

Subpixel Image Registration Using Circular Fiducials

ALON EFRAT*

*Department of Computer Science
Technion - Israel Institute of Technology
Haifa 32000, Israel*

and

CRAIG GOTSMAN

*Department of Computer Science
Technion - Israel Institute of Technology
Haifa 32000, Israel*

Received 23 July 1992

Revised 18 February 1994

Communicated by F. F. Yao

ABSTRACT

The design of fiducials for precise image registration is of major practical importance in computer vision, especially in automatic inspection applications. We analyze the subpixel registration accuracy that can, and cannot, be achieved by some rotation-invariant fiducials, and present and analyze efficient algorithms for the registration procedure. We rely on some old and new results from lattice geometry and number theory and efficient computational-geometric methods.

Keywords: Computer Vision, Registration, Fiducials, Lattice Geometry, Number Theory.

1. Introduction

The design of fiducials for accurate registration of images is a problem of major practical importance in computer vision. Imagine the following scenario: An electronic printed circuit board is to be automatically inspected for manufacturing defects. The system that does this is a computer equipped with a digital camera. A binary image of the board, obtained via the camera, is the input to the inspection algorithm, which then checks select areas of the image for the defects. The algorithm must first locate accurately these areas of the board in the image plane. To enable accurate location of points on the board, special patterns (“fiducials”) imprinted on the board are located in the image, and a select point in each fiducial serves as an “anchor” point, relative to which other points are referenced. Usually

*Current address: Dept. of Computer Science, School of Mathematical Sciences, Tel Aviv University, Ramat Aviv, Israel.

three fiducials suffice, as any other point in the image may be expressed uniquely in affine-invariant *barycentric* coordinates relative to these three anchor points (see Gotsman ¹ for a discussion on the optimal positions of these three points in the image plane). An error in the measured coordinates of any of the three anchor points induces an error in the location of any other image point (see Fig. 1), so it is crucial to locate the fiducials with the best accuracy possible. At first glance, it seems that positional location may be measured only up to the accuracy of physical pixel size. This is not necessarily true, as careful design of the fiducial pattern enables it, and consequently other points, to be positioned with subpixel accuracy. There is a tradeoff between the fiducial size (in pixel units), and the registration accuracy obtainable from the fiducial (also in pixel units). Bruckstein et al. ² prove an exponential lower bound $\Omega(2^{-n})$ on the accuracy obtainable by a fiducial pattern of n pixels, and design a pattern that achieves this bound. Unfortunately, their fiducial consists of rectilinear stripes of decreasing width (exponentially small), and relies on exact prior knowledge of both pixel size and orientation, which is unpractical in a real imaging system. In a realistic imaging setup, there is usually no reliable prior knowledge of physical pixel size or orientation relative to the fiducial pattern. The fiducial pattern should be *robust*, in the sense that it functions independently of this information, even though its exact performance may be affected by these parameters. To achieve this, practitioners have proposed the use of simple geometric patterns invariant under rotation, e.g. solid disks, as fiducials. The registration “anchor” point is the disk center. As different disks may have identical digital images, there is some uncertainty in retrieving this point from the image. The amount of this uncertainty determines the registration accuracy obtainable. Simple algorithms for estimating the disk center from its digital image have been proposed, ^{3,4} achieving subpixel accuracy, of which the most popular is the *centroid* algorithm. With this method, the disk center is taken to be the centroid of the dark “blob” representing the disk in the image. Hill ⁵ analyzes the precision obtainable from the centroid algorithm, based on mathematically incorrect assumptions. His main conclusion is that the registration error should decrease as the disk radius increases (or, equivalently, the imaging resolution increases). Computer simulations of Bose and Amir ⁴ show that his analysis is probably a good *approximation* of the truth, in the sense that increasing the disk radius does indeed increase the accuracy obtained by the centroid algorithm. In this paper we provide the first rigorous analysis of the centroid algorithm, showing that on the average (over all possible fiducial positions), subpixel accuracy of $O(r^{-1/2})$ pixel units is obtained, where r is the fiducial radius in pixel units.

Havelock ⁶ pointed out that instead of computing just a simple point estimate of the fiducial anchor point, it is more appropriate to compute the entire region of the plane where this point may lie. Calling these regions *locales*, Havelock lists a series of ideal geometric properties of locales (such as convexity). Under these (unbased) assumptions, he derives bounds on geometric measures of the locales. In this paper we investigate the locales of circular fiducials, providing precise bounds on their geometric properties, e.g. area and diameter. For example, we prove that the




Fig. 1. The image registration problem. (a) A printed circuit board to be inspected. There are two preferred orthogonal directions (vertical and horizontal). Three fiducial patterns (black disks), marked A,B and C are imprinted on the board. The disk centers induce an affine-invariant coordinate system on the image plane. The point P may be represented uniquely as the triple (α, β, γ) , where $\alpha + \beta + \gamma = 1$, meaning that $P = \alpha A + \beta B + \gamma C$. (b) A high-resolution binary image of the board when the sensor pixel orientation is not aligned with the preferred directions of the board. The positions of the fiducial centers are to be determined with subpixel accuracy. An error of ϵ in any of their coordinates induces an error of $O(\epsilon)$ in the calculation of the position of P . (c) A low-resolution binary image of the board at the same orientation as (b). The accuracy in locating P decreases relative to (b).

diameter of the locale of a solid disk of radius r is $\Omega((\log r)^{1/2} r^{-1})$ pixel units. We also provide an $O(r \log r)$ time algorithm for computing the region. This involves the investigation of the geometric structure of high order Voronoi diagrams of the unit lattice.

All the information in the digital image of a solid disk is contained in the pixels on its circumference. Because of this, O’Gorman et al. ⁷ suggest that more precision may be obtained by adding

more “circumferences” to the fiducial. This is achieved by using not a solid disk, but a “bulleye” pattern consisting of concentric rings of alternating colors. Each boundary between successive rings adds information which a registration procedure may utilize. Unfortunately, in order to achieve an order of magnitude improvement over the simple solid disk, we have to relax our requirement that the fiducial be scalable. Under this constraint, we construct a bulleye fiducial enabling image registration with $O(r^{-2} \log r)$ accuracy and provide an algorithm for computing its “locale”.

2. The Imaging Model

To facilitate our analysis, we make a few simplifying assumptions about the imaging process. First, we consider only binary images, containing black and white pixels. Second, the sensor *point-samples* the scene, essentially modelling the pixels as infinitely small points. An image pixel is black if a dark point of the scene covers the corresponding sample point, otherwise it is white, as illustrated by Fig. 2. Third, the imaging procedure is noise-free, so the pixels are exactly the points of the planar lattice Z^2 . The first assumption is a *worst case* assumption, as real images

usually contain greyscales, providing additional information which may improve the registration. The second and third assumptions are *best case*, as real sensors are noisy. Each sensor pixel integrates over an area in the scene, and noise is inevitably present.




Fig. 2. Digital image of shape. (a) Original shape. (b) Point-sampling noise-free imaging model. Pixels are modelled as lattice points. (c) Digital binary pixel image. (d) Effect of doubling imaging resolution.

3. Fiducial Design Considerations

Binary fiducials are black and white patterns. These patterns, when imprinted on objects to be inspected, enable precise location of select points of the pattern in their digital images. A major consideration in fiducial design is the area/accuracy tradeoff. A successful fiducial should be compact (occupy a small number of image pixels), and achieve a good registration accuracy. Bruckstein et al. ² designed a “one-dimensional” fiducial, shown in Fig. 3, occupying a line of n pixels, and achieving a registration accuracy of 2^{-n} pixel units. In practical terms, they are able to locate an interval of this size in which the edge of the fiducial must lie. In the same paper, they generalize this pattern to two dimensions with the same accuracy and prove that this is in a sense optimal, so no better accuracy may be achieved by a n -pixel fiducial. Their design, however, is tailored to the parameters of the image, requiring *exact* knowledge of the pixel size and orientation. Without this, the fiducial is useless. A more “robust” fiducial would not rely on this information. It would be essentially independent of pixel orientation (isotropic), and be “scalable” - not rely on prior knowledge of the pixel size, but gracefully trade off registration accuracy for area (e.g. resulting from an increase or decrease in sensor resolution). Here we concentrate on the design of these type of fiducials. We use solid disks as rotation and scale invariant fiducials, and bulleeyes as rotation (but not scale) invariant fiducials. The bulleye fiducial significantly improves the accuracy obtainable from the simple disk.

4. The Disk Registration Problem

Based on the criteria of Section 3, circular patterns, being isotropic, are good candidates to serve as fiducials. The simplest such fiducial pattern is a black disk on a white background. The center of the disk serves as the registration point. Let $D(\alpha, \beta, r)$ be a disk in the plane of radius r , centered at (α, β) , and Z^2 the unit planar lattice. Call the set of lattice points $\sigma(D) = D \cap Z^2$ the *signature* of D .




Fig. 3. The one dimensional n -pixel fiducial of Bruckstein et al. achieving a registration error of 2^{-n} pixel units. The pattern of black/white pixels obtained by the sensor (bottom line) is the binary expansion of $1 - d$ up to n digits (d is the distance from the left edge of the fiducial to the preceding lattice point). The left edge of the fiducial must therefore lie in the interval $[d, d + 2^{-n}]$.

This is the noise-free point-sampled binary image of D . D is said to be *consistent* with $\sigma(D)$. Denoting by $\text{Conv}(A)$ the convex hull of the set A , observe that all the information of $\sigma(D)$ is contained in the vertices of $\text{Conv}(\sigma(D))$ (a subset of $\sigma(D)$), in the sense that if these vertices are known, the entire signature σ may be reconstructed (by a “scan-conversion” process). The mapping of D to $\sigma(D)$ is not injective. Fig. 4(a) demonstrates that distinct disks may be consistent with one signature. This means that given a disk signature, it is usually not possible to determine exactly the original disk center. However, the larger the disk radius, the more constrained the disk center is, because of the decreasing continuous curvature of its circumference. This is not true for a square-shaped fiducial, as Fig. 4(b,c) demonstrates. At some orientations, increasing the size of the square does not reduce the uncertainty in the position of its center, given its signature. As the orientation of the fiducial relative to the pixel grid is arbitrary, the worst case must be assumed.

Assume the unit of distance is the image pixel size. The *disk registration problem* is the determination of the disk parameters with best possible accuracy from its digital signature. Denote the parameters of D by $\alpha(D)$, $\beta(D)$ (center) and $r(D)$ (radius) respectively. The disk registration problem can now be formulated as follows:

Disk Registration Problem: Given $\sigma(D)$ - determine $(\alpha(D), \beta(D), r(D))$.

Because of the uncertainty in the signature, there is no unique answer to the disk registration problem.




Fig. 4. Circular and square fiducials. (a) The solid disk fiducial. Many distinct disks may have identical signatures (the two small disks and the two large disks are *consistent* with the same signatures), so the disk registration problem may not be answered precisely. Constraints on disk center increase with disk radius r , so registration accuracy should improve with an increase in r . (b) Square fiducial: When the square is aligned with the pixel grid, the uncertainty in the position of its center is independent of size (0.5 pixel units), so squares are not good fiducials. (c) The square at a 45° angle orientation relative to the pixel grid. The uncertainty in this case may also reach 0.5 pixel units, independent of the size of the square.

5. Lower Bounds on Disk Registration

We now address the question of how accurate a registration may be achieved using a solid disk as a fiducial pattern. The accuracy is measured in pixel units. Let σ be a disk signature, and $F(\sigma)$ the locus of the centers of all disks consistent with this signature, called the *feasible region* of σ :

$$F(\sigma) = \{(x, y) \in \mathbb{R}^2 : \exists r [\sigma(D(x, y, r)) = \sigma]\} \quad . \quad (1)$$

Obviously, since any point in $F(\sigma)$ is a possible center of a disk with signature σ , examination of this region will provide lower bounds on the registration accuracy inherent in the signature. These regions are the *locales* of Havelock (Ref. 6).^a The definition (1) of $F(\sigma)$ is equivalent to:

$$F(\sigma) = \{(x, y) \in \mathbb{R}^2 : \exists r [\forall p \in \sigma (\|p - (x, y)\| \leq r) \wedge \forall p \in \sigma^c (\|p - (x, y)\| \geq r)]\}$$

where $\sigma^c = Z^2 - \sigma$ and $\|\cdot\|$ is the L_2 (Euclidean) norm. This, in turn, means that $F(\sigma)$ is the locus of all points in the plane closer to all lattice points in σ than to all other lattice points. These locii have been investigated in the field of computational geometry. Denote $k = |\sigma|$ (the cardinality of σ). $F(\sigma)$ is the *cell* corresponding to the point set σ in the order- k Voronoi diagram of the unit lattice Z^2 (see Lee⁸ for definitions of higher order Voronoi diagrams). This is known to be a convex region of the plane (Preparata and Shamos,⁹ Sect 6.3). In fact, this Voronoi diagram is a partition of the plane into convex cells, each cell a feasible region consistent with

^aActually, Havelock defines a locale $F_r(\sigma)$ for disks of a given radius r . Contrasting with $F(\sigma)$, $F_r(\sigma)$ is usually not convex.




Fig. 5. Order-23 Voronoi cells of the lattice points Z^2 in a unit square. This structure is periodic over the entire plane. (a) L_2 norm Voronoi cells. (b) L_1 norm Voronoi cells. (c) L_∞ norm Voronoi cells. The rich structure of the L_2 norm Voronoi diagram, compared to that of the L_1 and L_∞ norm diagrams, is consistent with the fact that disks are better fiducials than squares.

a different disk signature of size k . In Section 6.1 we show how to compute $F(\sigma)$ efficiently, given σ . Now we bound from below various geometric properties of $F(\sigma)$.

Note: The Voronoi cell formulation is consistent with the fact that squares are very bad fiducials. The order- k Voronoi diagram of the lattice defined using the L_1 norm (corresponding to the case of an axis-aligned square) or the L_∞ norm (corresponding to the case of a square at a 45° orientation relative to the pixel grid) has a very simple structure (see Fig. 5). There are very few distinct signatures of size k generated by squares, which are the unit “circles” induced by these norms, hence very few cells. Disks, which induce a L_2 norm Voronoi diagram with a rich structure, are good fiducials.

5.1. Bounding the Feasible Region Diameter

We now show that if a disk $D(x, y, r)$ induces a signature σ , it is possible to translate D in the plane by a measurable amount without changing σ . Define

$$d(D) = \min_{p \in Z^2, q \in \partial D} \|p - q\|_2$$

$$d(r) = \max_{\{D : r(D)=r\}} d(D)$$

The quantity $d(D)$ is the distance of the closest lattice point to the circumference of D , and $d(r)$ is the worst (largest) such distance among all disks of radius r . Our first theorem guarantees the existence of disks whose circumference is distant from all lattice points. Our proof considers only disks centered at lattice points.^b There may exist other disks improving this lower bound.

^bAfter this paper was submitted for publication, we learned that Theorem 1 was proved independently by Abbott.¹⁰

Theorem 1

$$d(r) = \Omega(\log(r)^{1/2}r^{-1})$$

Proof. Landau¹¹ obtained the following number-theoretic result: The number of integers less than n which can be represented as the sum of two squares is $\theta(n \log(n)^{-1/2})$. Call these the *Landau numbers* of n . Imagine an origin-centered disk D , whose circumference intersects the lattice point $(r(D), 0) = (\lfloor \sqrt{n} \rfloor, 0)$. Any lattice point $(x, y) \in D$ represents a Landau number of n , namely $x^2 + y^2$, since $x^2 + y^2 \leq \lfloor \sqrt{n} \rfloor^2 \leq n$. All lattice points representing the same Landau number lie on the same origin-centered circle. Landau's result implies that the number of such distinct circles centered at the origin with radius less than $r = \lfloor \sqrt{n} \rfloor$, intersecting at least one lattice point is $\theta(\log(r)^{-1/2}r^2)$. Hence, for large r , the number of such circles with radius between r and $r + 1$ is $\theta(\log(r)^{-1/2}r)$, so the difference in radius between at least one pair of successive circles in this set is $\Omega(\log(r)^{1/2}r^{-1})$, whence

$$d(r) = \Omega(\log(r)^{1/2}r^{-1})$$

□

Theorem 2 Define $F(\sigma)$ as in (1). Then

$$\begin{aligned} \text{diam}(F) &= O(r^{-1/2}) \\ \text{diam}(F) &= \Omega(\log(r)^{1/2}r^{-1}) \end{aligned}$$

Proof. The upper bound follows from an extension of the proof, due to Newman,¹² that $d(r) = O(r^{-1/2})$. The lower bound follows from Theorem 1. Indeed, for any r , there exists a D_r with radius of size $O(r)$, which may be increased (or decreased) by order of magnitude $d(D_r) = \Omega(\log(r)^{1/2}r^{-1})$ without changing its signature. The concentric disk between the circumference of D_r and the nearest lattice point of radius $r \pm d(D_r)/2$ may be translated (i.e. its center moved) in any direction the distance $d(D_r)/2$. The theorem follows. □

5.2. Bounding the Feasible Region Area

In this section we bound the area of the feasible region $F(\sigma)$.

Theorem 3

$$\text{Area}(F) = \Omega(r^{-2}) \quad .$$

Proof. It is known⁸ that the number of cells in an order- k Voronoi diagram of n points is $O(nk)$. For any integer m , and $n = m^2$, consider the portion of the unit planar lattice $Z^2 \cap [0, m]^2$. The order- k Voronoi diagram of this $\theta(n)$ point set behaves periodically over this truncated lattice (except at the boundaries), so its global structure is determined uniquely by its local structure in the unit square $[0, 1]^2$. There are $O(nk)$ cells in the entire structure, consisting of n unit squares, so there are $O(k)$ cells in one unit square. As the unit square has unit area, there must

be at least one cell with area $\Omega(1/k)$. The signature of a disk of radius r contains $k = \theta(r^2)$ lattice points, ^c implying the theorem. \square

Note that Theorem 3 implies the lower bound $\text{diam}(F) = \Omega(r^{-1})$, which is weaker than the bound obtained in Theorem 2.

6. Disk Registration Algorithms

This section is dedicated to the presentation and analysis of efficient disk registration algorithms. We distinguish between *exact* algorithms - those which take the signature σ as input, and supply the vertices of $F(\sigma)$ as output, and *approximate* algorithms - those which take the same signature σ as input and supply a single point, an approximation of the disk center, as output. In the former case, some point of $F(\sigma)$ has to be chosen as the registration point (usually some sort of center), but knowledge of the entire region may be useful.

6.1. Exact Algorithms

Notation: Let L be a subset of the lattice points Z^2 , and B the set of lattice points which are the vertices of lattice cells having vertices both in L and in L^c . Denote the *inner boundary* of L by $\partial_I L = B \cap L$ and the *outer boundary* of L by $\partial_O L = B \cap L^c$. The *boundary* of L is $\partial L = B$.

We now show how to compute the feasible area $F(\sigma)$ in $O(r \log r)$ time, where $r = r(D)$ is the radius of a disk D consistent with σ . ^d An algorithm for computing $F(\sigma)$ with this complexity has been proposed by Bhattacharya ¹³ in a different context, but is more complicated.

When analyzing the time complexity of an algorithm, we must first define the computation model. In our case, it is critical to specify how we receive the binary image data: After all, there is not much point developing a fast algorithm when just *loading* all the data into memory may dominate the run time. We adopt the (reasonable) assumption that the binary image is already present in memory, and we can, in $O(1)$ time, access any pixel value.

The following lemma formalizes what we have already mentioned, namely, that all information in the signature is conveyed by its boundary pixels.

Lemma 1 *When computing $F(\sigma)$, it suffices to consider only $\partial\sigma$.*

Proof. For any fixed r , the locus of centers of disks of radii r consistent with σ is given by

$$F_r(\sigma) = \left(\bigcap_{p \in \sigma} D(p, r) \right) \cap \left(\bigcup_{p \in \sigma^c} D(p, r) \right)^c$$

It is easy to see that $\bigcap_{p \in \sigma} D(p, r) = \bigcap_{p \in \partial_I \sigma} D(p, r)$, since the disk of radius r centered at p is just the locus of centers of all disks of radius r containing p , and a disk of radius r contains σ if and only if it contains $\partial_I \sigma$. If $p \in \sigma^c - \partial_O \sigma$, then for

^cThe exact number of points is $\pi r^2 + o(r^2)$, as Theorem 7 states.

^dNotice that it is not important, in terms of asymptotic analysis, which consistent disk is taken, as the radii are all of the same order of magnitude.

any $r \geq \frac{1}{2}$, $D(p, r)$ contributes nothing to $[\cup_{p \in \sigma^c} D(p, r)]^c$. Since $F(\sigma) = \cup_{r>0} F_r(\sigma)$, when σ is fixed only $\partial\sigma$ is important in the calculation of $F(\sigma)$ too. \square

Given any point of $\partial\sigma$, we can trace along the boundary of the disk signature to find all other points of $\partial\sigma$. In our computation model, we assume that at least one point of σ is known in advance. It is then possible to find a point of $\partial\sigma$ from this by advancing in any fixed direction. Since there are obviously $O(r)$ points in $\partial_I\sigma$ and $\partial_O\sigma$, there are also $O(r)$ points in $\partial\sigma$.

Following the methods in Preparata and Shamos, ⁹ Chap. 7, the collection of all disks consistent with σ may be represented by the set $Q'(\sigma) \subseteq \mathbb{R}^3$ (parameter space), where $(x, y, r^2) \in Q'(\sigma)$ if and only if

$$\forall (x_i, y_i) \in \sigma \quad (x - x_i)^2 + (y - y_i)^2 \leq r^2$$

and

$$\forall (x_i, y_i) \in \sigma^c \quad (x - x_i)^2 + (y - y_i)^2 > r^2 \quad .$$

It is obvious that $F(\sigma)$ is the projection of $Q'(\sigma)$ on the XY plane. Consider the bijective mapping $\vec{U} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ given by

$$\vec{U}((x, y, z)) = (x, y, z - x^2 - y^2)$$

By Lemma 1, it suffices to consider only $\partial\sigma$, so the point (x, y) is the center of a disk of radius r consistent with σ if and only if

$$\forall (x_i, y_i) \in \partial_I\sigma \quad \vec{a}_i \cdot \vec{U}((x, y, r^2)) \geq k_i$$

$$\forall (x_i, y_i) \in \partial_O\sigma \quad \vec{a}_i \cdot \vec{U}((x, y, r^2)) \leq k_i$$

where

$$\vec{a}_i = (2x_i, 2y_i, 1) \qquad k_i = x_i^2 + y_i^2$$

Let $Q := \vec{U}(Q')$. It follows that Q is the convex polyhedron obtained as the intersection of the appropriate half-spaces defined by the set of planes $\{\vec{x} \in \mathbb{R}^3 \mid \vec{x} \cdot \vec{a}_i = k_i\}$ corresponding to the lattice points $\{(x_i, y_i) \in \partial\sigma\}$.

Obviously, Q and Q' have the same projection on the XY plane, which is just $F(\sigma)$. The polyhedron Q can be computed, for instance, using Muller and Preparata's $O(n \log n)$ time algorithm for computing the convex polyhedron defined by n hyperplanes (described in Preparata and Shamos, ⁹ Chap. 7). As here $n = O(r)$, $Q(\sigma)$ is computable in $O(r \log r)$ time. Computing F from Q is trivial, so

Theorem 4 *Given σ , $F(\sigma)$ may be computed in $O(r \log r)$ time.*

If we would like just *one* consistent disk, it is possible to find the parameters of the disk of minimal radius consistent with σ in $O(r)$ time by simple linear programming, using e.g. the algorithms of Megiddo ¹⁴ or Matoušek et al. ¹⁵ The above analysis enables us to prove:

Theorem 5 *Given a signature σ , the locus of radii of the disks consistent with σ is an interval (in \mathbb{R}).*

Proof. The polyhedron $Q(\sigma)$ is convex, therefore its projection on the Z axis is an interval. The mapping \vec{U} is continuous, so the projection of $Q'(\sigma)$ on the Z axis is also an interval $[r_1^2, r_2^2]$. Hence, the locus of the radii of the disks consistent with σ coincide with the interval $[r_1, r_2]$. \square

6.2. Combinatorial Geometric Bounds

Denote by $\text{Comp}(P)$ (the *complexity* of P) the number of edges of a planar polygon P . Our previous results enable us to obtain an upper bound on the complexity of the cells of the order- k Voronoi diagram of the planar lattice Z^2 . This contrasts with the case of arbitrary planar point sets, where no non-trivial bound exists. As stated above, the polyhedron Q has $O(r)$ faces, and F , which is Q 's projection, is a cell in an order- $\theta(\pi r^2)$ Voronoi diagram. Conversely, any cell in any t -order Voronoi diagram is a feasible area for a disk signature of size t . Putting this together with Theorems 2 and 3 yields:

Theorem 6 *Let C be a cell of an order- k Voronoi diagram of the lattice Z^2 . Then*

$$\begin{aligned} \text{Comp}(C) &= O(k^{1/2}) \\ \text{Diam}(C) &= O(k^{-1/4}) \\ \text{Diam}(C) &= \Omega(k^{-1/2}(\log k)^{1/2}) \\ \text{Area}(C) &= \Omega(k^{-1}) \end{aligned}$$

6.3. Recognizing Digital Disks

Nakamura and Aizawa¹⁶ present an algorithm to determine whether a given set of lattice points are the digital signature of a disk. Section 6.1 implies a simpler algorithm achieving the same goal. A given set of lattice points L is the digital signature of a disk iff $F(L) \neq \emptyset$. As we have shown, given any point of L , $F(L)$ may be computed in $O(n \log n)$ time, where n is the size of ∂L .

6.4. Approximate Algorithms

In Section 6.1, we showed how to compute $F(\sigma)$ efficiently. The algorithm is, however, fairly complex, and for many applications, the entire region is not of interest. Usually just one point is required as an “estimate” for the fiducial center, along with an estimate of the radius. Algorithm Register is a simple (and essentially trivial) algorithm for estimating the radius and center of a disk given its signature:

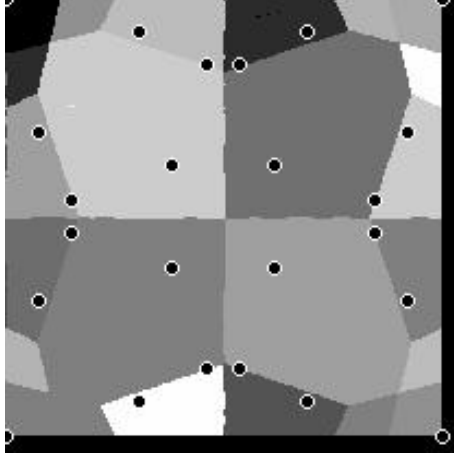


Fig. 6. Voronoi cells (feasible regions) $F(\sigma)$ and centroids $C(\sigma)$ associated with digital signatures of disks containing 13 lattice points. Note the variance between the sizes of the cells and the fact that not always $C(\sigma) \in F(\sigma)$ (C is not always consistent with F).

Register(Signature σ)

begin

$$r_{est} := \sqrt{|\sigma|/\pi};$$

$$(\alpha, \beta)_{est} := \text{centroid}(\sigma);$$

end ;

The cardinality of $\sigma(D)$ may be calculated efficiently from the lattice points on $\partial_I \sigma$ using Pick's theorem (Coxeter, ¹⁷ p. 208). Algorithm Register is straightforward enough. The natural question is whether the estimates it produces are *consistent*, i.e. whether $(\alpha, \beta)_{est} \in F(\sigma)$, and r_{est} is a possible radius of a disk consistent with σ . The answer to this is negative. Fig. 6 shows a centroid of a disk signature which lies outside the feasible region of the same signature. This already indicates that the algorithm is approximate. However, the next few theorems show that, although the Register algorithm does not always produce consistent estimates, it does perform well on the average. First some notation:

Notation:

$$\Delta(D) = |\sigma(D)| - \pi r(D)^2$$

$$\Delta(r) = \max_{\{D:r(D)=r\}} \Delta(D)$$

The quantity $\Delta(D)$ is the *discrepancy* of the disk D - the difference between its area and the number of lattice points in its interior. The quantity $\Delta(r)$ is the worst (largest) discrepancy among all disks of radius r . The following theorem summarizes some known bounds on $\Delta(r)$.

Theorem 7 1. (Iwaniec and Mozzochi ¹⁸) $\Delta(r) = O(r^{7/11})$.

2. (Hardy ¹⁹) $\Delta(r) = \Omega(r^{1/2}(\log r)^{1/4})$.

Kulpa ²⁰ reports the results of simulations aimed at estimating $\Delta(r)$. He was apparently unaware of the wealth of analytical results bounding this quantity (see Krätzel, ²¹ Chap. 3 for a historical survey), of which those quoted in Theorem 7 are, to the best of our knowledge, the best to date.

Kendall proved an average case result:

Theorem 8 (Kendall ²²)

$$\mathbf{E}_{\{D:r(D)=r\}}[\Delta(D)] = O(r^{1/2}) \quad (2)$$

where $(\alpha(D), \beta(D))$ is uniformly distributed in $[0, 1]^2$.

Kendall's result is a consequence of an estimate of the zero order moment of the distance of the points of $\sigma(D)$ from the center of D . The following lemma generalizes this to higher-order moments:

Lemma 2 Let D be a disk with center (α, β) and radius r . Define, for any non-negative integer k , the complex random variable

$$R_k(\alpha, \beta) = \sum_{p \in \sigma(D)} [p - (\alpha + i\beta)]^k$$

where $i = \sqrt{-1}$ and p is a lattice point interpreted as a complex number. Then

$$\begin{aligned} \mathbf{E}[R_k] &= \begin{cases} \pi r^2 & k = 0 \\ 0 & k > 0 \end{cases} \\ \mathbf{E}|R_k - \mathbf{E}[R_k]| &= O(r^{k+1/2}) . \end{aligned}$$

where (α, β) is uniformly distributed in $[0, 1]^2$.

Proof. See Appendix A. □

The special case $k = 0$ of Lemma 2 is Kendall's result. We shall need the $k = 1$ case in the sequel.

The bounds summarized in Theorem 7 may be used to quantify the accuracy of the estimate of the disk radius by Algorithm Register. The closer $|\sigma(D)|$ is to the true area of the disk D , namely, $\pi D(r)^2$, the more accurate the estimate r_{est} will be. Theorem 7 (1) implies that the estimate always satisfies (in the worst case)

$$|r - r_{est}| = O(r^{-4/11})$$

and by (2), on the average (over all disk centers consistent with σ)

$$\mathbf{E}|r - r_{est}| = O(r^{-1/2}) .$$

The accuracy of the centroid estimate of Algorithm Register for the disk center may be bounded too. Note that the centroid is determined by the signature only, so just as a signature σ defines a unique feasible region, it determines also a unique centroid. Analogously to the analysis of Theorem 3, the number of centroids

corresponding to signatures of size k is $O(k)$ per unit square. Fig. 6 shows some of these. Using the terminology of Lemma 2, the average distance between a disk center (α, β) consistent with σ and the centroid of σ is

$$\mathbf{E}||(\alpha, \beta) - (\alpha, \beta)_{est}|| = \frac{\mathbf{E}|R_1|}{\theta(\pi r^2)} .$$

Since, by Lemma 2

$$\mathbf{E}|R_1| = O(r^{3/2}) ,$$

we have

$$\mathbf{E}||(\alpha, \beta) - (\alpha, \beta)_{est}|| = O(r^{-1/2}) . \quad (3)$$

Our results mean that the improvement in average registration accuracy caused by an increase in disk radius behaves asymptotically like $r^{-1/2}$. This confirms and explains the empirical results of Bose and Amir. ⁴

7. The “Bulleye” Fiducial

O’Gorman et al. ⁷ mention that a more promising variant of the circular solid disk fiducial might be the “bulleye” - a number of concentric rings of alternating colors. They show, under the (plausible but mathematically incorrect) assumption that the signatures of concentric rings are “uncorrelated”, that the average registration accuracy of $O(r)$ concentric rings (where the width of each ring is a constant number of pixels) using a weighted centroid algorithm, is $O(r^{-1})$. Were it true, this would mean that an order of magnitude improvement in registration accuracy (relative to (3)) would be achievable without sacrificing any additional space beyond that of the simple disk. In this section, we design a bulleye fiducial of radius $O(r)$ achieving registration accuracy $O(r^{-2} \log r)$. This increase in accuracy relative to the simple solid disk is due to the fact that we assume prior knowledge of the pixel size. We do not require, however, knowledge of the pixel orientation. The bulleye is a circular version of a variant of the optimal one-dimensional fiducial of Bruckstein et al. (see Fig. 3). We prove that these fiducial patterns have a feasible area of diameter $\leq \frac{\log r}{r(r+\log r)}$, and show how to achieve this accuracy. The variant of the optimal one dimensional fiducial we use is robust to limited noise in the image.

7.1. A Robust One Dimensional Fiducial

Definition 1 A δ -noisy one dimensional lattice is an infinite set of points $\{p_i \in \mathbb{R} : i \in \mathbb{Z}\}$, such that $\forall i |p_i - i| \leq \delta$.

Let us see where the algorithm, described in Fig. 3, for the optimal one dimensional fiducial can fail as a result of image noise. Consider, for example, the first component of the fiducial, consisting of two blocks of size $\frac{1}{2}$, one white and one black. The algorithm may err only if a lattice point was close to the boundary between the blocks, and “moved” from one to the other as a result of noise. Assuming the noise is bounded by δ (a δ -noisy lattice), all we have to do to correct this is check whether the lattice point is in the interval $[\frac{1}{2} - 2\delta, \frac{1}{2} + 2\delta]$. This can be

achieved by appending a unit length window to the fiducial which is all white except for that interval, which is black. The position of this “error-correction” window in the fiducial has no importance, as long as its distance to the edge of the two-halves block is a integral number of pixels. We can correct noise in all other components of the fiducial by appending to them similar error-correction windows (see Fig. 7). Since it is impossible to achieve registration accuracy better than δ on a δ -noisy lattice, the length of the entire fiducial should be n pixels, where $2^{-n} \approx \delta$.

For $\delta = \frac{\log r}{r(r+\log r)}$, the length of the optimal one-dimensional fiducial achieving registration accuracy of δ (including error-correction windows) is $2 \times \log_2 \delta^{-1} \leq 2 \log_2 r$.




Fig. 7. Robust one-dimensional fiducial construction. (a) Four-quarters block of the optimal one-dimensional fiducial of Fig. 3. (b) The error-correction window corresponding to the block of (a). In general, it is all white, except for a strip of width 2δ at each interval boundary in (a) not covered in the windows corresponding to the previous blocks.

7.2. The Two Dimensional Bulleye

This circular fiducial is very simple. For $\delta = \frac{\log r}{r(r+\log r)}$, let $b_1, w_1, b_2, w_2, \dots, b_m, w_m$ be the lengths of the black and white strips (respectively) of the one-dimensional fiducial built for a noisy δ -noisy lattice (as described in Section 7.1). The two dimensional fiducial is constructed from a solid black disk of radius r , a white ring of width 2, and a sequence of black and white rings, where the i 'th black (white resp.) ring has width b_i (w_i resp.) (see Fig. 8). The radius of the resulting fiducial is $\leq r + 2 \log_2 r = O(r)$.




Fig. 8. To create the bulleye fiducial achieving a registration accuracy of $\delta = \frac{\log r}{r(r+\log r)}$, rotate the δ -robust one-dimensional fiducial around a solid black disk of radius r enclosed by a white ring of width 2.

7.3. Bounding the Center of the Bulleye

It is possible to compute the center of this fiducial from its digital image with accuracy δ . First locate the center to within one pixel accuracy. This is possible

because of the white ring of two pixels radius between the solid black disk and the rotated one-dimensional fiducial. Using the algorithms described earlier in the paper (e.g. the centroid algorithm) on the signature of the solid disk (which may be isolated), the center may certainly be found up to one pixel accuracy.

To find the center of the bulleye with accuracy δ , we use two lines of image pixels only - the horizontal row not above and closest to the center of the bulleye, and the vertical column not to the right and closest to the center of the bulleye. These horizontal and vertical lines are adjacent to the pixel just found. Call these lines the X -line and Y -line, respectively. Assume that the true fiducial center is δ_x pixel units above the X -line, and δ_y pixel units to the right of the Y -line. The square of size $\delta = \frac{\log r}{r(r+\log r)}$ centered at (δ_x, δ_y) contains the center of the bulleye. We now show how to compute (δ_x, δ_y) .

The pixels along the X -line form the image of a one-dimensional fiducial which has shifted to the left by s_x pixel units, in which the individual blocks have shifted relative to each other by no more than

$$\begin{aligned} \frac{\delta_x^2}{r} - \frac{\delta_x^2}{r + \log r} &\leq \frac{1}{r} - \frac{1}{r + \log r} \\ &= \frac{\log r}{r(r + \log r)} \\ &= \delta \end{aligned}$$

pixels. Since this "noise" is not more than δ , the δ -robust one-dimensional fiducial algorithm covers this case. Apply the "error-correction" algorithm on this pixel pattern to obtain the estimate d_x . Do the same for the pixels on the vertical Y -line to obtain d_y . The quantity s_y is defined analogously to s_x . By elementary geometry,

$$r^2 = (r - s_x)^2 + \delta_x^2 \qquad r^2 = (r - s_y)^2 + \delta_y^2$$

yielding

$$s_x = r - \sqrt{r^2 - \delta_x^2} \qquad s_y = r - \sqrt{r^2 - \delta_y^2} \quad .$$

As $\delta_x = s_x + d_x$ and $\delta_y = s_y + d_y$, this results in the quadratic equations for the two unknowns δ_x and δ_y :

$$\begin{aligned} \delta_x &= d_x + r - \sqrt{r^2 - \delta_x^2} \\ \delta_y &= d_y + r - \sqrt{r^2 - \delta_y^2} \end{aligned}$$

which may be solved easily.

8. Concluding Remarks

We have presented analytic bounds on the registration accuracy achievable in binary images using circular fiducials, along with efficient methods for performing the registration. For solid disks, $\Omega((\log r)^{1/2}r^{-1})$ is the best that can be hoped for,

the simple centroid algorithm achieving $O(r^{-1/2})$ accuracy on the average. Whether the lower bound may be achieved depends on improving the *upper* bound on the diameter of the feasible regions.

The non-scalable bulleye fiducial improves the registration accuracy of a solid disk by an order of magnitude to $O(r^{-2} \log r)$. It is not clear whether a similar order of magnitude improvement may be achieved for a scalable version (e.g. containing a constant number of rings). This will be resolved when the relation between the lattice points on the convex hulls of concentric ring signatures is determined. They seem to be “uncorrelated” in a sense, but this must be quantified more precisely.

A major open problem is whether an exponentially small registration accuracy may be achieved using some isotropic scalable fiducial pattern. The circular fiducials investigated here seem to be able to yield only a polynomially small registration accuracy.

Our analysis deals with the perfect model where the signature is not noisy. In real-world scenarios, it is usually impossible to eliminate noise from the imaging process. It would be interesting to determine how robust our methods are with respect to noise. Empirical results of O’Gorman et al. ⁷ indicate that circular fiducials are indeed robust.

Our methods could probably be improved if our assumption of a binary image were relaxed. Greyscale images contain more information than binary images, and are readily available. This case deserves separate investigation.

Acknowledgements

Thanks to Alfred Bruckstein for presenting these problems to us and to Ronny Roth for interesting discussions on number theory. Thanks also to Avigail Orni for many helpful remarks.

References

1. C. Gotsman. On the most robust affine basis. *Pattern Recognition Letters*, 14:647–650, 1993.
2. A.M. Bruckstein, L. O’Gorman, and A. Orlicsky. Design of shapes for precise image registration. Technical report, AT&T Bell Labs., 1989.
3. I. Amir. Algorithm for finding the center of circular fiducials. *Computer Vision, Graphics and Image Processing*, 49:398–406, 1990.
4. C.B. Bose and I. Amir. Design of fiducials for accurate registration using machine vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12:1196–1200, 1990.
5. J.W. Hill. Dimensional measurements from quantized images. In D. Nitzan et al., editor, *Machine Intelligence Research Applied to Industrial Automation*, pages 75–105. SRI, 1980.
6. D.I. Havelock. The topology of locales and its effects on position uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(4):380–386, 1991.
7. L O’Gorman, A.M. Bruckstein, C.B. Bose, and I. Amir. Subpixel registration using a concentric ring fiducial. In *Proceedings of Int. Conf. on Pattern Recognition*. IEEE, 1990.

8. D.T. Lee. On k-nearest neighbor Voronoi diagrams in the plane. *IEEE Transactions on Computers*, 31(6):478–487, 1982.
9. F.P. Preparata and I.M. Shamos. *Computational Geometry: An Introduction*. Springer-Verlag, New York, 1985.
10. H.L. Abbott. Approximation for the location of the center of a disk. *SIAM Review*, 34:658–661, 1992.
11. E. Landau. Über die einteilung der ... zahlen in 4 klassen .. *Arch. Math. Phys.*, 3(13):305–312, 1908.
12. D.J. Newman. A gradually turning curve must come near a lattice point. *Journal of Number Theory*, 6:7–10, 1974.
13. B.K. Bhattacharya. Circular separability of planar point sets. In G.T. Toussaint, editor, *Computational Morphology*, pages 25–39. Elsevier-Science, 1988.
14. N. Megiddo. Linear programming in linear time when the dimension is fixed. *Journal of the ACM*, 31:114–127, 1984.
15. J. Matoušek, M. Sharir, and E. Welzl. A subexponential bound for linear programming. In *Proceedings of the 8th Annual ACM Symposium on Computational Geometry*, pages 1–8, 1992.
16. A. Nakamura and K. Aizawa. Digital circles. *Computer Vision, Graphics and Image Processing*, 26:242–255, 1984.
17. H.S.M. Coxeter. *Introduction to Geometry (2nd Edition)*. Wiley and Sons Inc., New-York, 1980.
18. H. Iwaniec and J. Mozzochi. On the divisor and circle problems. *Journal of Number Theory*, 29:60–93, 1988.
19. G.H. Hardy. On Dirichlet’s divisor problem. *Proceedings of the London Mathematical Society*, 2(15):1–25, 1916.
20. Z. Kulpa. On the properties of discrete circles, rings and disks. *Computer Vision, Graphics and Image Processing*, 10:348–365, 1979.
21. E. Krätzel. *Lattice Points*. Kluwer Publishers, 1988.
22. D.G. Kendall. On the number of lattice points inside a random oval. *The Quarterly Journal of Math. Ser. B (Oxford)*, 19:1–25, 1948.
23. M. Abramowitz and I. Stegun. *Handbook of Mathematical Functions*. Dover Publications Inc., New-York, 1970.

Appendix A

Lemma 2 *Let D be a planar disk with center (α, β) and radius r . Define in the complex plane, for any non-negative integer k , the complex random variable*

$$R_k(\alpha, \beta) = \sum_{p \in \sigma(D)} [p - (\alpha + i\beta)]^k$$

where $i = \sqrt{-1}$ and p is a lattice point interpreted as a complex number. Then

$$\begin{aligned} \mathbf{E}[R_k] &= \begin{cases} \pi r^2 & k = 0 \\ 0 & k > 0 \end{cases} \\ \mathbf{E}|R_k - \mathbf{E}[R_k]| &= O(r^{k+1/2}) . \end{aligned}$$

where (α, β) is uniformly distributed in $[0, 1]^2$.

Proof. The complex random variable R_k is doubly periodic in α and β and hence can be expanded, at least formally, as a Fourier series.

$$R_k(\alpha, \beta) = \sum_{n,m} a_{n,m} e^{2\pi i(n\alpha + m\beta)}$$

Define the function $U_r : C \rightarrow C$ as

$$U_r(z) = \begin{cases} z^k & \text{if } |z| \leq r \\ 0 & \text{otherwise} \end{cases}$$

Then

$$R_k(\alpha, \beta) = \sum_{j,k} U_r((j - \alpha) + i(k - \beta)),$$

where (j, k) range over all lattice points. Now, by an inverse Fourier transform,

$$\begin{aligned} a_{n,m} &= \int_0^1 \int_0^1 R_k(\alpha, \beta) e^{-2\pi i(n\alpha + m\beta)} d\alpha d\beta \\ &= \sum_{j,k} \int_{j-1}^j \int_{k-1}^k U_r(u + iv) e^{-2\pi i(nu + mv)} dudv \end{aligned}$$

By transforming to polar coordinates, denoting $\gamma = (m^2 + n^2)^{1/2}$ and $\alpha = \arg(n + im)$:

$$\begin{aligned} a_{n,m} &= \int_0^{2\pi} \int_0^r x^k e^{ik(\theta - \alpha)} e^{-2\pi i\gamma x \cos \theta} x dx d\theta \\ &= e^{-ik\alpha} \int_0^r x^{k+1} \left[\int_0^{2\pi} e^{-2\pi i\gamma x \cos \theta} e^{ik\theta} d\theta \right] dx \end{aligned}$$

For $\gamma = 0$ ($n = m = 0$), direct computation yields

$$a_{0,0} = \mathbf{E}[R_k] = \begin{cases} \pi r^2 & k = 0 \\ 0 & k > 0 \end{cases}$$

For $\gamma > 0$

$$a_{n,m} = 2\pi i^k e^{-ik\alpha} \int_0^r x^{k+1} J_k(2\pi\gamma x) dx$$

where $J_k(z)$ is the k 'th order Bessel function. Using the integral identity for any non-negative integer k (Abramowitz and Stegun, ²³ Eq. 11.3.20)

$$\int_0^s x^{k+1} J_k(x) dx = s^{k+1} J_{k+1}(s)$$

we obtain

$$a_{n,m} = \frac{i^k e^{-i\alpha} r^{k+1}}{\gamma} J_{k+1}(2\pi\gamma r) \quad (n, m) \neq (0, 0)$$

It is well known (Abramowitz and Stegun, ²³ Eq. 9.2.1) that for all $k \geq 0$,

$$J_k(z) = O(z^{-1/2})$$

whence

$$|a_{n,m}| = (n^2 + m^2)^{-3/4} O(r^{k+1/2}) \quad (n, m) \neq (0, 0) .$$

Parseval's identity from Fourier analysis yields:

$$\begin{aligned} \mathbf{E}|R_k|^2 &= \int_0^1 \int_0^1 |R_k|^2 d\alpha d\beta \\ &= \sum_{n,m} |a_{n,m}|^2 \end{aligned}$$

Consequently

$$\begin{aligned} \mathbf{E}|R_k - \mathbf{E}[R_k]|^2 &= \mathbf{E}|R_k|^2 - |\mathbf{E}[R_k]|^2 \\ &= \sum_{(n,m) \neq (0,0)} |a_{n,m}|^2 \\ &= O(r^{2k+1}) \sum_{n,m} (n^2 + m^2)^{-3/2} \end{aligned}$$

Let $s(l)$ be the number of representations of the integer l as the sum of two squares.

It is well known that

$$s(l) = O(l^\delta)$$

for every $\delta > 0$, so $\sum_{l=1}^{\infty} l^{-3/2} s(l)$ is convergent, and

$$\begin{aligned} \mathbf{E}|R_k - \mathbf{E}[R_k]|^2 &= O(r^{2k+1}) \sum_{l=1}^{\infty} l^{-3/2} s(l) \\ &= O(r^{2k+1}) \end{aligned}$$

It follows that

$$\begin{aligned} \mathbf{E}|R_k - \mathbf{E}[R_k]| &\leq \sqrt{\mathbf{E}|R_k - \mathbf{E}[R_k]|^2} \\ &= O(r^{k+1/2}) , \end{aligned}$$

as required. □