

# A Shape Relationship Descriptor for Radiation Therapy Planning\*

Michael Kazhdan<sup>1</sup>, Patricio Simari<sup>1</sup>, Todd McNutt<sup>2</sup>, Binbin Wu<sup>2</sup>,  
Robert Jacques<sup>3</sup>, Ming Chuang<sup>1</sup>, and Russell Taylor<sup>1</sup>

<sup>1</sup> Department of Computer Science, Johns Hopkins University, USA

<sup>2</sup> Department of Radiation Oncology and Molecular Radiation Science,  
Johns Hopkins University, USA

<sup>3</sup> Department of Biomedical Engineering, Johns Hopkins University, USA

**Abstract.** In this paper we address the challenge of matching patient geometry to facilitate the design of patient treatment plans in radiotherapy. To this end we propose a novel shape descriptor, the Overlap Volume Histogram, which provides a rotation and translation invariant representation of a patient's organs at risk relative to the tumor volume. Using our descriptor, it is possible to accurately identify database patients with similar constellations of organ and tumor geometries, enabling the transfer of treatment plans between patients with similar geometries. We demonstrate the utility of our method for such tasks by outperforming state of the art shape descriptors in the retrieval of patients with similar treatment plans. We also preliminarily show its potential as a quality control tool by demonstrating how it is used to identify an organ at risk whose dose can be significantly reduced.

## 1 Introduction

In the treatment of patients with malignant tumors, the goal of intensity-modulated radiation therapy (IMRT) is to deliver a high dose of radiation to the tumor volume while sparing nearby organs at risk (OAR). In practice, a pre-treatment computed tomography (CT) scan of the patient is segmented to identify the tumor volume and OAR. The segmented scan is then used by a dosimetrist to guide the settings for each multi-leaf collimator (MLC) on a radiotherapy machine targeting the tumor. This step is referred to as IMRT planning.

The design of a high-quality IMRT plan is one of the most time-consuming and least automated steps of the treatment cycle. The dosimetrist must optimize the MLC settings to achieve a dose distribution that most closely meets a set of physician-driven constraints. For example, in a commonly used treatment protocol for head-and-neck cancer [1] at least 95% of the tumor volume should receive

---

\* This research was supported in part by the generosity of Paul Maritz, Philips Radiation Oncology Systems (Madison, WI), and by Johns Hopkins University internal funds.

a dose of at least 70 Gy, while no more than 50% of each parotid should receive more than 30 Gy, and no part of the spinal cord should receive more than 45 Gy. The objective function is non-convex and requires that the dosimetrist perform multiple refinement optimization steps in order to best meet the constraints. Furthermore, computing the dose distribution from a set of MLC settings requires a complex simulation involving the inhomogeneous attenuation and scattering properties within the patient volume, making each individual refinement step computationally expensive. As a result, the process of planning an individual patient usually takes many hours to complete, often resulting in a time-lag of several days between the time that a patient comes in for scanning and when the patient can return for treatment.


We will argue that through shape matching using an appropriate shape relationship descriptor, both the speed and quality of the treatment planning process can be increased. Using a database of previously treated patients, the segmented geometry of the new patient serves as a query into the database, returning the treated patients whose configurations of tumor volume to OAR most closely resemble those of the new patient. Using the treatment plans of these most similar patients, we can facilitate the treatment of new patients, either by directly suggesting a treatment plan for the new patient or by using the retrieved plans as seeds in the optimization. The key challenge in designing such a system is the definition of a shape descriptor that captures not only the geometries of the tumor volume and OAR, but also their configurations relative to each other. Intuitively, the closer an OAR is to the tumor, the harder it is to irradiate the tumor while sparing the OAR.

To this end, we introduce the *Overlap Volume Histogram* (OVH). For each OAR, this 1D histogram describes the distribution of distances of the organ’s volume relative to the tumor volume. Since the “spareability” of an organ strongly depends on its proximity to the irradiated tumor, this descriptor provides a simple shape signature that is well-suited to the challenge of treatment plan retrieval. In our results we compare the OVH to several existing state-of-the-art shape descriptors and show significantly improved performance in retrieving patients with similar treatment plans.

## 2 Related Work

Due to the large size of the databases and guided by the need for real-time performance, many recent techniques for shape retrieval have taken the *shape descriptor* approach. The goal has been to obtain a concise, robust, and discriminating abstraction of the shape geometry that is well-suited for efficient querying. Most typically, the shape descriptor is represented as a vector in a high-dimensional space, dissimilarity of two shapes is measured as the distance between the corresponding descriptors, and database retrieval reduces to finding nearest neighbors in a high-dimensional space.

In whole-object retrieval the challenge is to find the shapes in the database which, up to transformation, are most similar to the query. Since the entire shape



**Fig. 1.** Visualization of the tumor volume (black) and several OAR (left) and the corresponding OVHs in differential (center) and cumulative (right) forms

is used for retrieval, normalization techniques can be used to remove much of the transformational ambiguity in matching, allowing for the use of the center of mass for removing translational ambiguity, radial-variance or mean-/bounding-radius for removing scaling ambiguity, and principal axes for rotational ambiguity. These methods have included: 1D histograms capturing the distribution of points [2,3,4], crease angles [5], and curvature [6] over the surface; spherical functions characterizing the distribution of surface normals [7], axes of reflective symmetry [8], conformality [9], and angular extent [10]; 3D functions characterizing the rasterization of the boundary boundary points [11] and the distance transform [12]; and even 4D plenoptic functions characterizing the 2D views of a surface [13].

Partial object retrieval is more difficult. The goal is to retrieve shapes containing sub-regions that are a good match to a subset of the query. Because it is not known in advance which subsets of the shapes will be used for matching, global normalization techniques cannot be applied. Instead, partial matching methods represent a shape by a multitude of shape descriptors, each centered at a different point on the shape’s boundary and characterizing only the subset of the shape in its vicinity. Using these descriptors, retrieval can be performed by rating target models in terms of the number and quality of matches between the descriptors of the query and the descriptors of the target. Commonly used shape descriptors for partial shape matching have included spin images [14,15], shape contexts [16], and curvature maps [17].

What makes patient retrieval difficult is that the notion of shape similarity required for retrieving treatment plans is not one characterized by either of these two methodologies. The geometric relationship between the tumor and the surrounding anatomy is more important than the detailed shapes of individual structures. The volume of overlap has been suggested for characterizing the relationship between OAR and target [18]. However, this only provides a single value and becomes uninformative when the OAR and target do not overlap.

### 3 The Overlap Volume Histogram

In designing a shape descriptor, our goal is to capture the proximity of the different OARs to the tumor volume. To this end, we define the *Overlap Volume Histogram*

**Algorithm 1.** GetOVH( *BinaryVoxel*  $O$ , *BinaryVoxel*  $T$  )

---

**Require:** Binary voxel grids for organ  $O$  and tumor  $T$   
*Histogram* dOVH  $\leftarrow 0$   
*Voxel* DT  $\leftarrow$  GetSignedDistanceTransform( $T$ )  
**for all**  $o \in O$  **do** dOVH[DT( $o$ )] += 1/| $O$ ||  
*Histogram* OVH  $\leftarrow$  cumulative(dOVH)  
**return** OVH

---

(OVH). This is a one-dimensional distribution associated to each organ at risk, measuring its distance from the tumor.

For a tumor  $T$  and organ  $O$ , the value of the OVH of  $O$  with respect to  $T$  at distance  $t$  is defined as the volume of the subset of the organ a distance of  $t$  or less from the tumor:


$$\text{OVH}_{O,T}(t) = \frac{|\{p \in O | d(p, T) \leq t\}|}{|O|},$$

where  $d(p, T)$  is the signed distance of  $p$  from the tumor's boundary and  $|O|$  is the volume of the OAR.

In practice, the OVH of an organ with respect to the tumor is efficiently computed from the segmented CT scans. Using the segmented tumor volume, we compute its signed Distance Transform. We iterate over the voxels interior to the organ and for each voxel, we evaluate the distance transform of the tumor, splatting a unit-volume vote into the associated bin. This gives us the differential of the volume which can then be used to compute the final cumulative form in a single pass. Since the signed distance transform can be computed in linear time (e.g. using [19]) the total running time of the algorithm is linear in the size of the CT scan. The algorithm is given in Algorithm 1.

An example of a patient's OVH descriptors is shown in Fig. 1. The image on the left shows the geometry of the tumor volume (black), spinal cord (red), brainstem (blue), and right and left parotid glands (dark and light green respectively). Examining the OVHs, we can quickly identify properties of the geometric configuration of the organs relative to the tumor. For example, the fact that the OVH of the left parotid has non-zero values at negative distances indicates that part of the parotid is overlapped by the tumor volume and therefore it will not be possible to spare the parotid in its entirety when fully irradiating the tumor. Similarly, since the the OVH values for both the spinal cord and brainstem are zero for distance values smaller than one centimeter, we know that no point on the tumor can be within a centimeter of these organs. Therefore, a treatment plan keeping most of the radiation within a centimeter of the tumor is likely to spare them.

To use the OVH for retrieval we need to define a metric on the space of descriptors. Since the differentials of the OVH correspond to distributions, a natural metric is the Earth Mover's Distance (EMD) which measures the amount of work needed to transform one distribution into the other. Although computing this distance in general is computationally difficult, requiring the solution of a



**Fig. 2.** DVHs for the treatment plan of the patient in Fig. 1

bipartite graph matching problem, it is straightforward to compute in our case. As shown by Rubner *et al.* [20], in the case that the distributions are normalized and one-dimensional, the EMD between two distributions is the  $L_1$ -distance between their cumulative histograms. As a result, run-time retrieval only requires finding the target descriptor(s) in the database minimizing the  $L_1$ -distance to the query.

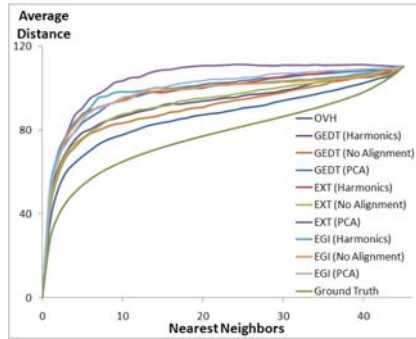
## 4 Experimental Results

To evaluate the effectiveness of a shape descriptor in patient retrieval, we measure the accuracy with which similarity of descriptors predicts similarity of treatment plans. The challenge in implementing such an experiment is that there does not exist a canonical metric for measuring plan similarity.

After a treatment plan is designed by a dosimetrist, a simulation of the IMRT is performed to determine the resulting dose distribution. In practice, the plan quality is evaluated by considering the *dose-volume histograms* (DVHs) [21] of the different organs and target volumes. These are 1D distributions, whose value at a specific dose is the volume of the organ or tumor that would receive at most that much dose under the proposed plan.

Fig. 2 shows the DVHs derived from the treatment plan for the patient shown in Fig. 1. Since the goal of the treatment is to kill the tumor, the plan results in a DVH for the tumor that has large values for all doses. For serial organs like the spinal cord and brainstem, the goal of the treatment is to ensure that no part of the organ receives a high dose, and the DVHs for both have zero value beyond 50 Gy. Since the parotids are parallel organs that remain functional even after a noticeable fraction of their volume has received high dose, the DVHs for both the left and right parotids show small volumes of the organ receiving doses larger than 60 Gy. Additionally, since the proximity of the left parotid to the tumor makes it hard to spare, the treatment results in more irradiation of this gland, with 10% of the organ receiving as much as 70 Gy.

We evaluate the quality of a descriptor by measuring the effectiveness with which it retrieves patients having similar DVHs. We do this by calculating the variation in DVH distances between a patient and the patient's  $k$  nearest neighbors (sorted by descriptor similarity).



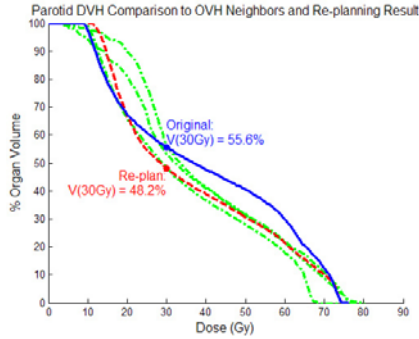
**Fig. 3.** Plots of the avg. distances from a patient’s DVH to the DVHs of its  $k$ -nearest neighbors as defined using several different shape descriptors

Given a patient, we compute the sum of squared  $L_1$ -distances from the DVH of patient to the DVHs of the patient’s  $k$ -nearest neighbors, summing over both the nearest neighbors and the different OARs. We repeat the experiment for all patients in the database and average the (root of the) sums. This gives a 1D distribution of the expected distance of a patient’s DVH from the DVH of its  $k$  nearest-neighbors. In general, we expect descriptors that better predict DVH similarity to give rise to distributions with smaller expected distances. Clearly, the best results are obtained when patients are sorted based on DVH similarity.

Fig. 3 compares the distance distribution obtained with our OVH descriptor with the distributions obtained using several other common shape descriptors, including Extended Gaussian Images (EGI) [7], angular extent functions (EXT) [10], and Gaussian Euclidean Distance Transforms (GEDT) [12]. As a baseline, the figure also shows the results when DVH similarity is used to sort the patients, (“Ground Truth” plot).

For each competing method, we obtained a representation of the organ-tumor relationship by computing the descriptor of the union of the organ and tumor geometries. We also addressed the problem of rotational alignment in three ways. **Harmonic:** We made the descriptors rotation invariant by storing only the magnitudes of the spherical frequencies. **PCA:** We used PCA to align each patient into a canonical pose prior to the computation of the descriptor. And, **No Alignment:** We used the given poses to allow for the possibility of consistent alignment across patients.

As the plots indicate, despite the quality of the traditional descriptors in general shape matching tasks, they are not well-suited patient retrieval, where it is the inter-organ relations that determine similarity. In contrast, our OVH descriptor is specifically designed to capture the proximity of different geometries. Thus it outperforms these methods, more often successfully using organ geometry to identify patients with similar treatment plans. The distance to the ground truth curve implies that, in practice, similarity of patient geometry does not always result in similar treatment. In fact, we have found that in some instances, two patients with nearly identical anatomies receive plans of markedly differing



**Fig. 4.** Parotid DVH (blue) along with the DVHs of its 3 nearest OVH neighbors (dotted green) and the DVH obtained after re-planning (dashed red.) Note the reduction of  $V(30\text{Gy})$  from 55.6% to 48.2%.

quality. This is because, given the difficulty of the planning process, physicians can unwittingly approve suboptimal plans.

The OVH can be used as a tool in plan quality assessment by searching the database for similar patients with better plans. As an example, Fig. 4 illustrates a preliminary result in which a parotid’s IMRT plan (blue) was automatically flagged as a candidate for replanning because the three database parotids with most similar OVHs (green) received much lower doses. The DVH obtained after replanning is illustrated in dashed red. With the new plan, the  $V(30\text{Gy})$  value to the patient’s parotid was reduced from 55.6% to 48.2%, crossing the 50% spareability mark in the RTOG protocol. This replanning result is preliminary, since the focus of this paper is the introduction of the OVH and its shape-retrieval performance evaluation (Fig. 3). In a related work [22], we focus specifically on the application of our OVH descriptor to IMRT planning. In this study based on sorting the mid-points of OVHs and comparing to the order of DVHs to find outliers, we were able to improve 13 out of 32 patient treatment plans.

## 5 Conclusion and Future Work

We have introduced the OVH, a novel shape relationship descriptor focused on the characterization of the inter-spacial relationship of different geometries. We have shown it can be computed efficiently and have demonstrated its practical efficacy by showing that it outperforms traditional shape descriptors in retrieving patients with radiotherapy treatment plans similar to the query using geometry alone. We also showed its potential as a quality control tool by demonstrating how it was used to identify an OAR whose dose could be significantly reduced. Based on these encouraging preliminary results, we have begun to explore the use of more complex OVH features and search methods for their application in patient retrieval for radiation therapy.

## References

1. Eisbruch, A., Chao, K.C., Garden, A.: Phase I/II study of conformal and intensity modulated irradiation for oropharyngeal cancer (RTOG 0022). Radiation Therapy Oncology Group of the American College of Radiology (2004)
2. Ankerst, M., Kastenmüller, G., Kriegel, H.-P., Seidl, T.: 3D shape histograms for similarity search and classification in spatial databases. In: Güting, R.H., Papadias, D., Lochovsky, F.H. (eds.) SSD 1999. LNCS, vol. 1651, pp. 207–228. Springer, Heidelberg (1999)
3. Osada, R., Funkhouser, T., Chazelle, B., Dobkin, D.: Matching 3D models with shape distributions. In: IEEE International Conference on Shape Modeling and Applications, pp. 154–166 (2001)
4. Osada, R., Funkhouser, T., Chazelle, B., Dobkin, D.: Shape distributions. *Transactions on Graphics* 21(4), 807–832 (2002)
5. Besl, P.: Triangles as a primary representation. In: Hebert, M., Boulton, T., Gross, A., Ponce, J. (eds.) NSF-WS 1994 and ARPA-WS 1994. LNCS, vol. 994, pp. 191–206. Springer, Heidelberg (1995)
6. Zhang, J., Siddiqi, K., Macrini, D., Shokouf, A., Dickinson, S.: Retrieving articulated 3-D models using medial surfaces and their graph spectra. In: *Energy Minimization Methods in Computer Vision and Pattern Recognition*, pp. 285–300 (2005)
7. Horn, B.: Extended Gaussian images. In: *Proceedings of the IEEE*, vol. 72, pp. 1656–1678 (1984)
8. Kazhdan, M., Chazelle, B., Dobkin, D., Finkelstein, A., Funkhouser, T.: A reflective symmetry descriptor. In: *European Conference on Computer Vision*, pp. 642–656 (2002)
9. Wang, S., Wang, Y., Jin, M., Gu, X.D., Samarasinghe, D.: Conformal geometry and its applications on 3d shape matching, recognition, and stitching. *IEEE PAMI* 29(7) (2007)
10. Vranic, D., Saupe, D.: 3D model retrieval with spherical harmonics and moments. In: *DAGM Symposium on Pattern Recognition*, pp. 392–397 (2001)
11. Gain, J., Scott, J.: Fast polygon mesh querying by example. *SIGGRAPH Technical Sketches*, 241 (1999)
12. Funkhouser, T., Min, P., Kazhdan, M., Chen, J., Halderman, A., Dobkin, D., Jacobs, D.: A search engine for 3D models. *ACM Transactions on Graphics* 22(1), 83–105 (2003)
13. Chen, D., Tian, X., Shen, Y., Ouhyoung, M.: On visual similarity based 3D model retrieval. *Computer Graphics Forum* 22(3), 223–232 (2003)
14. Johnson, A., Hebert, M.: Efficient multiple model recognition in cluttered 3-D scenes. In: *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 671–677 (1998)
15. Johnson, A.E., Hebert, M.: Using spin-images for efficient multiple model recognition in cluttered 3-D scenes. *IEEE PAMI* 21(5), 433–449 (1999)
16. Frome, A., Huber, D., Kolluri, R., Bulow, T., Malik, J.: Recognizing objects in range data using regional point descriptors. In: *European Conference on Computer Vision*, pp. 224–237 (2004)
17. Gatzke, T., Grimm, C.: Curvature maps for local shape comparison. In: *IEEE International Conference on Shape Modeling and Applications*, pp. 244–253 (2005)
18. Hunt, M.A., Jackson, A., Narayana, A., Lee, N.: Geometric factors influencing dosimetric sparing of the parotid glands using IMRT. *International journal of radiation oncology, biology, physics* 66, 296–304 (2006)



19. Saito, T., Toriwaki, J.: New algorithms for euclidean distance transformation of an n-dimensional digitized picture with applications. *Pattern Recognition* 27(11), 1551–1565 (1994)
20. Rubner, Y., Tomasi, C., Guibas, L.: The earth mover’s distance as a metric for image retrieval. *International Journal of Computer Vision* 40, 99–121 (2000)
21. Drzymala, R., Brewster, L., Chu, J., Goitein, M., Harms, W., Urie, M.: Dose-volume histograms. *International Journal of Radiation Oncology, Biology and Physics* 21, 71–78 (1991)
22. Wu, B., Ricchetti, F., Sanguineti, G., Kazhdan, M., Simari, P., Chuang, M., Taylor, R., Jacques, R., McNutt, T.: Patient geometry-driven information retrieval for IMRT treatment plan quality control. *Medical Physic.* (in submission)