

HIDDEN SEED RECONSTRUCTION FROM C-ARM IMAGES IN BRACHYTHERAPY

Ryan C. Kon, Ameet Kumar Jain, and Gabor Fichtinger

Johns Hopkins University
Engineering Research Center for Computer Integrated Surgical Systems and Technology
{jain, gabor}@cs.jhu.edu

ABSTRACT

There has been a pressing clinical need for adaptive intra-operative dosimetry in the delivery of prostate brachytherapy implants. The missing prerequisite is the robust matching of the seeds across multiple C-arm images. This is further aggravated since seeds are invariably hidden in each image. We present a solution to recover these hidden seeds in this paper. A network flow formulation of the problem is proposed, where the desired solution is obtained (in polynomial time) by computing the flow with minimum cost. Phantom experiments show that using four X-ray images, on an average 99.8% of the seeds are recovered correctly, while simulations indicate that our algorithm is robust to segmentation errors of up to 1 mm and hidden seed rate of at least 8%. The results show strong feasibility and clinical data collection is currently underway.

1. MOTIVATION AND BACKGROUND

With an approximate annual incidence of 220,000 new cases and 33,000 deaths, prostate cancer continues to be the most common cancer in men in the United States. The definitive treatment modality for low risk prostate cancer is permanent brachytherapy, which is performed on approximately 40,000 patients each year. In this treatment a large number of small ($\sim 1 \times 5$ mm) radioactive capsules are implanted into the prostate to kill the cancer by emitting radiation. According to a comprehensive review by the American Brachytherapy Society [1], *the preplanned technique used for permanent prostate brachytherapy has limitations that may be overcome by intraoperative planning*. However, the report continues, *the major current limitation of intraoperative planning is the inability to localize the seeds in relation to the prostate*.

Brachytherapy is typically performed transperineally under real-time transrectal ultrasound (TRUS) guidance and C-arm fluoroscopy is often used for gross visual observation of the implant. TRUS usually provides adequate imaging of soft tissue anatomy, but it fails to visualize implanted seeds,

while C-arm fluoroscopy can visualize seeds, but not soft tissues. By reconstructing the implanted seeds from C-arm fluoroscopy and registering them to ultrasound, intra-operative dosimetry becomes possible.

3D coordinates of the implanted seeds can be calculated by resolving the correspondence of seeds across multiple X-ray images. Five major obstacles need to be overcome: (a) C-arm calibration; (b) C-arm pose tracking; (c) Seed segmentation; (d) Seed matching and reconstruction; and (e) Registration between C-arm and TRUS. While adequate techniques are available for most of these problems [2] and several groups have published results supporting C-arm fluoroscopy for intra-operative dosimetry [3], this technique has yet to become a standard of care in hospitals. The last technological barrier appears to be robust matching and reconstruction of seeds across multiple C-arm images. This problem is still unsolved because implants are usually dense and all the seeds are hardly ever visible in any one image. Even in the case of hand segmentation, perfectly overlapping seeds cannot be segmented ($\sim 2 - 6\%$). This phenomenon is often referred to as the issue of "hidden seeds" and is the focus of this paper.




Fig. 1. Example of a hidden seed not detectable by the human eye (right) matching to two distinct seeds (left).

Many contemporary works have made a simplifying assumption that all the seeds are visible, which makes these algorithms infeasible for clinical use. Among previous works, Fast-CARS was extended to incorporate hidden seeds, but the new algorithm reconstructed a greater number of seeds than were actually present [4]. Another variant was proposed [5] by ordering the seeds using the epipolar constraints. Unfortunately, the algorithm required co-planar images (co-linear X-ray sources) and could not reconstruct undetected seeds if they existed in the same search restriction band and did not

THIS WORK HAS BEEN FINANCIALLY SUPPORTED BY NIH 1R43CA099374-01, NSF EEC-9731478, AND DOD PC050170.

extend to multiple images. An intensity-based method using tomosynthesis [6] and using Hough trajectories [7] has also been proposed. However they require an unfeasibly large number of images to achieve a stable reconstruction, yet do not offer accuracies better than $1mm$. Another technique [8] optimized on seed positions and camera parameters, by generating simulated images and iterating them until they match the observed images. This optimization method is prone to fall into local minima and was tested only on clean simulated images. Significant works as they are, the problem merits further research to produce a clinically viable solution.

We have previously proposed matching and reconstruction of brachytherapy seeds using the Hungarian algorithm (MARSHAL) [9]. In this paper, we extend MARSHAL to also tackle the hidden seed problem.

2. METHODS AND MATERIALS

We assume that the seeds are 3D points and convert the seed-matching problem to a network-flow-based combinatorial global optimization. In this formulation, any correspondence of the seeds is represented by an appropriate flow through the network. The goal is to find the flow with the minimum cost. In the case with no hidden seeds, two images do not have a unique solution, while using three or more images makes the problem NP-hard [9]. Since the case with hidden seeds is a generalization of the earlier problem, it is also NP-hard. This proves that a globally optimal polynomial time algorithm is not possible.

In contrast to the previously proposed heuristic methodologies, MARSHAL solves the problem in a more mathematically rigorous framework of combinatorial optimization. This formal approach allows better control of the behavior of the algorithm, as well as consideration of seed groups as a whole (global optimization) instead of analyzing seed groups with heuristic rules (local optimization).

A Network-flow-based Formulation: Let N be the number of seeds inserted and N_1 , N_2 , and N_3 be the number of seeds actually segmented in the acquired C-arm images I_1 , I_2 , and I_3 . Let s_{ij} be the position of the i^{th} seed in the j^{th} image. We construct a directed network as shown in Fig 2.

Sets A, B, C, and D, each containing N_1 , N_2 , N_3 , and N_1 , respectively, represent images I_1 , I_2 , I_3 , and I_1 , respectively. There are no edges between nodes of the same set, but there are directed edges (links) connecting every node in set A (left) to set B, set B to set C, and set C to set A (right). There are also links from the source S to every node in set A (left), and similarly from set A (right) to the sink T. A flow of value N originates at S and ends at T. Because any of the sets can contain fewer than N nodes, each link from S to A or from A to T must allow for a flow greater than 1. A flow of 2 at a node implies that this seed is a hidden seed. To enforce the constraint that each seed in image 2 & 3 is chosen at least

once but not more than twice, we add dummy internal links. Thus, we allow a maximum/minimum flow of 2/1 through each of the internal links and through nodes connected to the source and sink. The links connecting the sets allow for a maximum/minimum flow of 1/0.




Fig. 2. The flow network used to solve the matching problem.

The problem now reduces to efficiently computing a flow of N . To determine the optimal solution, we need to assign a cost C_{ij} to the link connecting seed s_{i1} to seed s_{j2} . C_{ij} represents the likelihood of seed s_{i1} matching seed s_{j2} , with the cost being 0 if they match perfectly and ∞ (infinity) if they do not match at all, with extensive discussions available [9]. Any feasible flow in the network has a net cost of $\sum C_i f_i$, where f_i is the flow in link i and C_i is the cost of sending a unit flow along that link. Thus the seed-matching problem is reduced to finding the flow with minimum cost.

The minimum-cost flow can be computed using many polynomial and pseudo-polynomial algorithms [10]. We implemented the cycle cancelling algorithm to compute min-cost flow, which starts with a feasible flow and then searches for negative cost cycles in the residual network and pushing a flow along the negative cycle. This is iterated until there is no negative cost cycle detected, at which point the minimum cost flow has been computed. Negative cost cycles can easily be computed using the Bellman-Ford algorithm, which runs in $O(Edges * Vertices)$ time. Thus the run-time for the min-cost flow evaluation is $O(cost\ of\ the\ initial\ flow) \times O(Edges * Vertices)$, which will be pseudo-polynomial. Although there are strongly polynomial time algorithms available, the cycle-cancelling algorithm proves to be sufficiently fast.

Overview of the Algorithm: Extended MARSHAL is illustrated in Fig 3. Since the problem is NP-hard the network is not able to constrain the hidden seeds in image 1. Nevertheless, hidden seeds in image 2 & 3 are very well constrained. Thus the min-cost flow nearly solves the problem, except for a few cases which exhibit some ambiguity. These cases usually incarnate themselves as self-consistent complete subsets of 2-5 seeds. Each of these subsets can be independently solved using an extremely fast brute-force type of algorithm.

In the case that the size of any subset is large, the subset can be resolved by constructing a miniature network and solving for the min-cost flow. It is unusual to encounter small subsets which have hidden seeds in all three images, and hence

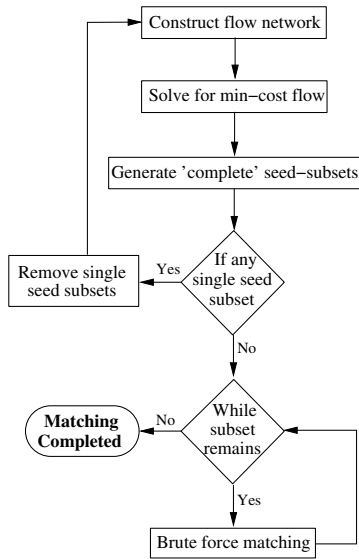


Fig. 3. The flowchart for extended MARSHAL.

these subsets can be resolved correctly by switching the order of the images. In the rare case that all three images in the subset have hidden seeds, then a choice of set A can be made arbitrarily and the network solved. Seeds solved as hidden in sets B/C can then be replicated and used as the outer set and the network solved again.

The robustness of the reconstruction can be significantly improved by taking another X-ray image, and thus we extend the algorithm to work with any number of images. The problem is well conditioned due to the epipolar constraints, so when four images are available, three images can be chosen as the main images, while one can be chosen as an assisting image (one with the most number of hidden seeds) affecting only the costs in the main network. This approach can be extended to any number of images. Thus extended MARSHAL always recovers the hidden seeds and resolves the correspondences in polynomial time (typically close to $O(N^3)$ [9]).

3. RESULTS AND DISCUSSION

Simulations: Data was generated to simulate a 55 cc prostate with a seed density of 2.0 seeds/cc. The algorithm was run on three different datasets using combinations of four images, varying segmentation error from 0 – 2 mm in increments of 0.25 mm. To test the sensitivity to the number of hidden seeds, the algorithm was also run varying the hidden seed percentage from 0 – 20% in increments of 2%. Averaged results are shown in Fig 4.

Phantom experiments: Experiments were conducted on a precisely fabricated acetol seed phantom. The FTRAC fiducial [2] was used to track the C-arm, and was attached to the phantom (Fig 5). The phantom comprises of twelve 5 mm




Fig. 4. (a) Sensitivity to segmentation error. (b) Sensitivity to hidden seed percentage.

thick slabs, each having one hundred holes with 5 mm spacing. Known implant constellations were created, with the number of seeds ranging from 40 to 100 in increments of 15, while keeping the seed density at 1.56 seeds/cc.

For a given constellation, 6 images within a 20° cone around the AP-axis were taken using an uncalibrated Philips Integris V3000 and dewarped using the pin-cushion test. Accurate ground truth for matching was computed from the known 3D seed locations. Matching was achieved with three and four images. All seeds closer than 1.2 mm were called hidden. Averaged results are displayed in Table 1.




Fig. 5. (a) An image of the phantom attached to the FTRAC fiducial. (b) A typical X-ray image of the combination.

Discussion: The simulation results show that the algorithm can nearly perfectly match all the seeds even with segmentation errors up to 1 mm. While increasing the segmentation error further decreases the matching percentage, at 2 mm error, the matching rate is still over 90%, with the reconstruction error of the mismatched seeds remaining below 2 mm. When the percentage of hidden seeds is varied, the algorithm can robustly match when the hidden seed percentage is at least 8%. Because the datasets for this case were generated by creating a threshold for closeness based only on the hidden percentage, significant error was introduced as the percentage increased, since the threshold became unrealistically large, creating much lower matching percentages than would normally be seen.

For the phantom data, using three images gives a good matching rate, but mismatched seeds reconstruct with a high error. Using four images gives substantially better results,

		Number of Seeds											
		3 Images					4 Images						
		40	55	70	85	100	40	55	70	85	100		
Matching Rate (%)	Match	98.3	100	99.9	98.2	98.0		100	100	100	99.9	99.5	
		Reconstruction Error (mean)	0.64	0.50	0.67	0.71	0.77		0.64	0.39	0.63	0.76	0.82
		Reconstruction Error (STD)	0.27	0.27	0.31	0.26	0.34		0.25	0.21	0.29	0.27	0.28
Reconstruction Error (mean)	Mismatch	7.46	-	0.80	12.47	4.99		-	-	-	14.13	1.73	
		Reconstruction Error (worst)	9.02	-	0.80	17.34	8.38		-	-	-	14.13	1.89
Reconstruction Error (relative)	All	0.52	0.32	0.40	0.54	0.53		0.32	0.30	0.38	0.34	0.41	

Table 1. Performance on phantom data.

with nearly perfect matching and a mostly low reconstruction error for mismatched seeds. Since it is readily feasible to obtain a fourth image in a clinical setting, our implementation is extremely viable for intra-operative dosimetry.

Conclusions and Future Work: In contrast to other proposed methods, we have formalized the seed matching problem and have extended a previously proposed polynomial time algorithm (MARSHAL) to resolve hidden seeds. A MATLAB 7 implementation runs in under 20 s in a typical implant using any number of images. Using 4 images, it matched over 99.8% of the seeds. Simulations indicate that MARSHAL is robust to various parameters. It can reconstruct an implant when three or more images are used, with a robustness, precision, and speed that promises to be sufficient to support intraoperative dosimetry in prostate brachytherapy.

MARSHAL is being combinatorially improved to increase the matching rate when only three images are used and also to reject spuriously segmented seeds. Moreover, other sources of information like seed orientation are being added to the cost metric to increase robustness. Clinical data is also being collected for further analysis.

4. REFERENCES

- [1] S Nag et al, "Intraoperative planning and evaluation of permanent prostate brachytherapy: report of the american brachytherapy society," *Int J Radiat Oncol Biol Phys*, vol. 51, no. 5, pp. 1422–30, Dec 2001.
- [2] A Jain et al, "A robust fluoroscope tracking (frac) fiducial," *Med. Phys.*, vol. 32, no. 10, pp. 3185–98, Oct 2005.
- [3] D Todor et al, "Intraoperative dynamic dosimetry for prostate implants," *Phys Med Biol*, vol. 48(9), pp. 1153–71, May 7 2003.
- [4] Y Su et al, "Prostate brachytherapy seed localization by analysis of multiple projections: identifying and addressing the seed overlap problem," *Med Phys*, vol. 31(5), pp. 1277–87, May 2004.
- [5] S Narayanan et al, "Three-dimensional seed reconstruction from an incomplete data set for prostate brachytherapy," *Phys Med Biol*, vol. 49(15), pp. 3483–94, Aug 2004.
- [6] I Tutar et al, "Tomosynthesis-based localization of radioactive seeds in prostate brachytherapy," *Med Phys*, vol. 30(12), pp. 3135–42, Dec 2003.
- [7] S Lam et al, "Three-dimensional seed reconstruction for prostate brachytherapy using hough trajectories," *Phys Med Biol*, vol. 49(4), pp. 557–69, Feb 2004.
- [8] M Murphy et al, "Demonstration of a forward iterative method to reconstruct brachytherapy seed configurations from x-ray projections," *Phys Med Biol*, vol. 50, pp. 2715–37, Jun 2005.
- [9] A Jain et al, "Matching and reconstruction of brachytherapy seeds using the hungarian algorithm (marshal)," *Med Phys*, vol. 32, no. 11, pp. 3475–92, Nov 2005.
- [10] R Ahuja et al, *Network Flows: Theory, Algorithms, and Applications*, Prentice Hall; 1 edition, 1993.