INTEREST POINTS FROM THE RADIAL MASS TRANSFORM

Paul B. Albee

Department of CS
Central Michigan University
Mount Pleasant, MI 48859
albee@cps.cmich.edu

George C. Stockman

Department of CSE
Michigan State University
East Lansing, MI 48824
stockman@cse.msu.edu

ABSTRACT

The radial mass transform (RMT) is defined to produce a rotation and translation invariant vector representation of the neighborhood structure of points. The RMT transforms 3D (or 2D) data sets into a 1D signal, where $m(\mathbf{p}, r)$ gives the total mass (or intensity) of sensed data at distance r from the point $\mathbf{p}$. A support vector machine can be trained on example signals to detect salient points in an entire image or volume. Results show the method to be effective in multiple applications. The method is computation intensive but highly parallelizable and feasible for high value data sets.

1. INTRODUCTION

Detection of salient, or interesting, points is a fundamental operation for image processing [1, 2, 3, 4] . Such points are needed for registration of images so that they can be overlaid for mosaicking, fusion, or change detection. These points are used for correspondence in stereo computation and are often used as the basis for aggregation of other structure in an image, such as corners, lines, or contours [5, 6].

The radial mass transform (RMT) characterizes each image neighborhood by a rotation and translation invariant vector. We demonstrated its use in interest point detection on phantom data, on microtomographic images of a bee stinger, and on a spinal MRI. We demonstrate the use of the RMT for interest point detection using a framework that can, in general, use any vector of image neighborhood features.

We use the term *pixel* generally to mean either a 3D volume element or 2D image element. Our actual applications and implementations use 3D volume images. Section 2 gives the definition of the radial mass transform and methods of computing it. Section 3 gives characteristics of the transform and demonstrates its use in detecting salient points. Performance is assessed in Section 4.

This work was supported in part by the U.S. Department of Energy under contract no. W-31-109-Eng-38.

2. DEFINITION OF THE RMT

The RMT is designed to provide a basis for extracting rotation and translation invariant features from a volume. It integrates the mass, or density, intensity, etc., at distance r from some arbitrary fixed point $\mathbf{p}$ in space. Let the mass or intensity at a 2D point $\mathbf{x}$ be $f(r, \theta)$, where (r, θ) are the polar coordinates of $\mathbf{x}$ relative to $\mathbf{p}$.

$$m(\mathbf{p}, r) = \int_0^{2\pi} f(r, \theta)\, r\, d\theta \tag{1}$$

For a 3D volume, we have another polar coordinate ϕ.

$$m(\mathbf{p}, r) = \int_{-\frac{\pi}{2}}^{+\frac{\pi}{2}} \int_0^{2\pi} f(r, \theta, \phi)\, r^2 \, \sin\phi \, d\theta \, d\phi \tag{2}$$

An intuitive abuse of notation is given below in Equation 3, where V denotes either a 3D volume or a 2D slice or image. Moreover, to reduce the variance and bias of downstream decisions, the RMT is normalized by the area of the shell or ring involved in the integral to scale the otherwise rapidly increasing values for mass as r increases.

$$m(\mathbf{p}, r) = \frac{\int_{\mathbf{x} \in V} [\, ||\mathbf{x} - \mathbf{p}|| \, \approx \, r] \, f(\mathbf{x}) \, dV}{\text{area}(r)} \tag{3}$$

The RMT captures a great deal of structure of the neighborhood $N_{\mathbf{p}}$ of a point and is an efficient representation for segmentation, clustering or determination of interest points. By transforming to 1D space, a large suite of efficient algorithms are available.

The RMT is related to the circular Radon transform and the circular Hough transform [7]. The circular Radon transform integrates mass under a circular (or spherical) impulse, yielding a scalar value, whereas the RMT produces a 1D signal for a continuum of impulses. The Hough transform distributes each unit of mass at $I[x, y]$ in the image space to the set of all circles $C[r, x_c, y_c]$ that pass through $\mathbf{p} = [x, y]$. The RMT for pixel $[x_c, y_c]$ can be computed from a Hough accumulator array [8] by collecting the mass for each radius r associated with $[x_c, y_c]$.

2.1. Radial Mass Transform Computation

The discrete RMT integrates voxels at a discrete set of radial distances from a central voxel. $(2r+1)^3$ voxels are input to yield an output vector of length $r+1$. The computation of each vector coordinate sums a unique discrete set of voxel masses. If the original volume has S slices, of R rows and C columns, then the transformed data will have $SRCD$ values, where D is the maximum radius of the transform. This transformed space must then be significantly culled via extraction of a representative set of salient *interest points*. Using a naive implementation of the RMT, $\frac{4}{3}\pi D^3$ elements are accumulated for each single voxel. The overall running time of this transform, assuming a stride of 1 in r, is $\mathrm{O}(D^3 \times S \times R \times C)$.

There is a simple way to compute a single RMT at a single point $\mathbf{p}_0$: visit each voxel in the neighborhood of $\mathbf{p}_0$, compute its distance d to $\mathbf{p}_0$ and contribute its mass to accumulator $rmt[d]$. This works well inside a bounding box of the neighborhood or if we want to transform only a small number of neighborhoods $\mathbf{p}_j$. Algorithm 1 is the algorithm that we used to compute the normalized RMT for any subset of voxels within a volume. For each radius r, a set of offsets $RMTlookup_r$ identifies the voxels in the shell at distance r whose mass will be integrated.

Algorithm 1: Discrete RMT, offset method.
Input: Volume V, Origin for transform p_0, Maximum radius $r_{\max}$
Output: Discrete radial mass transform RMT_{p_0}

$\mathrm{RMT}(V, p_o, r_{\max})$
for $r = 1$ **to** $r_{\max}$
  $\mathrm{RMT}_{p_0,r} = 0$
  foreach $\Delta p \in RMTlookup_r$
    $\mathrm{RMT}_{p_0,r} = \mathrm{RMT}_{p_0,r} + V_{p_0+\Delta p}$
  $\mathrm{RMT}_{p_0,r} = \mathrm{RMT}_{p_0,r}/\mathrm{area}(r)$
return RMT_{p_0}

The RMT can be computed for a set of voxels, as described above, for an entire volume, or for regions in a volume that have high texture as detected by an inexpensive texture measure. Computing the RMT for an entire volume or highly textured sub-volume is used when interest detection is done via a trainable classification system.

2.2. Finding interest points on a phantom

An example involves an interactive user defining interesting points of a synthetic object – a toy jack that has complicated limbs and their intersections (refer to Figure 1). An interactive user identified interesting points and non interesting points on only one octant of the object. The RMTs of these examples were fed as training data to a support vector machine (SVM). The SVM produced a classification procedure that was then used to classify every voxel of the original volume, resulting in the interest points highlighted in Figure 1. A single slice of the *classified volume* is shown in Figure 2. The intensity shown in the slice is actually a continuous value proportional to the distance of the RMT of that voxel to the decision boundary created by the SVM. This example gives good evidence of the rotational and translational invariance of our RMT implementation (a direct study is reported below), since training data was drawn only from one octant of the volume.

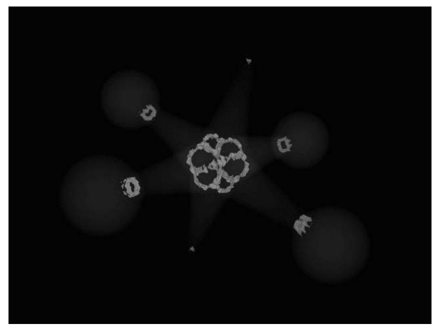

Fig. 1. Interest points detected by an SVM classifier trained on only one octant of positive and negative examples labeled by an interactive user. The original volume contained a synthetic image of a toy jack.

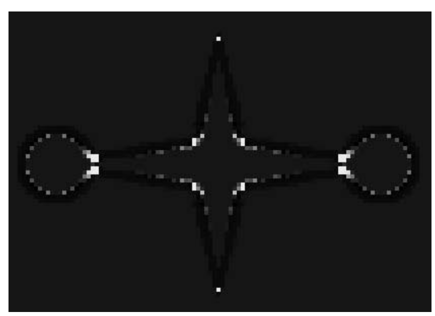

Fig. 2. One slice of interest points from Figure 1. Intensity of a voxel is proportional to the distance between the voxel RMT and the decision boundary computed by the SVM using the user supplied exemplars.

2.3. Features from the Discrete RMT

Additional features can be derived from the RMT: the RMT itself, RMT derivatives, and any other features, can be combined to determine the interest in a pixel neighborhood. Example RMT derived features are jaggedness, sum of squares of RMT, moments, etc. Due to space limitations, the reader is referred to [9] for details.

3. EXPERIMENTAL RESULTS

3.1. RMT Error

One source of error in the discrete RMT is the quantization of the spherical shells used to integrate the mass. Comparing the volume of the discrete RMT shells to the volume of continuous RMT shells of thickness one pixel, the largest relative error is 0.3 at radius $r = 1$ with an average relative error of 0.103 for $r = 1 \ldots 10$. Errors arise when the RMT includes non-uniform regions, since partially covered voxels contribute more or less than they should to the integral. The error does not appear to be a serious problem. The partial occupancy discretization error may be reduced but at a significant increase in computation cost.

3.2. Validation of RMT Rotational Invariance

To test how close the discrete RMT is to being rotationally invariant, a set of 421 cylindrical phantoms of radius 5 pixels were constructed at varied orientations. A second test was done using two sets of synthetic volumes that were generated from a real spinal MRI. The procedure for generating each phantom was to select a sub-region of the real volume and generate a series of rotated versions. The rotated volume was sampled to construct a new volume. For each generated volume, the RMT was computed and stored. For both types of synthetic volumes, the error in the RMT was computed using the Euclidian distance between the RMTs of matching points in each orientation and in the unrotated volume. The error results in Table 1 indicate that the largest error was 2.8%. Error in real data was about double that for the phantom cylinders.

Table 1. Results from RMT rotational invariance tests.

Data Set	Number of Orientations	Relative Error %		
		Mean	Std Dev	Max
Cylinder	441	1.17	0.21	1.43
Spine	8281	1.15	0.49	2.66
Vertebra	8281	1.64	0.65	2.82

3.3. Transformed Real Data

The process of detecting interest relies on the user identifying a set of training points for the classifier. In Figure 3a a cross section of a soil aggregate is shown where internal boundaries were considered interesting. A set of 14 interesting points and 13 uninteresting points were selected, shown in Figure 3b. Figure 4 shows the result of applying

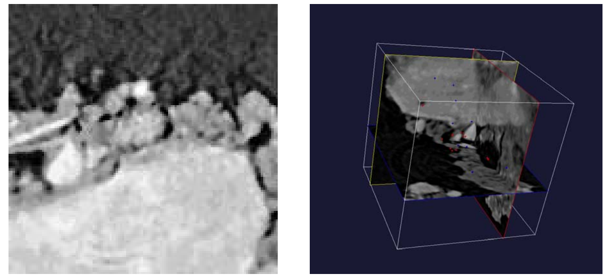

Fig. 3. Soil aggregate (a) cross section and (b) training point selection.

the interest classifier to the transformed version of the soil aggregate volume.

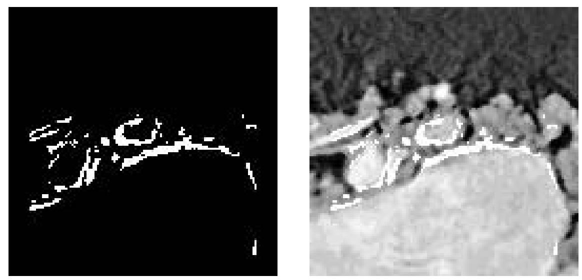

Fig. 4. Soil aggregate (a) detected interest points (b) overlain on original cross section.

Two of the many qualitative visual results appear in Figure 5, which shows good detection of the disks between vertebrae in the spine and good detail of the structure of a bee stinger.

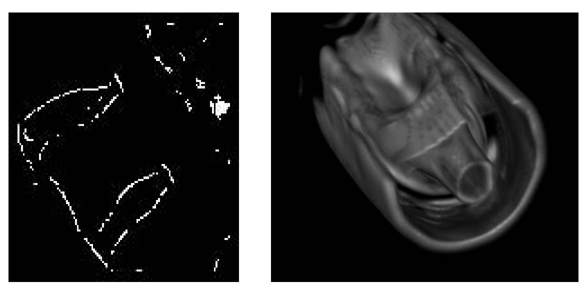

Fig. 5. Interesting points of (a) spine and (b) bee stinger.

4. PERFORMANCE

Similar to many scientific computations and visualizations, the computational cost of the RMT is high. One approach to minimizing the wall clock time is to compute the RMT in parallel and store it for later use, which we did by using 18 Sun Blade 100s networked with one Sun Ultra 10 and one Sun Fire V420.

The data sets ranged in size from 1.3MB (128^3) to 150MB (658 x 658 x 386). All of the phantoms had size 128^3 voxels. The computational effort for the RMT was constant for a given max radius ρ_{max} regardless of the underlying contents of the data.

Table 2. RMT Performance on phantom data using 18 hosts, $\rho_{max} = 10$. Times are given in Minutes:Seconds.

	CPU Time				Wall
	Volume			Slice	Time
Data Set	Load	Process	Store	Total	Total
ScutigeraSub	00:32	26:58	00:47	00:16	01:54
jack	00:31	27:00	00:47	00:16	01:53
spheres	00:31	26:59	00:47	00:16	02:00
uniformVol	00:32	26:59	00:47	00:16	02:00

Sample results for the real data sets are shown in Table 3. For data sets with the same slice dimension, the per-slice execution times were within 1 second and in general the execution times increased linearly in the number of voxels. Thus, it has been demonstrated that the RMT of a large volume can be computed on a network of simple machines within a couple of hours. Once RMTs have been computed and cached, they can be used by scientists in an interactive mode to explore the data. Input/output and graphical display, of course, remain as costly operations when the data set is large. Details are available in [9].

Table 3. RMT Performance on real data using 18 hosts, $\rho_{max} = 10$. Times are given in Hours:Minutes:Seconds.

	CPU Time				Wall
	Volume			Slice	Time
Data Set	Load	Process	Store	Total	Total
5atah1_a	0:22:49	39:03:25	1:41:59	0:08:29	2:34:34
5atah2_a	0:21:13	36:14:02	1:34:50	0:08:29	2:24:33
5atah3_a	0:18:04	30:51:50	1:20:38	0:08:29	2:07:28
5ath1_a	0:28:45	49:06:55	2:08:12	0:08:29	3:13:20
Scutigera	0:14:24	33:46:41	0:51:18	0:07:04	2:07:33
imp._bull_a	0:08:11	12:12:53	0:24:55	0:02:07	0:49:07
spine	0:02:21	5:28:50	0:08:23	0:07:04	0:23:10

5. DISCUSSION

Extensive experiments have been done with several kinds of data sets, but primarily microtomographic volumes from the Argonne National Laboratory. Results are very encouraging and work is continuing in using RMT products in the applications. One application is the analysis of the mechanics of the human spine, which depends on accurate segmentation of individual vertebra. Another application is the analysis of soil samples regarding the transport of carbon between pores in the soil aggregate and the atmosphere. The microtomographic volumes are very large, very noisy, and do not have a fixed calibrated range for the sensed material absorption. A preprocessing step has been used [9] to map (monotonically) widely ranging floating point density values into a fixed integer range $[0 \ldots k]$.

Interest points detected via the RMT and SVM classifier have been used in the automatic computation of a visual tour of the 3D volume. The goal of the visual tour is to produce a short smooth path of slices through the volume such that the slices reveal the most interest point stucture to the user. Our current software produces side-by-side view boxes, one displaying the original voxel values and one displaying the "interest value" of the voxels as produced by the SVM. Sample tours can be found on our website at `http://www.cps.cmich.edu/~albee/RMT/`.

6. REFERENCES

[1] C. Harris and M.J. Stephens, "A combined corner and edge detector," in *Alvey88*, 1988, pp. 147–152.

[2] Les Kitchen and Azriel Rosenfeld, "Gray-level corner detection," *Pattern Recognition Letters*, vol. 1, no. 2, pp. 95–102, December 1982.

[3] Karl Rohr, "On 3d differential operators for detecting point landmarks," *Image and Vision Computing*, vol. 3, no. 15, pp. 219–233, 1997.

[4] Karl Rohr, *Landmark-Based Image Analysis*, vol. 21 of *Computational Imaging and Vision*, Kluwer Academic Publishers, Dordrecht, February 2001.

[5] J. Canny, "Finding edges and lines in images," in *MIT AI-TR*, 1983.

[6] P. Hough, "Method and means for recognizing complex patterns," US Patent 3,069,654, 1962.

[7] Steven L. Tanimoto, *The Elements of Artificial Intelligence Using Common Lisp*, Computer Science Press, 1995.

[8] Linda G. Shapiro and George C. Stockman, *Computer Vision*, Prentice Hall, 2001.

[9] Paul Benjamin Albee, *Analysis and Visualization of Volumetric Data Sets*, Ph.D. thesis, Michigan State University, August 2004.