

Fast Random Walker with Priors using Precomputation for Interactive Medical Image Segmentation

Shawn Andrews, Ghassan Hamarneh, Ahmed Saad

Medical Image Analysis Lab, Simon Fraser University, Canada,
{sda56, hamarneh, aasaad}@sfu.ca

Abstract. Updating segmentation results in real-time based on repeated user input is a reliable way to guarantee accuracy, paramount in medical imaging applications, while making efficient use of an expert's time. The random walker algorithm with priors is a robust method able to find a globally optimal probabilistic segmentation with an intuitive method for user input. However, like many other segmentation algorithms, it can be too slow for real-time user interaction. We propose a speedup to this popular algorithm based on offline precomputation, taking advantage of the time images are stored on servers prior to an analysis session. Our results demonstrate the benefits of our approach. For example, the segmentations found by the original random walker and by our new precomputation method for a given 3D image have a Dice's similarity coefficient of 0.975, yet our method runs in $1/25^{th}$ of the time.

1 Introduction

Segmentation is a crucial task in medical imaging. Manual segmentation by an expert is accurate, but is also very time consuming, while fully automatic and accurate segmentation techniques are not yet a reality, thus semi-automatic techniques become a necessity. While many semi-automatic techniques assume only user initialization, repeated user interaction is necessary to guarantee the accuracy required for medical imaging. Therefore, it is critical to speed up these techniques, especially in 3D, in order to minimize the time spent waiting between a user inputting information and seeing the results [1, 2].

A full survey of semi-automatic algorithms is beyond the scope of this work [3, 4]. At a high level, semi-automatic algorithms can be divided into several classes. One class involves the specification of an approximate boundary, which evolves towards the correct segmentation by minimizing a cost function derived from shape priors and image information [5, 6]. Another class of algorithms requires the user to specify sequential points on or near the boundary, and then the boundary is filled in between these points using a minimal path approach [7, 8]. A third class of algorithms asks the user to provide seeds, or pixels within specific regions, and then uses these seeds as a basis for the segmentation [9, 10].

An example of the last class of algorithms is the seeded random walker (RW_SD) [10], which is a graph-based approach to image segmentation that, along with its extensions, has garnered hundreds of citations in only a few years. It boasts many advantages, including weak boundary detection, robustness to noise, trivial generalization

to simultaneous multi-region and 3D segmentations, a globally optimal solution ensuring repeatability, a probabilistic segmentation that can be very useful in directing a user to areas of uncertainty, and the straightforward user input method of providing seed pixels - all important features in medical imaging. On the downside, once seeds are given, RW_SD computes the segmentation by solving a large system of equations, which can be slow. In [1], this problem is alleviated by introducing RW with precomputation (RW_PREC). Since medical images usually exist “offline” on servers for some time before they are segmented, some precomputation can be done before user input that allows a fast approximation to the segmentation once seeds are given, or “online”. The speedup RW_PREC provides allows for user interaction with RW_SD in real-time.

Unfortunately, RW_SD has some limitations, specifically the segmentation is calculated based only on localized image data and disconnected regions must be seeded individually. These problems are addressed in [11], where regional intensity priors are introduced into the formulation. The priors result in more accurate segmentations and the ability to segment disjoint regions easily. However, in RW with priors (RW_PR), the image graph is not completely known offline, since priors are usually derived from the seeds and precomputation must be performed before seeds are given. This obsoletes the methods introduced in [1], which require the graph to be known. An algorithm with the robustness of RW_PR and the online speedup of RW_PREC would be a very useful interactive segmentation tool.

In this paper, we make the following contributions. First, starting with the random walker equations from [11], we derive an offline precomputation and an online approximation that allows for a significant online speedup that can be used in conjunction with priors. Secondly, we derive some additional precomputations that are performed offline to further speed up the online segmentation. Combining the robustness of RW_PR and the online speedup of RW_PREC, we create a useful interactive segmentation tool applicable to a more general class of problems than RW_PREC. Code demonstrating our method is available from <http://mial.cs.sfu.ca>.

2 Methods: Random Walker Improvements

We begin by giving a brief review of existing RW algorithms for later reference. In the following derivations, we consider binary segmentations, but we note that our methods extend trivially to multiple labels, just as all previous RW algorithms discussed do. RW_SD [10] constructs a graph and then defines L as the graph’s Laplacian matrix and x as a vector of the probabilities that each node belongs to the object being segmented. We define N , S , $U = N - S \gg S$ as the numbers of nodes in the graph, seeded nodes, and unseeded nodes, respectively. In RW_PR [11], λ is introduced as a vector of the prior probabilities for each node, weighted by a scalar γ . In RW_PREC [1], the first K eigenvalue/eigenvector pairs of L are found and stored in matrices A and Q respectively, so $L \approx QAQ^T$. These are used to construct a pseudo-inverse for L , $E = QA^{-1}Q^T$. For all three algorithms, all of the variables involved are expressed in terms of their

components corresponding to seeded (S) and unseeded (U) nodes:

$$L = \begin{bmatrix} L_S & B \\ B^T & L_U \end{bmatrix}, \quad x = \begin{bmatrix} x_S \\ x_U \end{bmatrix}, \quad E = \begin{bmatrix} E_S & R \\ R^T & E_U \end{bmatrix}, \quad Lx = f = \begin{bmatrix} f_S \\ f_U \end{bmatrix}, \quad g = \begin{bmatrix} g_S \\ g_U \end{bmatrix}, \quad (1)$$

where g is a constant eigenvector corresponding to the zero eigenvalue of L .

In RW_PR, x_U , the vector of RW probabilities for the unlabeled pixels, is obtained by solving a $U \times U$ system of equations (with I defined as the identity matrix):

$$(L_U + \gamma I_U)x_U = -B^T x_S + \gamma \lambda, \quad (2)$$

where $\gamma = 0$ gives RW_SD. In RW_PREC, several equations of size $S \times S$ are given whose solutions can be used to approximate x_U in the case that $\gamma = 0$. This method [1] is reviewed in [12]. When priors are added to the RW formulation ($\gamma > 0$), the graph is no longer known offline, and this precomputation method is rendered useless.

2.1 Precomputation with Priors

Following [1], our goal is to derive a method for using offline precomputation to speed up the online computation for the more general RW_PR, where the image graph is altered after seeds are given. We define L and x as in RW_PR [11]. Following preliminary steps similar to RW_PREC, we derive

$$(I_U + \gamma E_U)x_U - g_U \alpha = R^T f_S + \gamma E_U \lambda. \quad (3)$$

The details of the derivation are available in [12]. To proceed in our derivation, we will assume for the moment that we can calculate $J_U^{-1} = (I_U + \gamma E_U)^{-1}$ and define $\hat{B} = B J_U^{-1}$. Now, replacing B with \hat{B} in the derivation and taking $f_S = \hat{f}_S + \bar{f}_S \alpha$ and $\hat{P} = (I_S - \hat{B} R^T)$, we derive

$$\hat{P} \hat{f}_S = L_S x_S + \gamma \hat{B} E_U \lambda, \quad \hat{P} \bar{f}_S = \hat{B} g_U, \quad \alpha = \frac{g_S^T \hat{f}_S + \gamma g_U^T \lambda}{\gamma - g_S^T \bar{f}_S}. \quad (4)$$

The details are again in [12]. Now that we have f_S we plug it back into (3) to get

$$x_U = J_U^{-1}(g_U \alpha + R^T f_S + \gamma E_U \lambda). \quad (5)$$

The issue still exists of how to compute J_U^{-1} efficiently, as defining J_U requires the seeds and inverting a matrix of size U is too expensive of an operation to perform during the online phase. We will define J to be the extension of J_U to size N given by

$$J = (I + \gamma E) \approx (Q Q^T + \gamma Q \Lambda^{-1} Q^T) \approx Q(I_K + \gamma \Lambda^{-1}) Q^T \quad (6)$$

$$\Rightarrow J^{-1} \approx Q(I_K + \gamma \Lambda^{-1})^{-1} Q^T \quad (7)$$

since $Q Q^T \approx I_N$, which implies that $Q_U Q_U^T \approx I_U$. Thus

$$Q_U(I_K + \gamma \Lambda^{-1})^{-1} Q_U^T J_U \approx I_U \quad \Rightarrow \quad J_U^{-1} \approx Q_U(I_K + \gamma \Lambda^{-1})^{-1} Q_U^T. \quad (8)$$

Since $(I_K + \gamma \Lambda^{-1})$ is a diagonal matrix with all positive entries, it is easily invertible, and can be inverted during the offline phase. The simple multiplication of Q_U to both sides during the online phase produces an adequate approximation to J_U^{-1} , fulfilling our goal of a RW formulation combining priors and precomputation (RW_PR_PREC).

2.2 Increased Speed using Extended Precomputation

By using the proposed RW_PR_PREC, we can reduce the online solving of (2) from RW_PR, an equation of size $O(U)$, to solving (4), two equations of size $O(S)$. However, due to the low connectivity of the graph, L_U is very sparse, and (2) can be solved in $O(U)$ time. The online phase of RW_PR_PREC must take $O(U)$ time also, since it returns a probability vector of size U . We want to perform as few $O(U)$ cost computations as possible by minimizing the number of matrix multiplications between matrices of size $O(U)$. Analysis and optimizations of these asymptotic run times were not considered previously in [1], which we do here by analysis of matrix operations.

Currently, in the offline phase, we compute Q and Λ^{-1} , but now we will precompute additional matrices to be used to speed up the online phase. The speedup will come from being able to retrieve the components of these matrices corresponding to the seeded nodes in $O(S)$ time. We note that Q is an $N \times K$ matrix and for storage space considerations we do not want our precomputed matrices to be larger than that. The details of the additional precomputations are in Algorithm 1, denoted OPT_RW_PR_PREC. By precomputing these 4 matrices, we save $(SK + K + K^2)U$ scalar multiplications, with the detailed calculations in [12]. We note that when not using priors, with RW_PREC, A_1 can still be used for additional speedup.

Algorithm 1 OPT_RW_PR_PREC:

Offline:

- 1: Calculate Q and Λ^{-1} from L
- 2: $A_1 = LQ$
- 3: $A_2 = Q^T Q$
- 4: $A_3 = (I_K + \gamma \Lambda^{-1})^{-1}$
- 5: $A_4 = Q^T g$

Online:

- 6: $\hat{P} = (I_S - (A_{1S} - L_S Q_S) A_3 (A_2 - Q_S^T Q_S) \Lambda^{-1} Q_S^T)$
 - 7: $\hat{P} \hat{f}_S = L_S x_S + \gamma (A_{1S} - L_S Q_S) A_3 (A_2 - Q_S^T Q_S) \Lambda^{-1} (Q_U^T \lambda)$
 - 8: $\hat{P} \hat{f}_S = (A_{1S} - L_S Q_S) A_3 (A_4 - Q_S^T g_S)$
 - 9: $\alpha = \frac{g_S^T \hat{f}_S + \gamma g_U^T \lambda}{\gamma - g_S^T \hat{f}_S}$
 - 10: $f_S = \hat{f}_S + \hat{f}_S \alpha$
 - 11: $x_U = Q_U A_3 ((A_4 - Q_S^T g_S) \alpha + (A_2 - Q_S^T Q_S) \Lambda^{-1} (Q_S^T f_S + \gamma (Q_U^T \lambda)))$
-

3 Results

We would like our results to show that using our precomputed data we make the RW online phase fast enough for interactive segmentation without compromising much accuracy. Therefore, we present results showing our high speed gains on real 2D and 3D data while maintaining negligible (and controlled) reduction in accuracy. We note that the speed increase allows much quicker seed editing and thus will translate to much improved accuracy per time spent by user.

The experiments here were performed using unoptimized MATLAB code run on an Intel Core 2 Duo (2.4GHz) with 4GB of RAM. The algorithms were implemented by the authors, utilizing Grady’s MATLAB Graph Toolbox [<http://www.cns.bu.edu/~lgrady>]. A negative exponential function was used for the edge weights, $w_{ij} = \exp(-\beta(|i_a - i_b|))$, where i_a is the intensity of pixel a . All experiment data was collected over 100 trials, and all parameters were chosen empirically and fixed across all compared methods. The only parameter effecting the speed of our method is K , the number of retained eigenvectors.

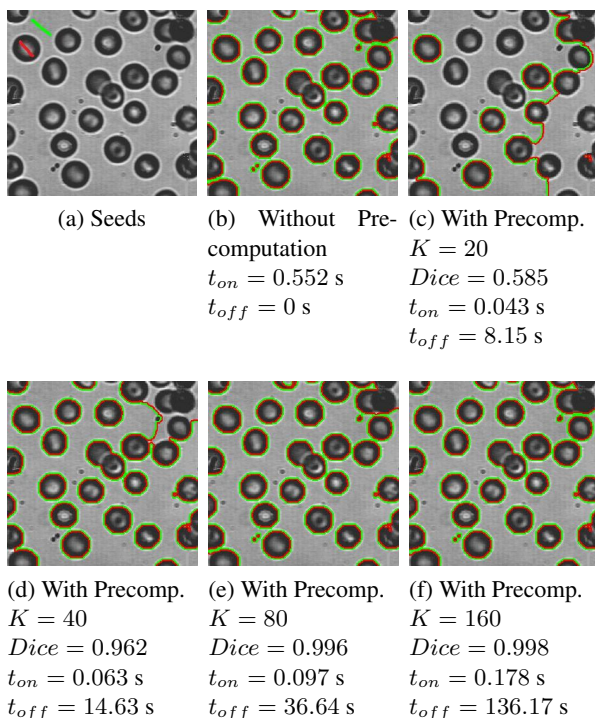


Fig. 1: (color figure) Comparison of results with and without precomputation for segmentation using priors on an image of size $N \approx 72,000$ pixels. For K , the number of eigenvectors used, we report $Dice$, the Dice similarity coefficient between RW_PR’s segmentation and OPT_RW_PR_PREC’s segmentation, t_{on} , the online time taken, and t_{off} , the offline time taken. We note that we are only concerned with t_{on} , and with $K = 80$, our method achieves excellent results in less than a fifth of the time taken when not using precomputation. Red and green correspond to different region boundaries.

The accuracy of the segmentations generated by our algorithms are evaluated by their similarity to the segmentations generated by RW_PR; the accuracy of RW_PR is well justified in other works [11]. We note that the speed and accuracy of our algorithms depend on the image only through K , and while we leave analytical methods for finding optimal K as future work, it was reported in [1] that $K = 40$ – 80 is often enough for 2D images, as our results in Section 3.1 corroborate, and we needed no more than $K = 350$ for larger 3D images as is seen in Section 3.2.

The resolution and noisiness of an image affect how large of a K is needed, and different noiseless images at the same resolution can require different values of K to be accurately segmented using precomputation. As seen in Fig. 1, even larger values of K provide online runtimes much faster than can be achieved without precomputation, so choosing K large enough to guarantee accuracy is our prime concern. As different

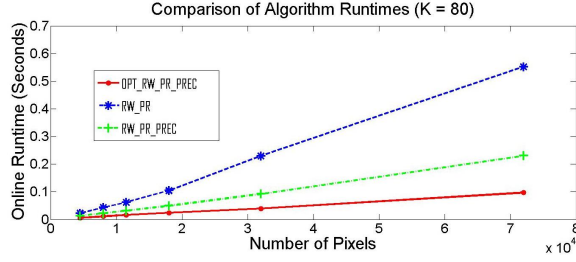


Fig. 2: Comparison of the runtimes of the original RW_PR (blue) and our proposed methods RW_PR_PREC (green) and OPT_RW_PR_PREC (red) for different resolutions of the image in Fig. 1. Note the standard deviation for the online runtimes are all under 0.05 seconds.

images affect the speed and accuracy of our algorithms only through how large K needs to be, and since we do not yet have an image dependent way to choose K (except based on resolution and noise), results for a variety of images would be redundant. Thus we focus our results on single 2D and 3D images at varying resolutions and with varying levels of noise. We note that offline runtime increases with K , but does not affect the application to interactive segmentation.

3.1 2D Results

Tests were performed on the 2D image in Fig. 1 of size $N = 265 \times 272 \approx 72,000$ pixels with an 8-connected image graph, $\beta = 30$, $\gamma = 0.001$, and two regions, where one region was divided into multiple disconnected sections and seeds were only put in one of these sections. These segmentation times do not include calculating the priors, an efficient step which was performed online, and is similar in all cases. The priors were calculated using a non-parametric density estimation with a Gaussian kernel [11]. Fig. 1 shows the Dice similarity coefficient and the average runtimes in seconds for both the online and offline phases of OPT_RW_PR_PREC for different values of K and compares the results to RW_PR, showing excellent speedup and minimal accuracy lost. Fig. 2 compares the runtimes of the different methods for different sized resolutions of the image in Fig. 1, again showing our precomputation gives excellent speedup.

3.2 3D Results

Tests were performed on a 3D CT image of the knee in Fig. 3 of size $N = 55 \times 55 \times 36 \approx 109,000$ voxels, a 26-connected image graph, and two regions, bone and non-bone. The bone region consists of 3 disconnected subregions (the femur, tibia, and patella). We tested the algorithms using priors by segmenting all the bones but placing seeds only in the tibia. We used RW_PR and OPT_RW_PR_PREC with $\beta = 100$ and $\gamma = 0.01$ and compared their average runtimes and the Dice similarity coefficient of their resulting segmentations. The average runtime of RW_PR was about 40.5 seconds, and when $K = 350$ eigenvectors are used, the average runtime of OPT_RW_PR_PREC was about 1.56 seconds. The Dice similarity coefficient between RW_PR's segmentation and OPT_RW_PR_PREC's segmentation was 0.975. Thus our method achieved a speedup of 25 times over RW_PR while maintaining excellent accuracy. The standard deviation of the runtimes of RW_PR was less than 1.0 s, and the standard deviation of the runtimes of OPT_RW_PR_PREC was less than 0.1 s.




Fig. 3: The bones of a knee segmented with $S = 100$ seeds in one of the bones. (a) was found using RW_PR in 40.5 seconds and (b) was found using OPT_RW_PR_PREC with $K = 350$ in 1.56 seconds. Dice's similarity coefficient between the two is 0.975.

3.3 Robustness to Noise

Here, we test the robustness to noise of OPT_RW_PR_PREC. We measured the similarity of the segmentations provided by the exact and approximate algorithms using the Dice similarity coefficient. The pixel intensities in our test images range from 0 to 1 and various levels of Gaussian noise with standard deviations $\sigma \in [0, 1]$ were added to the 2D image in Fig. 1 of size $N = 265 \times 272 \approx 72,000$ pixels with an 8-connected image graph, $\beta = 30$, and $\gamma = 0.001$. From Fig. 4a we see that OPT_RW_PR_PREC still provides good segmentations for small amounts of noise up to $\sigma = 0.2$ (with Dice's similarity coefficient > 0.95) if a large enough K is used. As the noise increases to $\sigma = 0.7$, Dice decreases. We can see the same trend in Fig. 4b, where $K = 200$ eigenvectors are used in the precomputation and the noise ranges from $\sigma = 0$ to 1.

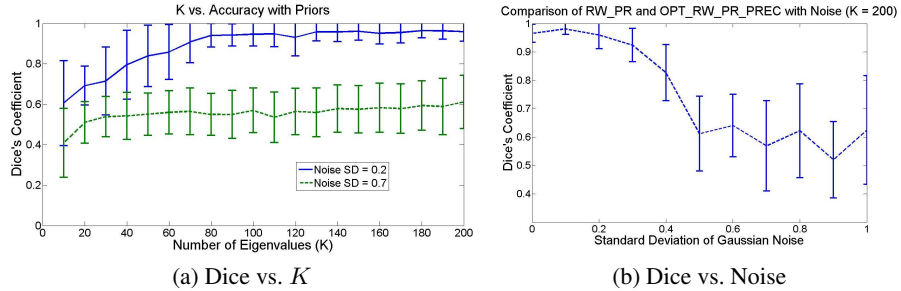


Fig. 4: Effect of K and noise on segmentation accuracy. (a) compares the Dice similarity coefficient between the segmentations found using RW_PR and OPT_RW_PR_PREC. Results are shown for two levels of noise and for multiple numbers of eigenvectors. (b) shows the Dice similarity coefficient between the segmentations at varying levels of noise with $K = 200$ eigenvectors and 20 trials for each level of noise. We see that large enough K lets us account for reasonable amounts of noise.

4 Discussion and Conclusions

The above tests give some strong results. We see from Figs. 1 and 2 that the additional precomputation of OPT_RW_PR_PREC greatly outperforms RW_PR in 2D, achieving a segmentation in about one fifth of the time and with over 99% similarity. From Fig. 3, we see the results are more pronounced in 3D, with OPT_RW_PR_PREC achieving speedups of a factor of 25 over RW_PR while still finding an almost identical segmentation with over 97% similarity. Furthermore, Fig. 2 shows that all the algorithms appear to increase linearly in runtime with the number of pixels, as predicted in Section 2.2.

Overall, we have derived a way of combining both precomputation and the use of priors into the popular RW algorithm. This allows RW to perform much faster segmentations when seeds and priors are either given or changed. Additionally, we've shown that some precomputations can be performed in addition to finding Q and Λ^{-1} that can greatly speed up the online phase of the algorithm. These improvements in speed provide a feasible way to enable the real-time editing of a wider variety of 2D and 3D images than was previously possible by allowing updating of both seeds and priors. This allows the user to ensure the accuracy of complex segmentations with minimal effort. Thus our contributions increase the usability and effectiveness of RW algorithms.

Future work will relate to using information from the image to automatically determine K , which needs to be set high enough to maintain accuracy. However, since the effect of a larger K is seen mostly in the offline phase which doesn't effect interactivity, we currently simply err on the side of caution when selecting K .

References

1. L. Grady and L. Kemal Sinop. Fast approximation random walker segmentation using eigenvector precomputation. *IEEE Trans. PAMI*, 2008.
2. Christopher J. Armstrong, Brian L. Price, and William A. Barrett. Interactive segmentation of image volumes with live surface. *Computers & Graphics*, 31(2):212–229, 2007.
3. S. D. Olabariaga and A. W. M. Smeulders. Interaction in the segmentation of medical images: A survey. *Medical Image Analysis*, 5(2):127–142, 2001.
4. Y. Kang, K. Engelke, and W. A. Kalender. Interactive 3D editing tools for image segmentation. *Medical Image Analysis*, 8(1):35–46, 2004.
5. M. Kass, A. Witkin, and D. Terzopoulos. Snakes: Active contour models. *Int'l J on Computer Vision*, 1(4):321–331, 1987.
6. T. F. Chan and L. A. Vese. Active contours without edges. *IEEE Trans. on Image Processing*, 10(2):266–277, 2001.
7. L. Cohen and R. Kimmel. Global minimum for active contour models: A minimal path approach. *International Journal of Computer Vision*, 24:57–78, 1997.
8. W. Barrett and E. Mortensen. Interactive live-wire boundary extraction. *Medical Image Analysis*, 1:331–341, 1997.
9. Y. Boykov and M. P. Jolly. Interactive graph cuts for optimal boundary and region segmentation of objects in n-D images. *Proc. ICCV*, pages 105–112, 2001.
10. L. Grady. Random walks for image segmentation. *IEEE TPAMI*, 28(11):1768–1783, 2006.
11. L. Grady. Multilabel random walker image segmentation using prior models. *IEEE Conf. CVPR*, 1:763–770, June 2005.
12. S. Andrews, G. Hamarneh, and A. Saad. Fast random walker with priors using precomputation. Technical Report TR 2010-07, Simon Fraser University, June 2010.