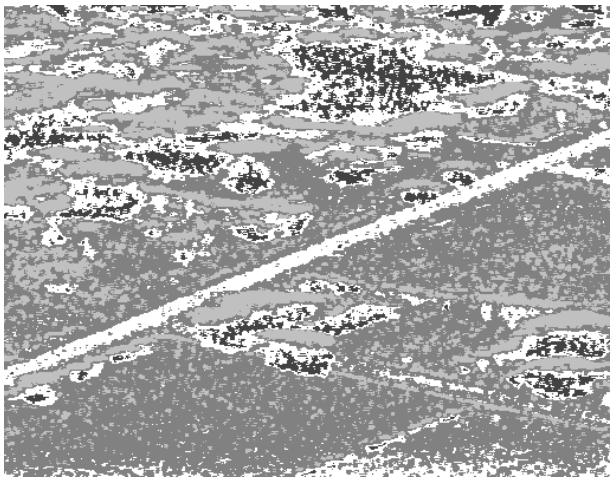

(a) ML Segmentation.

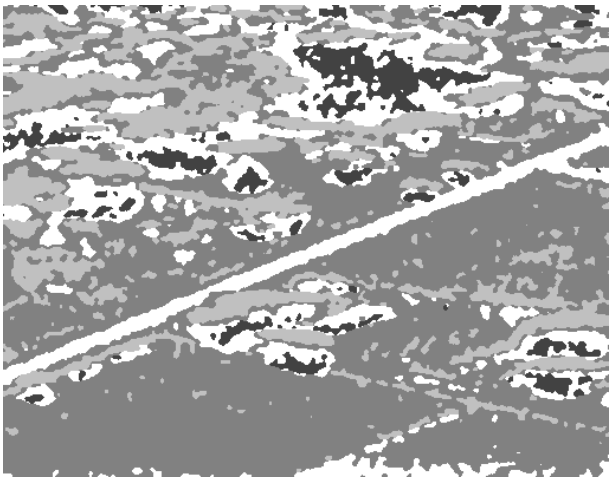

(b) Segmentation after Simulated Annealing.

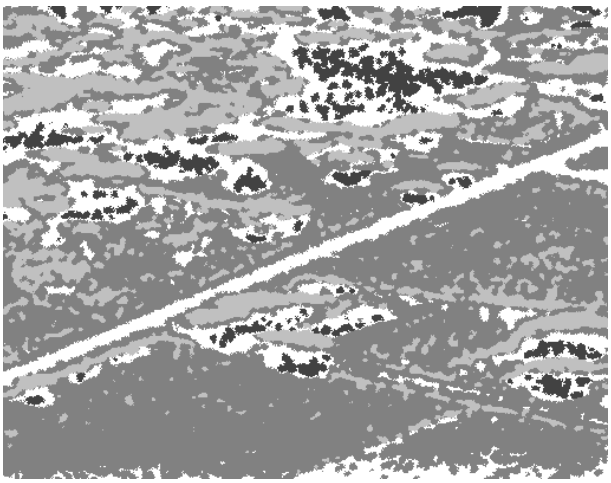

(c) After Urn Process.

Figure 1: Segmentation ADTS SAR Image after 10 iterations of SA and Urn Process

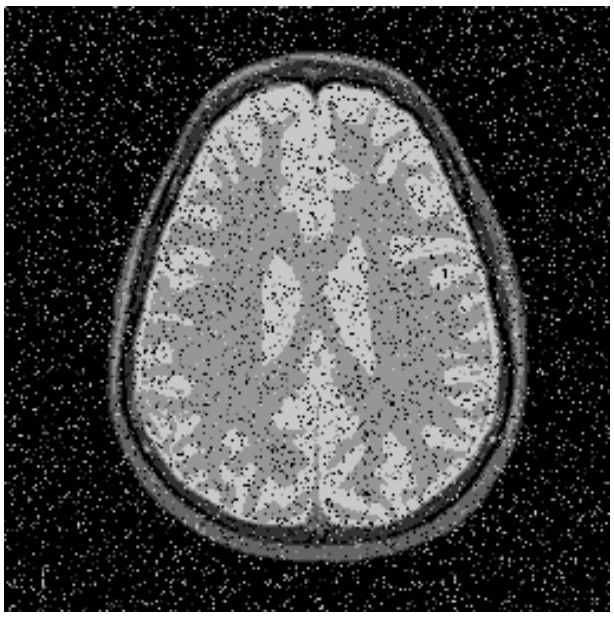

(a) Noisy NMC segmentation.

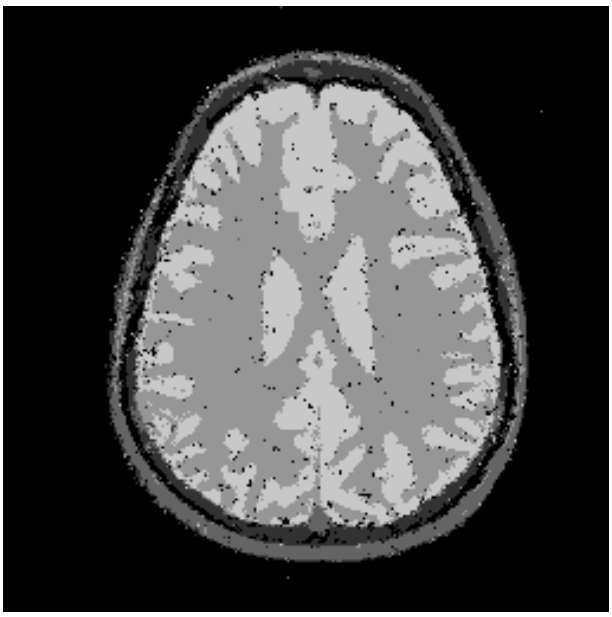

(b) After 1 iteration of Urn Process.

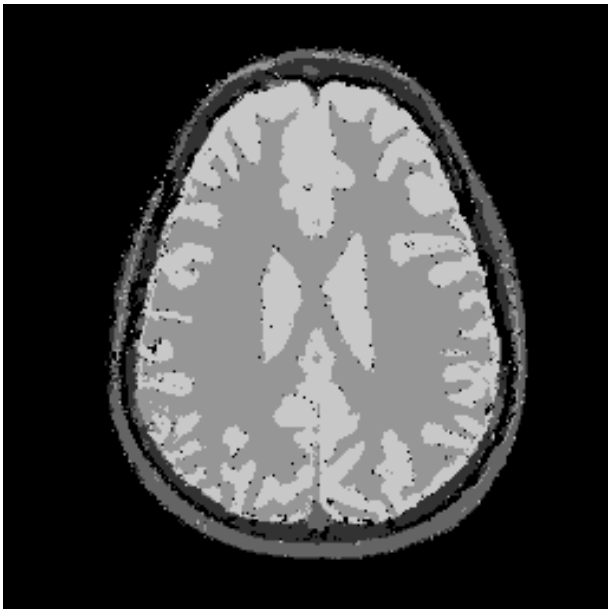

(c) After 10 iterations of Urn Process.

Figure 2: Segmentation of MR Images using Urn Process with Inhibition