Published in final edited form as:

Proc IEEE Int Conf Acoust Speech Signal Process. 2010 March 14; 2010: 453–456. doi:10.1109/ICASSP.2010.5495723.

CONVERGENCE BEHAVIOR OF THE ACTIVE MASK SEGMENTATION ALGORITHM

Doru C. Balcan¹, Gowri Srinivasa², Matthew Fickus³, and Jelena Kovačević⁴

¹ School of Interactive Computing, Georgia Institute of Technology, Atlanta, USA

² Dept. of Info. Sc. and Eng., and Center for Pattern Recognition, PES School of Engineering, Bangalore, India

³ Dept. of Mathematics and Statistics, Air Force Institute of Technology, Wright-Patterson AFB, USA

⁴ Dept. of Biomedical Eng., Electrical and Computer Eng. and Center for Bioimage Informatics, Carnegie Mellon University, Pittsburgh, USA

Abstract

We study the convergence behavior of the Active Mask (AM) framework, originally designed for segmenting punctate image patterns. AM combines the flexibility of traditional active contours, the statistical modeling power of region-growing methods, and the computational efficiency of multiscale and multiresolution methods. Additionally, it achieves experimental convergence to zero-change (fixed-point) configurations, a desirable property for segmentation algorithms. At its core lies a voting-based distributing function which behaves as a majority cellular automaton. This paper proposes an empirical measure correlated to the convergence behavior of AM, and provides sufficient theoretical conditions on the smoothing filter operator to enforce convergence.

Index Terms

active mask; cellular automata; convergence; segmentation

1. INTRODUCTION

Recent advances in fluorescence microscopy have made it possible for biologists to study particular proteins of interest in cells and conduct live cell imaging. Manually processing the enormous amount of collected data is extremely tedious, and so the need for accurate and efficient automated processing methods has become imperative. One of the first tasks after image acquisition is segmentation—delineation of cells (or objects of interest within them) from each other and the background. Since fluorescence microscope images of proteins in cells lack edges, most of the traditional segmentation algorithms in image processing or computer vision literature fail to produce meaningful results.

Segmentation with active contours

The class of algorithms generically called *active contours* is considered state-of-the-art in biomedical image segmentation. Its success resides in the ability to incorporate a general and flexible set of forces to drive the segmentation. In previous work [1], we successfully adapted the active contour framework to segment punctate patterns of fluorescence

microscope cell images. Adding multiscale and multiresolution techniques into the mix enabled us to exploit the computational structure of the problem efficiently: we first perform segmentation on a coarse approximation of the image and later refine the outcome by running only a few iterations of the algorithm at high resolution [2]. Finally, by combining the flexibility of active contours with the computational efficiency of multiscale methods, and the statistical modeling power of region-growing methods, we obtained a high-performance segmentation algorithm: the *Active Mask* (AM) framework [3].

Segmentation with active masks

The basic idea of the AM framework is to represent a segmented image ψ via a collection of M binary masks¹ μ_m . The AM iteration rule is:

$$\psi^{(i+1)}(n) = \underset{1 \leq m \leq M}{\operatorname{argmax}} \{V_g(\mu_m^{(i)})(n) + R_m(n)\}, \quad (1)$$

where n is the spatial position of a pixel, i is the iteration number, V_g is a smoothing filter, and R_m is a region-based distributing function (a soft threshold of the image) that skews the vote in mask m . For segmentation of cell images, R_1 is designed to distinguish the background region from the foreground and $R_m = 0$, for all $m \neq 1$. Thus, background pixels are represented by the first mask, while foreground pixels are distributed to remaining masks through *voting*. Eventually, the procedure halts if/when the pixels cease to change mask membership, or equivalently, when $\psi^{(i+1)} = \psi^{(i)}$. This typically happens when there is a smooth boundary between the masks; the degree of smoothness depends on the shape and scale parameter of the smoothing filter g used in the voting process.

Motivation and objectives

Experiments on images containing punctate patterns of cells (or specific proteins) demonstrated that AM successfully delineates regions of interest (see Fig. 1) even in cases where many existing image segmentation algorithms fail (for a thorough discussion, see for instance [3].) Nevertheless, as is the case with so many novel automated processing methods, understanding at a deeper level the true source of such flexibility and efficiency is sometimes very difficult. A question that was not conclusively answered by the earlier work was, “why does the AM algorithm always reach a zero-change state?”. This could be followed by “is convergence primarily a property of the smoothing filter?”. If so, it would also be good to know “how to design more flexible filters for AM algorithms using such “basic properties” as guiding principles?”. The general goal of this paper is to shed light onto these issues, but its primary focus will be to answer the first question. Namely, we explain the convergence behavior of the AM algorithm both empirically and theoretically. We show how AM can be very naturally framed as a type of cellular automaton, and we exploit this framework by identifying simple sufficient conditions on the structure of the filter which will ensure convergence. Due to the relevance of majority cellular automata in science, the results afforded by this setting enable us to establish beneficial connections with very diverse fields, ranging from social sciences [5] to statistical physics [6,7].

The paper is organized as follows: In Section 2, we present the AM framework in more detail, and provide an empirical measure tightly correlated with the convergence behavior. In Section 3, we describe how the Active Mask algorithm can naturally be framed as a Majority Cellular Automaton. Section 4 contains a summary of sufficient conditions for the

¹Masks are defined as: $\mu_m(n) = 1$ if $\psi(n) = m$, and 0 otherwise.

convergence of majority automata. Finally, we conclude with a discussion of the results in Section 5.

2. DISTRIBUTING FUNCTIONS AND EMPIRICAL CONVERGENCE

As described above, the active mask framework is instantiated by defining appropriate distributing functions. These partition the image by distributing the pixels to different masks and are akin to forces in the active contour framework.

Majority voting based distributing function

We start with image $f \in \ell^2(\mathbb{Z}^d)$, such that $f(n) = 0$ iff $n \notin \Omega$; $\Omega := \prod_{k=1}^d [0, N_k)$ is a rectangular subset of the d -dimensional integer lattice $\mathbb{Z}^d$. We define a collection of M masks ψ , as a function that assigns each pixel $n \in \Omega$ a value $\psi(n) \in \{1, \dots, M\}$. Then, V_g , the crux of the majority voting based function is defined by the nonstandard noncommutative convolution,

$$V_g(\mu_m)(n) := \begin{cases} \frac{(\mu_m \star g)(n)}{(\chi_\Omega \star g)(n)}, & n \in \Omega; \\ 0, & n \notin \Omega; \end{cases} \quad (2)$$

where χ_Ω is the characteristic function of Ω , and g a lowpass filter with a tunable scale parameter a . Given a random initialization of ψ , the majority voting based distributing function is the evolution of ψ as described in (1).

Region-based distributing function

Since (2) does not use any information directly derived from the image, the iterative voting process could produce different partitions of the image. To drive the segmentation to a meaningful configuration, we introduce a distributing function R derived from the image. It could be based on any set of features that distinguish region(s) of interest from each other. (For a more detailed explanation, and a concrete example, the reader is invited to [3].) Since we are interested in delineating cells and assume their statistical properties are relatively uniform and different from those of the background, here we consider only one nonzero R_m (arbitrarily chosen as R_1). In general, there could be as many nonzero R_m 's as there are regions of interest with distinct properties. As mentioned earlier, R is some function of the image (in our case, a soft threshold that separates the foreground and background regions). This restricts the voting to a specified segment in the image: in our case, the voting after the application of R 's is restricted to the foreground (see Fig. 1(c–d)).

Zero crossings and convergence

Throughout the experiments reported in our work on fluorescent microscope images (1) the iteration process always came to a halt, which lead us to postulate that there should exist a measure that is monotonically decreasing with iteration. While convergence properties of algorithms similar to (1) have been studied in the continuous domain, it is not easy to find a reliable discrete equivalent of monotonically decreasing functions [6,7]. This is because the argmax function governing the discrete process is neither linear, nor continuous. However, the action of iterative smoothing using lowpass filters such as g is well studied as the maximum (or minimum) principle [8] (a discrete version of the maximum principle is presented in [9]). Based on the diffusion equation, conditions under which a lowpass filter produces a coarse version of the image have been presented in [10]. It has been rigorously proved that a version of the maximum principle from the theory of parabolic differential

equations is equivalent to the condition of applying such a lowpass filter on an image so as not to produce any new zero-crossings [11].

In the case of active masks, for each pixel we count one zero crossing for each of its 8 immediate neighbors that belong to a different mask. For very general conditions on the filter g and its scale (or spread) parameter a , this number is empirically seen to decrease monotonically. Fig. 2 depicts such an example for 2 masks, using a Gaussian filter under three different scale conditions; each plot contains the evolution of 5 runs of the algorithm, initialized randomly and independently. We would like to point out the extremely abrupt decrease of the “empirical potential”, depicted here through a 4-times application of the natural logarithm to the number of zero-crossings. By inspecting the curves, we notice that the steepness of the average slope ($\omega_3 \approx -0.125$, $\omega_5 \approx -0.064$, and $\omega_8 \approx -0.019$) decreases monotonically with scale, in the regime where the filter is neither too small to make filtering trivial, nor too large to render the influences of distributing functions superfluous.

Next, we shall describe the mathematical details and working assumptions used when framing AM as a Cellular Automaton.

3. ACTIVE MASK AS MAJORITY CELLULAR AUTOMATON

Cellular automata are self-evolving discrete dynamical systems [12]. and they have found application in various fields like particle physics, computational biology, and social sciences. Tremendous amount of work in this area has focused on studying the convergence behavior of various types of automata. Since the crux of AM involves a self-evolving dynamical system via the voting function, in this section we shall formulate AM as a majority cellular automaton in order to facilitate our understanding of its convergence behavior.

We begin by expressing the convolution operation defined in (2) as the result of a linear operation acting on a vector: $\mathbf{f} \star \mathbf{g} = \mathbf{A}\mathbf{f}$. This will later enable us to prove the convergence of the AM algorithm via the structure of matrix $\mathbf{A}$ alone.

Based on (1), the “contribution” corresponding to mask m in deciding the outcome at location n at iteration i shall be

$$V_m^{(i)}(n) = (\mu_m^{(i)} \star g)(n) = (\mathbf{A} \cdot \mu_m^{(i)})(n). \quad (3)$$

Henceforth, we refer to $\mathbf{A}$ as the *influence* or *weight matrix*.

In case of a tie in (1), we can decide to break it by choosing the smallest m corresponding to a maximal element. In other words:

$$\psi_i(n) := \min \left(\arg \max_m V_m^{(i-1)}(n) + R_m(n) \right) \quad (4)$$

For any fixed cell n and iteration i , let $\xi_m = V_m^{(i)}(n) + R_m(n)$. This allows us to interpret (4) as follows: first we inspect all values $\xi_1, \xi_2, \dots, \xi_M$; then we collect *all indices* corresponding to the maximum value $\xi_{\max}$ of this sequence; finally, we assign the value of the *smallest* such index to $\psi_i(n)$.

Influence matrices in image processing

When using cellular automata on regular lattices, especially in image processing applications, a typical and intuitive assumption is that pixels' relative influences should depend primarily on their value and distance. As the voting functions are filtering/convolution operators, this symmetric influence assumption reflects upon the filter g , and equivalently upon the weight matrix $\mathbf{A}$, being symmetric. For example, by assuming periodic convolution and filter symmetry, matrix $\mathbf{A}$ itself will be symmetric and circulant (in 1-D) or block-circulant with circulant block² structure (in 2-D).

Whenever standard assumptions like periodicity are not appropriate, alternative convolution operators are preferred which apply different edge-handling schemes and thus scale the result of the filtering according to the spatial position of each pixel. For instance, the AM convolution operator used in (2) is implemented by zero-padding and local scaling.) In such situations, the corresponding linear transform is represented by a quasi-symmetric matrix [5].

Definition

A square matrix $\mathbf{B}$ is *quasi-symmetric* iff it can be written as $\mathbf{B} = \Lambda \mathbf{A}$, where $\mathbf{A}$ is symmetric and Λ is diagonal.

In the following section, we will characterize the convergence behavior of general Majority Cellular Automata, and in particular of AM, in terms of structural properties of the weight matrix.

4. SUFFICIENT CONDITIONS FOR CONVERGENCE

Convergence behavior results

Theoretical guarantees about the convergence behavior of Majority Cellular Automata with symmetric influences (although without additive terms) exist in the literature since the early 80's (see [12], [14], and references therein). A generalization to quasi-symmetric matrices, together with an instructive proof based on Lyapunov functionals, was recently published in [5]. We further extend these previous results to cover the general AM setting where the non-zero region-based additive term R can critically influence the result of the voting. In fact, we can prove the result even if a nonzero term R_m exists for each mask.

Before stating the theorem, let us make the observation that since there is a finite number of configurations and the production rules are deterministic, regardless of the initial configuration the automaton is bound to reach a state it has previously "visited". This means that its attractors are either fixed points or cycles of length at least 2. This rather trivial property is independent of the weight matrix. However, using more particular matrices results in a more "easily predictable" behavior of the automaton.

Theorem 4.1

If matrix $\mathbf{A}$ is quasi-symmetric, then the attractors of the corresponding AM Cellular Automaton can be either fixed points or cycles of length 2.

Proof sketch

For simplicity, we only address the symmetric case; the quasi-symmetric one is very similar. Without losing generality, we can assume that the starting configuration $\psi^{(0)}$ belongs to an

²Also referred in the literature as *BCCB* [13].

attractor, that is to a cycle of period K , and we will safely consider that iterations are residues modulo K . Suppose that $K > 2$; then, there exists a cell n and an iteration $i \in \mathbb{Z}_K$ such that $\psi_{i-2}(n) > \psi_i(n)$. By the definition of the AM automaton we have:

$$\left[\mathbf{A} \cdot \mu_{\psi_i(n)}^{(i-1)} \right](n) + R_{\psi_i(n)}(n) > \left[\mathbf{A} \cdot \mu_{\psi_{i-2}(n)}^{(i-1)} \right](n) + R_{\psi_{i-2}(n)}(n), \quad (5)$$

while for all the other values of cells and iterations we have a similar inequality, although not necessarily strict. After summing all inequalities over i and n , and some algebraic manipulations, we get

$$\sum_{i \in \mathbb{Z}_K} \sum_{n=1}^N \left[\mathbf{A} \cdot \mu_{\psi_i(n)}^{(i-1)} \right](n) > \sum_{i \in \mathbb{Z}_K} \sum_{n=1}^N \left[\mathbf{A} \cdot \mu_{\psi_{i-2}(n)}^{(i-1)} \right](n). \quad (6)$$

However, both of the left and the right term are equal to the same value, namely

$$\text{tr} \left[\sum_{i \in \mathbb{Z}_K} \mathbf{M}^{(i-1)^T} \mathbf{A} \mathbf{M}^{(i)} \right], \text{ where } \mathbf{M}^{(i)} \text{ is the matrix having the binary mask vectors } \mu_m^{(i)} \text{ as its columns. This produces a contradiction and so } K \leq 2.$$

One immediate consequence of Theorem 4.1 is that the convergence behavior of the AM Automaton remains the same, regardless of local scaling and additive R terms. Obviously, the final configurations will not generally be the same; this is precisely the reason behind their being employed by AM in the voting process: to bias the vote towards more “informative” configurations.

Fixed-point guarantees

The convergence result above is consistent with extensive empirical evidence reported in this paper and in previous work on Active Masks [3]; namely for any random starting state, the algorithm *always* converges to some fixed point when a Gaussian filter is used by the voting functions. However, this also brings up another interesting issue: why is AM with a Gaussian smoothing filter always converging *only to fixed points*? Since a Majority Cellular Automaton with symmetric weights may, in general, also have oscillatory attractors (see the “box-filter” example described later on), we can conclude that the (quasi-)symmetry condition is too weak to explain *exclusive* fixed-point convergence.

It turns out that a sufficient condition for the convergence of the Majority Cellular Automaton with symmetric weights to fixed points only is the positive semi-definiteness of the weight matrix. This result is stated by the following theorem. Due to lack of space we omit the proof. However, this and subsequent results regarding the convergence of AM (*e.g.*, generalization to the quasi-symmetric case) will be included in a future paper.

Theorem 4.2

Let $\mathbf{A}$ be a symmetric, positive semidefinite (PSD) real weight matrix. Then the corresponding Majority Cellular Automaton necessarily converges to fixed points.

Since any symmetric matrix $\mathbf{A}$ has real but not necessarily nonnegative eigenvalues the positive semi-definiteness condition truly is stronger than symmetry.

“Box-filter” and related (counter)examples

The presence of oscillating patterns for the 3-tap “box-filter” automaton (with periodic convolution, and an even number of points) is explained by the fact that the corresponding filter matrix is not PSD. Indeed, for a circulant matrix the set of eigenvalues is given by the Fourier transform of its first row [13]. Equivalently, they will be the result of evaluating a polynomial (in this case, $X^{N-1} + X + 1$) at the roots of unity. However, if $N \geq 4$ is even, then (-1) is a root of unity, and by evaluating the above polynomial at (-1) we get $(-1) < 0$, which proves our claim that matrix $\mathbf{A}$ is not PSD. By similar arguments, we can show that the “cross” filter (von Neumann’s automaton) in 2D presents oscillating states, namely the checkerboard patterns. Finally, the 3×3 box filter in 2D (also known as the Moore automaton) also has oscillating states, namely the 2×2 checkerboard patterns.

Gaussian periodic case

Let us point out that when the periodicity assumption is enforced (*i.e.*, when the matrix is circulant or BCCB respectively), simultaneous symmetry of the filter g with respect to both coordinate axes together with the nonnegativity of its Fourier transform, guarantees the positive semidefiniteness of the matrix and as a result, the corresponding AM algorithm will necessarily converge to fixed points. Such is the case of toroidal space with a Gaussian convolution filter. An extreme particular case is the trivial one: when the “scale” parameter of the Gaussian is tiny, g becomes a delta function and $\mathbf{A}$ is the identity matrix; then the algorithm converges in one step no matter from which state it may start. Of course, the value of such an “extreme” AM setup is vacuous; like the experiment in Fig. 2 this suggests that a “reasonable size” of the filter is important for the quality of the segmentation.

Diagonally dominant case

Finally, another interesting example is the case when the weight associated to each pixel (*i.e.*, the “central” filter tap) is larger than the sum of the absolute values of the weights of all other influences. For instance, in 1D the corresponding filter matrix $\mathbf{A}$ will be (symmetric and) diagonally dominant; by Gershgorin’s Circles Theorem the eigenvalues of $\mathbf{A}$ will be (real and) nonnegative, which then implies fixed-point convergence.

5. CONCLUSION

We presented an alternative, yet natural, view of the Active Mask segmentation algorithm [3] within the framework of Cellular Automata. By viewing AM as a special Majority Automaton, we could provide theoretical guarantees for its convergence behavior in terms of structural properties of the smoothing filter g or influence matrix $\mathbf{A}$. Besides a simple generalization of known convergence theorems whose object is the quasi-symmetric case, we were able to find sufficient conditions on $\mathbf{A}$ (*e.g.*, PSD) to guarantee convergence to fixed points only, which confirms empirical and experimental evidence.

We also studied whether a fixed mathematical measure (for our discrete, square lattice case) can be reliably associated with convergence behavior, that is, it varies monotonically as AM converges. We empirically found that the number of zero-crossings matches this profile. The connection to cellular automata literature brought up a very interesting analog: the Lyapunov functional associated to an automaton [12]. Although the number of zero-crossings may be a Lyapunov functional, a rigorous proof in this sense remains out of reach. Regardless, we will continue to ponder the richness of this problem and expect that our better understanding it will result in more efficient and intelligent biomedical image processing systems.

Acknowledgments

This work was supported in part by NSF through award DMS-0405376, by NIH through award R03-EB008870, and by AFOSR through grant FIATA09125G003.

References

1. Coulot, L.; Kirschner, HE.; Chebira, A.; Moura, JMF.; Kovačević, J.; Osuna, EG.; Murphy, RF. Topology preserving STACS segmentation of protein subcellular location images. Proc. IEEE Int. Symp. Biomed. Imaging; Arlington, VA. Apr. 2006; p. 566-569.
2. Srinivasa, G.; Fickus, MC.; Kovačević, J. Multiscale active contour transformations for the segmentation of fluorescence microscope images. Proc. SPIE Conf. Wavelet Appl. in Signal and Image Proc; San Diego, CA. Aug. 2007; p. 1-15.
3. Srinivasa G, Fickus M, Guo Y, Linstedt A, Kovačević J. Active mask segmentation of fluorescence microscope images. IEEE Trans Imag Proc. 2009
4. Linstedt Lab at CMU. <http://www.cmu.edu/bio/faculty/linstedt.html>
5. Melatagia Yonta P, Ndoundam R. Opinion dynamics using majority functions. Mathematical Social Sciences 2009;57:223–244.
6. Ruuth SJ. Efficient algorithms for diffusion-generated motion by mean curvature. Journ Comp Phys 1998;144:603–625.
7. Ruuth SJ, Merriman B. Convolution-thresholding methods for interface motion. Journ Comp Phys 2001;169:678–707.
8. Nirenberg L. A strong maximum principle for parabolic equations. Commun Pure and Appl Math 1953;6:167–177.
9. Perona P, Malik J. Scale-space and edge detection using anisotropic diffusion. IEEE Trans Patt Anal and Mach Intelligence 1990;12(7):629–639.
10. Koenderink JJ. The structure of images. Journ Bio Cybern 1984;50(5):363–370.
11. Hummel, RA. Morgan Kaufmann Readings Series archive. Morgan Kaufmann Publishers, Inc; CA, USA: 1987. Readings in computer vision: issues, problems, principles, and paradigms, chapter Representations based on zero-crossings in scale-space; p. 753-758.
12. Golés, E.; Martínez, S. Neural and Automata Networks. Dynamical Behavior and Applications. Kluwer; 1990.
13. Davis, PhJ. Circulant Matrices. 2. AMS Chelsea Publishing; 1994.
14. Poljak S, Sura M. On periodical behaviour in societies with symmetric influences. Combinatorica 1983;3(1):119–121.

Fig. 1.

Active mask segmentation of punctate patterns of proteins [3]. (a) Original image (courtesy of Linstedt Lab [4]) (b) Initialization using 256 random masks. (c) After one iteration of (1), the background is separated from the foreground by the region-based distributing function. (d) After 12 iterations, cells in foreground have been assigned to $M = 12$ distinct masks, which do not change with further iterations for the chosen scale a of filter g and at that resolution. (e) On gradually changing the scale parameter to merge spurious masks (multiple masks corresponding to one cell) and successively refining the segmentation outcome at higher resolutions to include relevant details in the corresponding cell masks, we obtain the final segmentation outcome of $M = 10$ masks at iteration $i = 179$ for the given random initialization.

Fig. 2. Decrease in the number of zero-crossings with iteration. Scale of the Gaussian filter is inversely correlated with convergence speed.