

Chapter 3 Two-Boundary Medial Primitives (Core Atoms)

In this chapter the *core atom* is defined, and the process by which one is formed is described in detail. Two methods for statistical analysis of populations of core atoms are then developed. The first of these operates on uncluttered targets in 2D images, with analysis proceeding upon the entire population of core atoms in the image, undifferentiated by location. The core atoms in this method are divided into sub-populations by their scale, resulting in histograms or *spectra* of scale. An analysis of orientation is performed using the technique of *angle doubling*. This approach represents an early step in the research that yielded important insight into the behavior of core atoms as they describe simple geometric objects. The second method of statistical analysis operates on core atoms formed on 3D images and differentiates populations by location. This method uses eigenanalysis instead of angle doubling and proves to be of greater practical value, as will be described in subsequent chapters.

3A. Defining the Core Atom

A *core atom* is defined as two boundary points $\mathbf{b}_1$ and $\mathbf{b}_2$ that satisfy particular requirements guaranteeing that the boundaries face each other. The boundary points, in either 2D or 3D, carry information about location and orientation. The orientation of a boundary point is normal to the object's boundary at that location.

A core atom can be represented by a single vector $\vec{\mathbf{c}}_{1,2}$ from the first boundary point to the second. A core atom is said to be "located" at the *center* point mid-way between the boundary points (see Fig. 3.1A). The *diameter* or *scale* of the core atom is defined as the length $|\vec{\mathbf{c}}_{1,2}|$. The *medialness* at the center point is high because the *boundariness* at both boundary points is high and because the boundary normals face each other. Core atoms carry information about orientation, scale and position, permitting populations of core atoms to be analyzed in these terms. Unlike some medial models, the angle between the lines from the center point to each respective boundary point is fixed at 180° for a core atom, as with Brady (Brady 1983).




Fig. 3.1 **A.** A core atom consists of two boundary points that face each other across an acceptable distance, and a center point at which the core atom is said to be *located*. **B.** The search area (gray) for boundary point b_2 depends on boundary normal $\hat{n}_1$ and the expected distance between the boundaries.

A certain flexibility in relative orientation of the associated boundaries away from parallel is allowed. Boundariness is regularly sampled throughout the image to select a population of boundary points b_i at locations $\mathbf{x}_i$ with orientations $\hat{n}_i$. Any kind of boundariness, including that which is based on gradient, variance, or texture analysis, can be used to form core atoms, provided an orientation is established for each boundary point. In general the aperture of the boundariness detector is held proportional to the expected scale of the core atoms.

Core atoms are created from a population of candidate boundary points by finding pairs that satisfy the following three criteria:

- (1) The scale of the core atom vector $\vec{c}_{1,2}$ must be in the specified distance range $[c_{\min}, c_{\max}]$.

$$\vec{c}_{1,2} = \mathbf{x}_2 - \mathbf{x}_1 \quad c_{\min} \leq |\vec{c}_{1,2}| \leq c_{\max} \quad (3.1)$$

The core atom vector can be oriented either way between the boundary points, since the order of the boundary points is arbitrary.

- (2) The boundary points must have sufficient *face-to-faceness* F defined as

$$F(b_1, b_2) = f_1 \diamond f_2 \quad f_1 = \hat{c}_{1,2} \diamond \hat{n}_1 \quad f_2 = \hat{c}_{2,1} \diamond \hat{n}_2 \quad (3.2)$$

("^" denotes normalization, $\hat{\mathbf{v}} = \mathbf{v}/\|\mathbf{v}\|$.) Since f_1 and f_2 are normalized to lie between +1 and -1, their product F must also lie between +1 and -1. Values for F near +1 occur when the boundaries face towards (or away from) each other across the distance between them. A threshold for acceptable face-to-faceness is set within some error e_f such that $F(\mathbf{b}_1, \mathbf{b}_2) > 1 - e_f$.

(3) $F(\mathbf{b}_1, \mathbf{b}_2) > 0$ implies that f_1 and f_2 are both positive, or both negative. The sign of f_1 (or f_2) is called the *polarity*. The appropriate polarity is either + or - depending on whether the expected target is lighter or darker than the background.

A single boundary point can be involved in a number of core atoms, linking each to a different partner on the other side of the object. Although at first glance the search for pairs of boundary points appears to be $O(n^2)$, hashing individual boundary points beforehand by location yields a large reduction in computation time (see Fig. 3.1B). The search area for boundary point $\mathbf{b}_2$ is limited to a solid sector surrounding the orientation $\hat{\mathbf{n}}_1$ of boundary point $\mathbf{b}_1$, and to a range between $c_{\min}$ and $c_{\max}$. The width of the sector depends on e_f .

In the following sections, two methods of statistical analysis of core atom populations are described. Core atoms tend to concentrate in clouds, forming in large numbers that often surpass the number of boundary points. This provides for robust statistics, compensating for the somewhat arbitrary nature of inclusion and exclusion in a cloud caused by thresholding.

3B. Unlocalized Spectra of Scale in 2D

This section describes a method for statistical analysis of core atom populations in 2D images. It was developed in the initial stages of this doctoral research and is included because it provided early insights into the behavior of core atoms (Stetten, Landesman et al. 1997). Core atoms are formed as described in the previous section, on a 2D image of a single object. The resulting core atoms are divided into sub-populations by their scale. Each sub-population can then be analyzed to yield a set of properties as a function of scale that are invariant to translation and rotation of the underlying object. Three such properties are derived.

- (1) *magnitude* , the number of core atoms at each scale,
- (2) *directionality*, a measure of whether the core atoms at each scale tend to be oriented in a particular direction (directionality ≈ 1) or evenly distributed in orientation (directionality ≈ 0) ,
- (3) *orientation*, the predominant orientation of the core atoms at each scale, assuming directionality ≈ 1 .

Plotting such properties as functions of scale yields *spectra* analogous to those used in infrared spectroscopy to provide clues about the configurations of chemical atoms within molecules. Clues about the configuration of core atoms in a population can be derived in a similar fashion from the spectra of magnitude, directionality, and orientation. Translational invariance is inherent in these spectra (as it is in infrared spectroscopy) since location is ignored. The use of scale spectra for local image properties has been described by others (Low and Coggins, 1990).

The magnitude spectrum offers a statistical representation of the width of an object. Since only the lengths of core atoms are involved and not their orientations, the magnitude spectrum demonstrates rotational invariance in addition to translational invariance. An example of a magnitude spectrum is shown in Figure 3.4A, in which a dark circular figure has been subjected to the formation of core atoms. The resulting core atoms are superimposed on the circle with centers shown in white and boundary points shown as black "+" symbols. The corresponding magnitude spectrum exhibits a peak at the diameter of the circle. (The log of magnitude is plotted to increase the dynamic range of the display).

Below magnitude are displayed two other spectra in Fig 3.4A, those of *directionality* and *orientation*. These are calculated by first determining an *axis vector* for the population of core atoms at each scale. The axis vector represents a composite of individual core atom orientations, indicating whether there is a predominant orientation, and if so, what that orientation is. Two constraints on the calculation of the axis vector are the following: (1) If the population has no predominant orientation, then each core atom must be canceled, on average, by another core atom orthogonal to the first. (2) The order in which one selects the boundary points in a given core atom must be arbitrary, so the orientations of two core atoms 180° apart must be indistinguishable. Both of these constraints are met by the process of angle-doubling, in which the *phase angle* (angle with respect to the x axis) of each core atom is doubled. After angle-doubling, simple vector addition of the normalized core atom orientations can be applied to compute an axis vector for the population. If the difference in orientation

between two core atoms is originally 90° (or 270°), it becomes 180° causing them to cancel under addition. If the difference is originally 180° , it becomes 360° (or 0°) making boundary order irrelevant in terms of contribution to the axis vector.

The mathematics of calculating the axis vector by angle-doubling is as follows (see Fig. 3.2):




Fig. 3.2 Angle doubling of the orientation vector $\hat{c}$.

Given a core atom whose orientation vector $\hat{c}$ has a phase angle a , the unit vector $\hat{d}$ whose phase angle is $2a$ can be determined, using standard trigonometric identities:

$$\cos 2a = \cos^2 a - \sin^2 a \quad \sin 2a = 2 \sin a \cos a \quad (3.3)$$

$$\hat{d}_x = \hat{c}_x^2 - \hat{c}_y^2 \quad \hat{d}_y = 2\hat{c}_y \hat{c}_x \quad (3.4)$$

where $(\hat{d}_x, \hat{d}_y)$ are the components of the vector $\hat{d}$, etc. We sum over the entire population of n core atoms at a given scale, dividing by n to yield the vector $\mathbf{b}$,

$$\mathbf{b} = \frac{1}{n} \sum_{i=1}^n \hat{\mathbf{d}}_i, \quad (3.5)$$

where $\hat{\mathbf{d}}_i$ indicates the i^{th} vector $\hat{\mathbf{d}}$. The vector $\mathbf{b}$ may be said to have a phase angle of $2g$ which, when halved, produces the axis vector $\mathbf{a}$ whose phase angle is. Halving the phase angle is accomplished using the standard trigonometric identities:

$$\sin^2 g = \frac{1 - \cos 2g}{2} \quad \cos^2 g = \frac{1 + \cos 2g}{2} \quad (3.6)$$

$$\hat{a}_y = \pm \sqrt{\frac{1 - \hat{b}_x}{2}} \quad \begin{array}{ll} \hat{a}_y \geq 0 & \text{if } \hat{b}_y \geq 0 \\ \hat{a}_y < 0 & \text{if } \hat{b}_y < 0 \end{array} \quad \hat{a}_x = \sqrt{\frac{1 + \hat{b}_x}{2}} \quad (3.7)$$

Constraints on the sign of the square root in calculating $\hat{a}_y$ restrict $\hat{\mathbf{a}}$ arbitrarily to the two right-hand quadrants. Orientation is expected to be meaningful only over 180 degrees because the boundary points in a core atom are order-independent.

Because $\mathbf{b}$ has been normalized by n , the modulus of $\mathbf{b}$ is a scalar between 0 and 1 that measures the *directionality* of the core atom population at that scale. If all the core atoms all have the same orientation, $|\mathbf{b}| = 1$. For a uniform distribution of orientations $|\mathbf{b}| = 0$, since each core atom has, on average, another core atom orthogonal to it. This is shown in figure 3.3A, where core atoms are evenly distributed around a circular figure. The corresponding experimental data (see figure 3.4A) shows an directionality close to 0, especially at scales where the number of core atoms is greatest. The dip toward zero in directionality corresponding to the maximum number of core atoms is predictable due to random variation in orientation according to $\sqrt{n}$.




Fig. 3.3 The two extremes of directionality. **A.** Core atoms with evenly distributed orientations. **B.** Core atoms aligned along a 1D ridge.

At the other extreme, a directionality of 1 is shown in figure 3.3B. Here, the core atoms have the same orientation, resulting in an axis vector $\mathbf{a}$ with a modulus of 1.

When $\mathbf{a}$ has a non-zero length, its orientation may also be reported. The underlying direction of the ridge will be orthogonal to the orientation of the core atoms. This results in the third spectrum, that of *orientation*. Although the orientation spectrum shifts vertically when the figure is rotated as shown in Fig. 3.5, the difference between orientations as a function of scale remains fixed, constituting a useful form of rotational invariance.




Fig. 3.4 Objects (grey), core atom centers (white) and boundary points (black "+").




Fig. 3.5 Rectangular object showing effects of rotation.




Fig. 3.6 Objects showing the ability to remove noise and differentiate shapes.




Fig. 3.7 Spade showing the effects of rotation.

Figure 3.4B shows an ellipse. The magnitude spectrum is now split into two distinct modes. Within each mode, the directionality is near 1, and the orientations of the two modes are 90° apart. These modes represent the minor and major axes of the ellipse, with the minor axis appearing at the smaller scale, i.e., the left end of the spectrum. Some core atoms cross the ellipse but are not oriented along the major or the minor axis. These core atoms have an intermediate scale (shorter than core atoms of the major axis and longer than those of the minor axis) and fall into two distinct populations in terms of their orientation. Therefore they exhibit a directionality of less than 1.

The orientation spectra in Figs. 3.4-3.7 are duplicated modulo 180° and wrap around from -180° to $+180^\circ$. The orientations are printed lightly (as in Fig. 3.4A) at scales where minimal directionality or insufficient magnitude exists. The total number of core atoms is included above each image. A large number of core atoms can form even on a small object, because each boundary point may form multiple links. For example, 4016 core atoms formed for the circle, which is close to the total number of pixels in the image! Large populations of core atoms are desirable for the purposes of statistical analysis.

Figure 3.4C shows a shape known as a *paisley*. The curved core of this figure smoothly changes orientation as a function of scale from the tip to the belly of the paisley. The slope of the orientation spectrum would have the opposite sign if the mirror image of the paisley were used. Eccentricity is near 1 throughout, since at each scale the sub-population is uniformly oriented across the paisley.

A vertical rectangle is shown in 3.4D, resulting in two distinct core atom populations, each with eccentricities approximately equal to 1. The two populations correspond to separate vertical and horizontal cores within the rectangle. The first population consists of small-scale core atoms with 0° orientation forming on the vertical core. The second population consists of larger-scale core atoms with 90° orientation forming on the horizontal core. The two cores cross in the center of the figure, although they miss each other in *scale space*, where an extra dimension is added to the ambient space of the image to represent scale. The direction of each core is orthogonal to the orientation of its core atoms. Subsequent reduction in the rectangle's size in figures 3.4E and 3.4F results in shifting the spectra to the left and compressing them along the scale dimension. If the spectra were plotted with $\log(\text{scale})$ along the abscissa, only shifting would occur, demonstrating a form of scale invariance.

Figures 3.5A-F show the effect of rotation on the rectangle. The magnitude and directionality spectra are essentially unaffected. The orientation spectrum shifts upward as the figure is rotated in a counter-clockwise direction. Note that the change in orientation as a function of scale remains unaffected; the long and short axes of the rectangle remain 90° apart.

Figure 3.6A shows the effect of Gaussian noise added to the rectangle figure. Core atoms are dispersed throughout the image as a result of the added noise. Figures 3.6B and 3.6C show the results of clustering core atoms by requiring that the maximum distance between centers of any two core atoms and between the scales of those two core atoms be below certain thresholds. The two largest clusters are shown in 3.6B and 3.6C, corresponding to the two cores of the rectangle described above. The combined spectra of 3.6B and 3.6C approximately equal that of 3.6D with a reduction in noise. The idea that location can be incorporated into the analysis of core atoms is developed extensively in the remainder of this dissertation.

Figs. 3.6E and 3.6F demonstrate the ability of the spectra to differentiate between a heart and a spade, which are essentially identical except for the stem of the spade. The spade in figure 3.6F contains a population of small scale core atoms across the stem with an directionality near 1 and an orientation of near 0° . These are missing from the spectra for the heart in Fig 3.6E.

Figures 3.7A-F show the effect of rotating the spade. The magnitude and directionality spectra remain essentially unchanged, while the orientation spectra progress upward. The strongest orientation is at small scale, where the stem accumulates a cluster with high directionality.

The encouraging results achieved with this method suggest its adaptation to the analysis of the cardiac ventricle using RT3D ultrasound, the stated goal of this dissertation. It is possible to apply the method to 2D slices from the 3D data set, but some 3D shapes simply do not lend themselves well to slicing. Although the LV is, in fact, fairly easy to slice, the core atom method would be more powerful if it could handle 3D shapes in general. True 3D core atoms can easily be produced using 3D boundariness, face-to-faceness calculated in 3D, etc. Analysis of core atoms in 3D, however, cannot proceed using angle-doubling. As described in the following section, angle-doubling can be replaced with eigenanalysis, which can in fact operate in any number of dimensions.

Another limitation of the method developed above is that, except in Fig. 3.6B-C, location information has not been used in the analysis. All of the core atoms in the image are lumped together into a single population. This is effective only for isolated targets without clutter or noise, and even those targets must have relatively simple shapes. As will be seen in Chapters 4 and 5, more complicated shapes can be analyzed by sorting core atoms into local clusters, each cluster providing for the analysis of a small piece of the overall shape. There are generally far fewer core atoms in each local cluster than in the entire image, and sub-dividing each cluster further by scale leaves too few at each scale for statistical analysis. Therefore, the concept of scale spectra will be abandoned for the remainder of this dissertation, and each local cluster will

be analyzed as a whole. What is lost is the ability to differentiate two cores of different scale that cross at a given location (such as the center of the rectangle in Fig 3.4D-F). What is gained is the ability to differentiate local concentrations of medialness from one another, both within one object and between different objects in a given image.

3C. Extracting Local Medial Properties in 3D

Consider now core atoms formed in 3D, with boundariness being the likelihood of there being a surface (where in 2D it was a curve). Observe that collections of core atoms in 3D can group in three basic ways corresponding to the fundamental geometric shapes shown Fig. 3.8. The surfaces of the objects are shown in dark grey with the corresponding cores shown in light gray. Under each object is shown the population of core atoms that would be expected to form with such objects, the core atoms now being depicted as simple line segments.




Fig. 3.8 Fundamental shapes (dark gray), corresponding cores (light gray), core atom populations (line segments) and eigenvectors $\hat{\mathbf{a}}_1$, $\hat{\mathbf{a}}_2$ and $\hat{\mathbf{a}}_3$.

The sphere generates a "Koosh ball" configuration of core atoms with spherical symmetry, with the core atom centers clustered at the center of the sphere. The cylinder generates a "spokes-of-a-wheel" arrangement with radial symmetry along the axis of the cylinder and the core atom centers clustered along the axis of the cylinder. The slab results in a "bed-of-nails"

configuration across the slab, with core atom centers clustered in the mid-plane of the slab. The cores of these basic objects are the point, the line, and the plane. As shown in Fig. 3.8, a system of shape-specific coordinate axes, namely $\hat{\mathbf{a}}_1$, $\hat{\mathbf{a}}_2$, and $\hat{\mathbf{a}}_3$, can be assigned in each case, although not all the axes are unique given the symmetries involved. For example, in the slab, $\hat{\mathbf{a}}_1$ and $\hat{\mathbf{a}}_2$ can rotate freely about $\hat{\mathbf{a}}_3$. Such a set of coordinate axes can be found for any population of core atoms using eigenanalysis, as will be shown below. Furthermore, the extent to which a core atom population resembles one of the three basic configurations depends on the corresponding eigenvalues. Given a population of m core atoms $\vec{\mathbf{c}}_i$, $i = 1, 2, 3, \dots, m$, the analysis of a core atom population begins by separating each core atom vector $\vec{\mathbf{c}}_i$ into its magnitude c_i and its orientation $\hat{\mathbf{c}}_i$. We continue to ignore, for the moment, the location of the core atom. The analysis of magnitude c_i over a population of core atoms yields a mean and standard deviation for the measurement of width in the underlying figure. The orientation $\hat{\mathbf{c}}_i$ of core atoms in a population lends itself to eigenanalysis, yielding measures of dimensionality and overall orientation for the population. The eigenanalysis is developed here in m dimensions, although for the remainder of the paper m will generally be equal to 3.

Given the population of n vectors in m dimensions, it is possible to find an m -dimensional vector $\hat{\mathbf{a}}_1$ that is most orthogonal to that population as a whole by minimizing the sum of squares of the dot product between $\hat{\mathbf{a}}$ and each individual core atom orientation $\hat{\mathbf{c}}_i$.

$$\hat{\mathbf{a}}_1 = \underset{\hat{\mathbf{a}}}{\operatorname{argmin}} \frac{1}{n} \sum_{i=1}^n (\hat{\mathbf{a}} \cdot \hat{\mathbf{c}}_i)^2 = \underset{\hat{\mathbf{a}}}{\operatorname{argmin}} (\hat{\mathbf{a}}^T \mathbf{C} \hat{\mathbf{a}}) \quad \text{where} \quad \mathbf{C} = \frac{1}{n} \sum_{i=1}^n \hat{\mathbf{c}}_i \hat{\mathbf{c}}_i^T \quad (3.8)$$

The exterior product $\hat{\mathbf{c}}_i \hat{\mathbf{c}}_i^T$ has the same affect as angle doubling did, namely making the sense of $\hat{\mathbf{c}}_i$ irrelevant but encoding its direction. Then as with angle doubling, the $\mathbf{C}$ matrix averages the results encoding both the directionality and orientation of the core atom population.

The $\mathbf{C}$ matrix is positive definite, symmetric, and has a unit trace. Therefore, its eigenvalues are positive and sum to 1, and its eigenvectors are orthogonal. If the eigenvalues of $\mathbf{C}$ are sorted $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_m$, the corresponding eigenvectors $\hat{\mathbf{a}}_1 \dots \hat{\mathbf{a}}_m$ are the axes of a coordinate system in which $\hat{\mathbf{a}}_1$ is the most orthogonal to the population $\hat{\mathbf{c}}_i$ as a whole. For example, $\hat{\mathbf{a}}_1$ would be the axis of the cylinder in Fig. 3. Furthermore, the eigenanalysis guarantees that $\hat{\mathbf{a}}_2$ is the most orthogonal to the population $\hat{\mathbf{c}}_i$ among those directions that are orthogonal to $\hat{\mathbf{a}}_1$. This process can be repeated until $\hat{\mathbf{a}}_m$ remains the *least* orthogonal to the population $\hat{\mathbf{c}}_i$, representing a form of *average* orientation for $\hat{\mathbf{c}}_i$. The axes $\hat{\mathbf{a}}_1 \dots \hat{\mathbf{a}}_m$ are thus ordered from *codimensional* (orthogonal to the vector set) to *dimensional* (collinear with the

vector set). In 3D the codimensional space is that of the core itself. That is, the space most orthogonal to the core atoms is the point, line, or plane of the core as shown in Fig. 3.8.

Returning now specifically to 3D, the previous analysis yields three eigenvalues which describe the dimensionality of the core.

$$l_i \geq 0 \quad l_1 + l_2 + l_3 = 1 \quad (3.9)$$

An eigenvalue of 0.0 means that the corresponding eigenvector is perfectly orthogonal to every core atom $\hat{c}_i$. Such is the case for $\hat{a}_1$ in the cylinder, and for both $\hat{a}_1$ and $\hat{a}_2$ in the slab. In the sphere none of the eigenvectors is completely orthogonal to the core atom population. Given the symmetries of the three basic shapes, the eigenvalues shown in Fig. 3.9 result.




Fig. 3.9 The lambda triangle defines the domain of possible eigenvalues that determine medial dimensionality.

Since l_3 is dependent on the other two, the system may be viewed as having only two independent variables: l_1 and l_2 . Because of constraints already mentioned, possible values for l_1 and l_2 are limited by $l_1 \leq l_2$ and $l_2 \leq (1 - l_1)/2$ which define a triangular domain called the *lambda triangle* (Fig. 3.9).

The vertices of the lambda triangle correspond to the three basic shapes in Fig. 3.8, and all possible eigenvalues fall within the triangle. A rather crude simplification of dimensionality is possible by dividing the triangle into three compartments to provide an integer description of dimensionality. Arbitrary thresholds of $l_1 = 1/5$ and $l_2 = 1/3$ will be used to divide the triangle into such areas of integer dimensionality to clarify the experimental results. However, it should be remembered that the underlying parameterization of dimensionality is not an integer or even a single scalar, but rather two independent scalars, l_1 and l_2 , whose values are constrained to be within the lambda triangle.