The Generalised Radon Transform: Sampling and Memory Considerations

C.L. Luengo Hendriks, M. van Ginkel, P.W. Verbeek and L.J. van Vliet

Pattern Recognition Group, Delft University of Technology,
Lorentzweg 1, 2628 CJ Delft, The Netherlands
{cris,michael,piet,lucas}@ph.tn.tudelft.nl

Abstract. The generalised Radon transform is a well-known tool for detecting parameterised shapes in an image. Applying the Radon transform to an image results in a parameter response function (PRF). Curves in the image become peaks in the PRF. The location of a peak corresponds to the parameters of a shape, and the amplitude to the amount of evidence for that shape. In this paper we discuss two important aspects of the Radon transform. The first aspect is discretisation. Using concepts from sampling theory we derive a set of sampling criteria for the Radon transform. The second aspect concerns a projection-based algorithm to reduce memory requirements.

1 The Radon transform

The (generalised) Radon transform is a technique for detecting parameterised shapes. Given a model of the shape, it defines a mapping from the image space onto a parameter space. The axes of the parameter space correspond to the parameters of the model. When applied to an image, the Radon transform yields a parameter response function (PRF) defined on the parameter space. A shape in the image becomes a peak in the PRF. The location of the peak corresponds to the parameters of the shape. Shape detection is thus reduced to peak detection. We discuss two aspects of the Radon transform: its discretisation and an algorithm to reduce its storage requirements. We focus on the Radon transform for (hyper-)spheres, but the discussion of the discretisation holds for arbitrary shapes. In its most general form, the Radon transform is

$$P(\boldsymbol{p}) = \int_{\boldsymbol{s} \text{ on } \boldsymbol{c}(\boldsymbol{p})} I(\boldsymbol{s}) d\boldsymbol{s}, \tag{1}$$

with $P(\boldsymbol{p})$ the PRF, $I(\boldsymbol{s})$ the image and $\boldsymbol{c}(\boldsymbol{p})$ the shape for parameter vector $\boldsymbol{p}$.

There are two common approaches to the discretisation of the Radon transform. The first is a straight-forward discretisation of the integral using standard numerical algorithms. It chooses a point in parameter space and computes its value by integrating the image over all the points belonging to the curve. In the second algorithm we choose a point in the image and add its contribution to all the appropriate points (or bins) of the PRF, a process known as voting. This

to be published in Proceedings of CAIP 2003, Groningen, The Netherlands,
by Springer-Verlag, Berlin, 2003.

approach is traditionally known as the Hough transform. We stress that the two are identical in the continuous domain [3]. An advantage of the Radon paradigm is that it gives us control over how we visit the points in parameter space. Later we use this property to reduce the memory requirements for storing the PRF.

2 Sampling the Radon transform

Several authors have addressed discretisation-related effects, in particular binning effects, of the Hough transform [1, 7] and references therein. Here, we pursue a different approach: we apply the principles of sampling theory to the Radon transform, thus trying to avoid discretisation errors altogether. Sampling theory gives us the conditions under which a signal can be sampled without loss of information. The same theory is also applicable to the discretisation of an arbitrary linear *operator*, including the Radon transform. We start by writing the Radon transform in the form of a general linear operator:

$$P(\boldsymbol{p}) = \int_{\mathbb{R}} C(\boldsymbol{p}, \boldsymbol{s}) I(\boldsymbol{s}) d\boldsymbol{s}. \tag{2}$$

This equation becomes a Radon transform by choosing the function $C(\boldsymbol{p}, \boldsymbol{s})$ appropriately. If a subset $\boldsymbol{x}$ of the parameters $\boldsymbol{p}$ represents the spatial position of the shape, then C has a special form:

$$C(\boldsymbol{p}, \boldsymbol{s}) = K(\{\boldsymbol{p} \setminus \boldsymbol{x}\}, \boldsymbol{x} - \boldsymbol{s}), \tag{3}$$

where $\{\boldsymbol{p} \setminus \boldsymbol{x}\}$ denotes all the parameters in $\boldsymbol{p}$ except those in $\boldsymbol{x}$. Along the dimensions $\boldsymbol{x}$ the integral reduces to a convolution [8, 9]. Using the convolution property of the Fourier transform, the operation reduces to a multiplication in the Fourier domain resulting in a huge speed-up.

Following [4], we will now investigate sampling criteria for equation (2). There are two aspects. Keeping $\boldsymbol{p}$ fixed, we will first consider under which conditions we may replace I and C by sampled (along $\boldsymbol{s}$) versions and the integral by a summation. If these conditions are satisfied, we may compute $P(\boldsymbol{p})$ for an arbitrarily chosen $\boldsymbol{p}$. We must then show that it is possible to sample the PRF $P(\boldsymbol{p})$, so that we only need to evaluate $P(\boldsymbol{p})$ on a discrete set of points. For simplicity we restrict ourselves to a one-dimensional example: $\boldsymbol{p} \to p$ and $\boldsymbol{s} \to s$. The Fourier axes corresponding to p and s are denoted by $\tilde{p}$ and $\tilde{s}$ respectively. The sampling distance along s is Δs. The discrete coordinate corresponding to s is n, i.e. the sampled version of I is $I(n\Delta s)$. We first investigate under which conditions the following is true:

$$P(p) = \int_{\mathbb{R}} C(p, s) I(s) ds = \Delta s \sum_{n \in \mathbb{Z}} C(p, n\Delta s) I(n\Delta s). \tag{4}$$

We denote the band-limit (along s) of the product $C(p, s)I(s)$ by $\mathrm{b}_s\{CI\}$. With p fixed, the sampling criterion for the computation of this integral, $\tilde{s}_s > \mathrm{b}_s\{CI\}$,

is a relaxed version of the Nyquist criterion [10]. Also the band-limit of CI can be expressed in that of C and I:

$$\mathrm{b}_s\{CI\} \leq \mathrm{b}_s\{C\} + \mathrm{b}_s\{I\}. \tag{5}$$

It follows that both the operator function C and the image I must be band-limited to allow discretisation. Proper sampling of the image I is a prerequisite for any image analysis and poses no specific problem. This is not true for the operator function C, which is not band-limited in general. We must impose a band-limit on C. This clearly leads to a different Radon transform, but this reflects a conscious choice with well-understood consequences. The alternative, sampling C without imposing a band-limit first, leads to aliasing effects.

We can compute $P(p)$ for an arbitrary value of p. If $P(p)$ is band-limited, we may sample $P(p)$ at the correct (Nyquist) rate. We determine whether $P(p)$ is band-limited by computing its Fourier transform:

$$\mathcal{F}\{P\}(\tilde{p}) = \mathcal{F}\left\{\int_{\mathbb{R}} C(p,s)I(s)ds\right\}(\tilde{p}) = \int_{\mathbb{R}} \mathcal{F}_p\{C(p,s)\}(\tilde{p},s)I(s)ds. \tag{6}$$

If C is band-limited along the p axis with band-limit $\mathrm{b}_p\{C\}$, then the integral above evaluates to zero for $\tilde{p} > \mathrm{b}_p\{C\}$, which means that $P(p)$ is band-limited as well.

The discussion above also holds for the complete multi-dimensional operation: our argument holds for each spatial dimension s_i separately and for each parameter dimension p_j as well. The same ideas also extend trivially to other sampling schemes, such as the hexagonal grid.

Kiryati and Bruckstein [6] have proposed band-limitation of the Hough transform. Their approach consists of replacing the sinusoids that are stamped in parameter space by band-limited versions (see also the references in [6]), in essence the same technique used by the Parzen estimator. In our formalism, this corresponds to imposing a band-limit on $C(\boldsymbol{p}, \boldsymbol{s})$ along the $\boldsymbol{p}$ axes. The difference between their and our approach lies mainly in the model for the input data: in their case a set of mathematical points in a continuous space, in our case a sampled image.

2.1 Band-limiting the operator function $\mathbf{C}(\boldsymbol{p}, \boldsymbol{s})$

The Gaussian filter is approximately band-limited with critical sample spacing σ [11] and corresponding band-limit $b = \frac{1}{2}\sigma^{-1}$. Its properties, in particular good simultaneous frequency and spatial localisation, make it a good choice for band-limiting $C(\boldsymbol{p}, \boldsymbol{s})$: We obtain C_b, a band-limited version of C, as follows:

$$C_b(\boldsymbol{p}, \boldsymbol{s}) = C(\boldsymbol{p}, \boldsymbol{s}) * G(\boldsymbol{p}, \boldsymbol{s}; \boldsymbol{\Sigma}). \tag{7}$$

The diagonal covariance matrix $\boldsymbol{\Sigma}$ reflects that we impose band-limitation along each dimension separately.

The function $C(\boldsymbol{p}, \boldsymbol{s})$ is in general very sparse. This follows directly from equation (1): for any given $\boldsymbol{p}$, the points $\boldsymbol{s}$ which belong to the shape span some

curve or manifold in $C(\boldsymbol{p}, \boldsymbol{s})$. The Radon transform for hyper-spheres provides a convenient example to investigate the structure of $C(\boldsymbol{p}, \boldsymbol{s})$ and the effects of band-limitation in some detail. The parameter vector $\boldsymbol{p}$ consists of the centre $\boldsymbol{x}$ of the D-dimensional sphere and its radius r: $\boldsymbol{p} = (x_1, \cdots, x_D, r)$. The operator function C becomes

$$C(\boldsymbol{p}, \boldsymbol{s}) = K(r, \boldsymbol{x} - \boldsymbol{s}) \quad \text{with} \quad K(r, \boldsymbol{\xi}) = \delta\big(\frac{1}{2}\sqrt{2}(\|\boldsymbol{\xi}\| - r)\big) \tag{8}$$

for a sphere. The function K represents a cone. If we consider a sufficiently small surface patch of the cone, we may consider it as a plane. Along the normal to this plane, the function K should behave like a Dirac delta. Hence the factor $\frac{1}{2}\sqrt{2}$ in (8). With this normalisation we have that $\int_0^R \int_{\mathbb{R}^D} K(r, \boldsymbol{\xi}) d\boldsymbol{\xi} dr$ equals the surface area of the truncated cone with a base of radius R.

What is the effect on K of the Gaussian smoothing applied to C? Let us first consider the effect of the smoothing applied along the $\boldsymbol{p}$ axes. All parameters share the same units and it is therefore logical to use the same σ_K along each dimension. The effect on a local surface patch, if it can be considered planar locally ($\sigma_K \ll r$), is that the Dirac profile is substituted by a Gaussian profile

$$K_b(r, \boldsymbol{x} - \boldsymbol{s}; \sigma) = K(r, \boldsymbol{x} - \boldsymbol{s}) *_{\boldsymbol{x}, r} G(\boldsymbol{x}, r; \sigma_K) \approx G(\frac{1}{2}\sqrt{2}(\|\boldsymbol{x} - \boldsymbol{s}\| - r); \sigma_K\big). \tag{9}$$

The next step is to apply a Gaussian along the $\boldsymbol{s}$ axes; but this is, in fact, unnecessary. The structure of K along axes $\boldsymbol{x}$ and $\boldsymbol{s}$ is not independent. The Gaussian smoothing along $\boldsymbol{x}$ implies a Gaussian smoothing along $\boldsymbol{s}$, as is evident from (9). It is unnecessary to apply the Gaussian smoothing along the $\boldsymbol{s}$ axes, unless the required smoothing σ_s along the $\boldsymbol{s}$ axes is larger than that required along the $\boldsymbol{p}$ axes.

The Gaussian smoothing has been chosen to allow a sampling distance of σ_K along the normal to the plane. The actual sampling will be along $\boldsymbol{\xi}$ and r. When using a rectangular sampling grid, the off-axis band-limit is larger than the on-axis band-limit. In the case of our cone, this means that we can reduce the size of the Gaussian by a factor of $\sqrt{2}$:

$$K_b(r, \boldsymbol{\xi}) = \sqrt{2}\, G(\|\boldsymbol{\xi}\| - r; \sigma_K). \tag{10}$$

The consequences of the imposed band-limitation are as follows: as long as the Gaussian is small with respect to the curvature of the manifolds represented by the operator function C, the effects of the Gaussian are negligible. In fact, it is possible to interpolate the PRF and obtain sub-pixel accuracy. High-curvature patches of C correspond either to highly curved shapes or to shapes which vary rapidly as a function of the parameters. In neither case is it reasonable to expect good results.

3 Reducing memory requirements

The parameter space for the Radon transform typically has more dimensions than the input image. This implies that the PRF might not fit into the available

computer memory. This constraint has traditionally prevented wide-spread use of these transforms for 3D images.

Many authors have tackled this problem in a variety of manners. Most notably, Ballard and Sabbah [2] propose to partition the parameter space into two or more spaces with independent parameters, which can be computed sequentially. Hsu and Huang [5] apply this method to detect 3D ellipsoids (with 6 parameters, the axes are supposed to lie on the grid). They split the 6D parameter space into two 4D parameter spaces, which have to be combined to find the objects.

Another method often employed involves splitting the parameter space into overlapping regions, from which the maxima are extracted. This does not involve a reduction of dimensionality, but incurs a penalty in computational cost because of the overlap. In the case of a sphere, it is natural to split the parameter space along the r-axis, since a slice $P(r_i, \boldsymbol{x})$ is computed by a single convolution. We will call this method the Sliding Window method (SW).

We propose a different approach to reduce the memory requirements. Spheres can be detected very efficiently by storing only the maximum projection along the r-axis of P, together with the location of these values on the r-axis (if one is prepared to ignore concentric spheres). That is, we keep

$$S(\boldsymbol{x}) = \max_r \{P(r, \boldsymbol{x})\} \quad \text{and} \quad R(\boldsymbol{x}) = \arg\max_r \{P(r, \boldsymbol{x})\}. \tag{11}$$

The local maxima in $S(\boldsymbol{x})$ indicate the location of the center of the spheres, and $R(\boldsymbol{x})$ gives the corresponding radii. Both of these can be computed by a small modification of the Radon algorithm. Instead of storing all the $P(r_i, \boldsymbol{x})$ slices, we propose to take the point-wise maximum of each slice with the previously computed intermediate result. This does not add any computational cost to the algorithm, since finding the local maxima needs to be done anyway. This maximum projection even simplifies this task. We call this method the Maximum Projection method (MP), and should be both faster and much less memory-hungry than the SW method.

The resulting PRF $S(\boldsymbol{x})$ is not band-limited. But, if the spheres are clearly identifiable and well separated, it turns out to have nicely-shaped peaks (i.e. the neighbourhoods of the local maxima are band-limited or nearly so). Thus, it is still possible to obtain the center of the sphere with sub-pixel accuracy. However, the r-axis at each location has been discretised to sampling locations. The accuracy to which r can be estimated depends on the number of samples taken, not the band-limitation of $K_b(r, \boldsymbol{x})$ along the r-axis.

It is possible to implement such a Radon transform for other shapes as well, in which the maximum projection can be taken over more than one dimension. That is, only the spatial dimensions need to be kept, all other dimensions can be collapsed into one maximum image and one or more maximum position images, of which there are as many as parameter dimensions are reduced.

4 Results

To demonstrate the claims made in the previous sections, we computed the Radon transform of 25 synthetically generated, 3D test images, 128^3 pixels in size, each containing 20 spheres of different radii (between 6 and 18 pixels) at random, sub-pixel locations. Some of the spheres were touching, but none were overlapping. These spheres had a Gaussian profile (with $\sigma = 1$), thereby approximating band-limitness. We computed the Radon transform with the two methods explained above (setting $\sigma_K = 1$ in equation (10)): SW and MP. The SW method uses a window of 7 slices in the radius direction, from which 2 slices overlap other regions. It required five times as much memory, and took three times as much time to finish, as the MP method. Apparently the local maxima algorithm we used is relatively expensive compared to the convolutions themselves. We evaluated both methods by computing the differences between the true parameters of the spheres and the computed ones. The average error and the standard deviation give a quantitative performance measure for the algorithm. These results are summarized in Table 1.

We found that both methods found the location of the spheres with the same accuracy (actually the parameters found for an individual sphere were very similar). The bias is very small, not significant in relation to the standard deviation. Both methods underestimate the radius in the same way, but the MP method found the rounded values of the radii found by the SW method. The underestimation of the radius depends on r, and it is possible to correct for it by increasing the radius of the functional $K_b(r, \boldsymbol{x})$.

We added Gaussian noise with standard deviation σ_N such that the signal-to-noise ratio $SNR = \frac{\max I(\boldsymbol{s})}{\sigma_N} = 2$ (with $I(\boldsymbol{s})$ the uncorrupted image). The standard deviation in the error of the spatial coordinate increases by about 50% for this noise level, but is still very small. This shows that the projection method is a good approximation with or without noise, and shows that the Radon transform itself is very insensitive to noise.

4.1 Ballotini

As a demonstration application, we used a rather poor quality X-ray micro-CT image of ballotini (small, hollow glass beads, see Fig. 1, the two images on the top left). Some of the glass walls give a very wide response in the imager (probably caused by refraction or reflection). In one such region many small spheres can be fitted. To avoid this, we replaced the kernel K_b by a kernel K'_b that penalizes for high grey-values inside the sphere:

$$K'_b(r, \boldsymbol{x}) \;=\; K_b(r, \boldsymbol{x}) - K_b(r - 4, \boldsymbol{x}), \tag{12}$$

with K_b the original kernel as given in Eq. (9).

By requiring that the inner part of the sphere be empty, the discriminating abilities of the transform (for these images) are greatly enhanced (see Fig. 1). The computational cost is increased minimally, since only generating the image

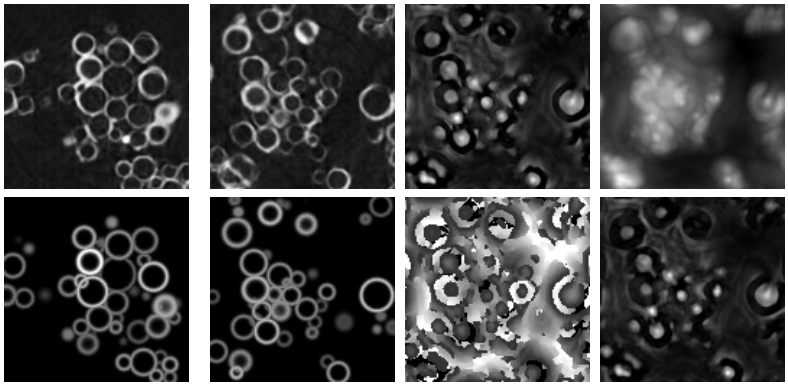

Fig. 1. Two slices of the 3D ballotini image and the results of the Radon transform. Top to bottom, left to right: Slice of the input image; corresponding slice of the image reconstructed with the found parameters. Another slice of the input image; corresponding slice of the output. Corresponding slice of $S(\boldsymbol{x})$; slice of $R(\boldsymbol{x})$. $S(\boldsymbol{x})$ without the inner sphere; $S(\boldsymbol{x})$ when the inner sphere has a diameter 2 pixels smaller than the outer sphere (instead of 4 as actually used).

for the functional $K'_b(r, \boldsymbol{x})$ is more expensive. The "magic number" 4 used in this functional was chosen such that the sphere $K_b(r, \boldsymbol{x})$ was not affected too much ($2 \cdot 2\sigma_K = 4$), since that would cause a heavier underestimation of the radius. In the synthetic test images used in above, this setting leads to an average underestimation of the radius of 0.3 pixels (v.s. 0.2 pixels for the transform with K_b as the kernel). As before, this systematic error can be corrected for.

To find the spheres in the PRF $S(\boldsymbol{x})$, a threshold is used to decide which local maxima are important enough to represent a sphere in the input images. More complex decision rules could be used, but are outside the scope of this paper. Figure 1 shows the results for two different slices from the 3D image.

5 Conclusions

We have given the conditions under which the Radon transform can be computed free of discretisation errors. In general these conditions must be imposed by actively band-limiting the operator function C. This has no consequences for sufficiently smooth shapes. The PRF that results is band-limited, allowing interpolation, and sub-pixel accuracy in the estimated parameters.

The Radon transform reduces to a convolution for position-type parameters, yielding a large speed-up. We propose a memory-efficient implementation computing a single r slice of $P(r, \boldsymbol{x})$ (through convolution) at a time. We keep track of the maximum projection and the argument-maximum projection along the r axis as we compute the slices. We argue that this approach can be used for other shapes as well.

We have applied this modified Radon transform to a 3D image of glass hollow beads. To compute its PRF we have employed a convolution kernel that contains

Table 1. Error made when estimating parameters of spheres in synthetic 3D images. The error in the position ($\delta x = \hat{x} - x$) and the error in the radius ($\delta r = \hat{r} - r$) are shown separately (the units are pixels for both). The error in the position considers the first spatial coordinate only.

	SNR = ∞		SNR = 2	
	MP Method	SW Method	MP Method	SW Method
$E(\delta x)$	−0.00130	−0.00127	−0.00356	−0.00385
$std(\delta x)$	0.02917	0.02938	0.04478	0.04605
$E(\delta r)$	−0.20706	−0.21850	−0.20906	−0.21986
$std(\delta r)$	0.30502	0.07142	0.30683	0.07326

not only a sphere, but also a second, smaller, concentric sphere with negative grey-values. The resulting PRF has a much higher discriminating ability than that which would result from the same computation with a single sphere.

Acknowledgements

The authors wish to thank Scott Singleton and Dave Rowlands at Unilever R&D Colworth (UK), for their hospitality and permission to use the ballotini image in this paper. This work was partially supported by the Dutch Ministry of Economic Affairs through their IOP program and by Unilever R&D Vlaardingen (NL).

References

1. A.S. Aguado, E. Montiel, and M.S. Nixon. Bias error analysis of the generalised Hough transform. *Journal of Mathematical Imaging and Vision*, 12(1):25–42, 2000.
2. D.H. Ballard and D. Sabbah. Viewer independent shape recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5(6):653–660, 1983.
3. S.R. Deans. Hough transform from the Radon transform. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3(2):185–188, March 1981.
4. M. van Ginkel. *Image Analysis using Orientation Space based on Steerable Filters*. PhD thesis, Delft University of Technology, Delft, The Netherlands, 2002.
5. C.-C. Hsu and J.S. Huang. Partitioned Hough transform for ellipsoid detection. *Pattern Recognition*, 23:275–282, 1990.
6. N. Kiryati and A.M. Bruckstein. Antialiasing the Hough transform. *CVGIP: Graphical Models and Image Processing*, 53(3):213–222, May 1991.
7. J. Princen, J. Illingworth, and J. Kittler. Hypothesis testing: a framework for analyzing and optimizing Hough transform performance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(4):329–341, April 1994.
8. J. Sklansky. On the Hough technique for curve detection. *IEEE Transactions on Computers*, 27(10):923–926, October 1978.
9. G.C. Stockman and A.K. Agrawala. Equivalence of Hough curve detection to template matching. *Communications of the ACM*, 20(11):820–822, 1977.
10. P.W. Verbeek. A class of sampling-error free measures in oversampled band-limited images. *Pattern Recognition Letters*, 3:287–292, 1985.
11. L.J. van Vliet. *Grey-Scale Measurements in Multi-Dimensional Digitized Images*. PhD thesis, Delft University of Technology, Delft, The Netherlands, October 1993.