

Technical Report TR97-006, University of North Carolina, Department of Computer Science
SPIE Medical Imaging Conference, Feb., 1997, vol. 3034, part 2, pp. 642-652

Core atoms and the spectra of scale

George D. Stetten, Roxanne N. Landesman, Stephen M. Pizer

Department of Biomedical Engineering, Duke University, Durham, NC

NSF/ERC for Emerging Cardiovascular Technologies, Duke University, Durham, NC

Department of Biomedical Engineering, UNC, Chapel Hill, NC

Medical Imaging Display and Analysis Research Group, UNC, Chapel Hill, NC

ABSTRACT

Our purpose is to characterize figures in medical images as a first step toward finding and measuring anatomical structures. For clinical use, we require complete automation and reasonably short computation times. We do not require that a sharp boundary be determined, only that the structure be identified and measurements taken of its size and shape.

Our method involves the detection and linking of locations within an image that possess high "medialness," i.e., locations that are equidistant from two opposing boundaries. The method produces populations of *core atoms*, each core atom consisting of a center point (where medialness is high) and the two associated boundary points. We can cluster core atoms by the proximity of their centers and by the similarity of their size (scale). We generate statistical signatures of clusters to identify the underlying figure. In particular, we compute three spectra vs. scale for a cluster, including (1) *magnitude*: the number of core atoms, (2) *eccentricity*: their aggregate directional asymmetry, and (3) *orientation*: their aggregate direction.

We illustrate the production of these spectra for various graphical test images, demonstrating translational, rotational, and scale invariance of the spectra, as well as specificity between targets. We observe the effects of image noise on the spectra and show how clustering reduces these effects. Early results suggest that the scale spectra of core atoms provide an efficient and robust method for identifying figures, suitable for practical application in medical image analysis.

Keywords: medialness, core, ridge, core atom, scale, shape, spectrum, skeleton, image processing, image analysis.

2. BACKGROUND AND MOTIVATION

We describe in this paper a new method of describing objects in an image according to the *medial axes* of their underlying figures. The medial axes here are related to the *core*. The lineage of the core may be traced to the medial axis (otherwise known as symmetric axis or skeleton) introduced by Blum and developed by Nagel, Nackman, and others.^{1, 2, 3} Pizer extended the medial axis to the core by including scale, linking the radius of the medial axis to the aperture size of the boundariness operator, and producing a graded measure of behaving like a medial axis, a measure called *medialness*.^{4, 5} Methods involving cores and medialness have proven particularly robust against noise and variations in target shape.

Relating locations with high medialness to the underlying core has previously been accomplished through analyzing the geometry of loci resulting from ridge extraction. In contrast, our method depends on analyzing the statistics of populations of locations with adequately high medialness. These locations are identified by first searching for individual boundary points throughout the image in an initial sweep, and then matching pairs of boundary points to form what we call *core atoms*. Core atoms tend to cluster along the medial axis and, being a fuzzy representation of the medial ridge, allow for statistical analysis of the underlying figure.

An added benefit of our approach is efficiency. Ridge tracking usually depends on filtering for two boundary points simultaneously and has proven computationally expensive because it must cover the multiple dimensions of position, orientation and scale. It therefore typically has required initialization by a user or a predefined model. Our method of extracting cores is computationally more efficient and does not require any initialization by a user or model.

A *core atom* is composed of two *boundary points* and one *center point* midway between the boundary points (see Figure 1). Unlike previous core operators whose *object angle* (between the two radial arms) was permitted to be other than 180° , the object angle of a core atom is fixed at 180° . The *diameter* of a core atom, i.e., the distance between its boundary points, is proportional to the *aperture size* of the convolution kernel used to find the boundary points. Together, the diameter and aperture size constitute the *scale* of the core atom. The process of finding core atoms begins by determining a set of boundary points whose gradient magnitude is greater than some threshold. These boundary points -- their scale, location, and orientation -- form the basic feature set of the core extraction. The boundary points are checked against each other, and core atoms created from pairs that satisfy three criteria: *diameter*, *face-to-faceness*, and *polarity*, which are described in detail below.

Once the core atoms are found, they can be clustered, and the clusters analyzed to yield descriptors that are invariant to scale, translation, and rotation. These descriptors are in the form of *scale spectra*; that is, some statistical measure is plotted with scale along the abscissa. These statistical measures are therefore actually computed on subsets of the cluster, each subset containing core atoms of a given scale. The statistical measures are

- (1) *magnitude*, the number of core atoms at that scale,
- (2) *eccentricity*, a measure of whether the core atoms at that scale tend to be oriented in a uniform direction ($\text{eccentricity} \cong 1$) or evenly distributed in orientation ($\text{eccentricity} \cong 0$), and
- (3) *orientation*, the predominant orientation of the core atoms at that scale.

These scale spectra lend themselves quite readily to analysis of the underlying shape of the core manifold in a manner invariant to translation, rotation, and scale.

3. CORE ATOMS

A *core atom* consists of two *boundary points*, and a *center point* midway between them. The distance between the boundary points is said to be the *diameter* of the core atom.




Figure 1. A core atom contains two boundary points and one center point. The aperture σ of the boundary detection kernels is proportional to the distance between the two boundary points.

The definition of "boundariness" here is simply gradient magnitude $|\nabla I(\mathbf{x}, \sigma)|$, where I is the intensity at location \mathbf{x} measured at scale σ , although other measurements of boundariness could be used. The aperture size σ of the convolution kernel used to measure gradient is proportional to the diameter of the core atom, resulting in stability against image disturbances of a small scale relative to the object width. The image is sampled at an interval proportional to the aperture, and boundary points with gradient magnitude greater than some arbitrary threshold are selected to be candidates for core atom formation. Every possible combination of two candidate boundary points at a given scale is checked to see if it should form a core atom. Three criteria must be met:

1. The *diameter* vector $\mathbf{d}_{1,2}$ from boundary point \mathbf{b}_1 to boundary point \mathbf{b}_2 ,

$$\mathbf{d}_{1,2} = \mathbf{b}_2 - \mathbf{b}_1 \quad (1)$$

must have a Euclidean norm between d_{\min} and d_{\max} , which are proportional to scale.

2. The two boundary points must exhibit sufficient *face-to-faceness*, that is, their gradients must face toward (or away from) each other across the distance between them. The face-to-faceness F of two boundary points is defined as the product of two scalars, f_1 and f_2 :

$$F(\mathbf{b}_1, \mathbf{b}_2) = f_1 \cdot f_2 \quad (2)$$

where

$$\begin{aligned} f_1 &= \hat{\mathbf{d}}_{1,2} \cdot \hat{\nabla} I(\mathbf{b}_1, \sigma) \\ f_2 &= \hat{\mathbf{d}}_{2,1} \cdot \hat{\nabla} I(\mathbf{b}_2, \sigma) \end{aligned} \quad (3)$$

Here we use " \wedge " to denote the normalized vector, e.g.,

$$\hat{\mathbf{v}} = \frac{\mathbf{v}}{|\mathbf{v}|} \quad (4)$$

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 of F near +1 correspond to boundary points facing toward (or away from) each other, across the distance between them. Acceptable pairs of boundary points require *face-to-faceness* to be within some small error ϵ of +1.

$$F(\mathbf{b}_1, \mathbf{b}_2) > 1 - \epsilon \quad (5)$$

3. Assuming $F(\mathbf{b}_1, \mathbf{b}_2) > 0$, it follows that f_1 and f_2 are either both positive, or both negative. The sign of f_1 (or f_2) is called the *polarity*, and measures whether boundaries face toward or away from each other, the exact correspondence being determined by whether the target is lighter or darker than the background. The pair of boundary points must have the appropriate polarity.

We apply these three tests -- *diameter*, *face-to-faceness*, and *polarity* -- to each pair of boundary points at each scale and produce a population of core atoms with a diversity of scales and orientations upon which to perform statistics.

4. SPECTRA OF CORE ATOMS

We develop three spectra for a population of core atoms: *magnitude*, *eccentricity*, and *orientation*. For a given population, the histogram of the number of core atoms at each scale (diameter) is called the *magnitude* spectrum. The magnitude spectrum offers a straightforward representation of the width of the medial axis covered by the cluster. Since only the length of core atoms is involved, the magnitude spectrum demonstrates rotational invariance. An example of a magnitude spectrum is shown in figure 4A, where a dark circular figure has been subjected to the formation of core atoms. The resulting core atoms are superimposed on the dark circle, with the centers shown in white and boundary points shown in black. The 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).

Beyond magnitude, we can develop two other spectra by first calculating the *axis vector* at each scale. The axis vector represents a composite of individual core atom diameter vectors, indicating whether there is a predominant orientation in the core atom population, and if so what that orientation is. Two constraints on the calculation of the axis vector are clear: (1) If the population has no predominant orientation, then each core atom must be canceled, on average, by another core atom perpendicular to the first, and (2) the order in which one selects the boundary points in a core atom is arbitrary, so the orientations of two core atoms 180° apart must be indistinguishable. Both of these constraints are met by the process of summation after angle doubling, in which the *phase angle* (angle with respect to the x axis) of each core atom is doubled. Therefore if the difference in orientation between two core atoms is originally 90° (or 270°), it becomes 180° causing them to cancel, and if the difference is originally 180° , it becomes 360° (or 0°) making them indistinguishable.

The process of calculating the axis vector is as follows:




Figure 2.

Given a core atom whose diameter vector \mathbf{d} has a phase angle α we solve for the vector \mathbf{d}' whose phase angle is 2α , using standard trigonometric identities:

$$\cos 2\alpha = \cos^2 \alpha - \sin^2 \alpha \quad \sin 2\alpha = 2 \sin \alpha \cos \alpha \quad (6)$$

$$\frac{d'_x}{|\mathbf{d}'|} = \frac{d_x^2}{|\mathbf{d}|^2} - \frac{d_y^2}{|\mathbf{d}|^2} \quad \frac{d'_y}{|\mathbf{d}'|} = 2 \frac{d_y}{|\mathbf{d}|} \cdot \frac{d_x}{|\mathbf{d}|} \quad (7)$$

Since $|\mathbf{d}'| = |\mathbf{d}|$, these reduce to

$$d'_x = \frac{d_x^2 - d_y^2}{|\mathbf{d}|} \quad d'_y = \frac{2d_y d_x}{|\mathbf{d}|} \quad (8)$$

We normalize each \mathbf{d}' and sum over the entire population of n core atoms at a given scale, dividing by n to yield the vector \mathbf{a}' , whose components are

$$a'_x = \frac{1}{n} \sum_{i=1}^n \frac{d'_{x,i}}{|\mathbf{d}_i|} \quad a'_y = \frac{1}{n} \sum_{i=1}^n \frac{d'_{y,i}}{|\mathbf{d}_i|} \quad (9)$$

The vector \mathbf{a}' may be said to have a phase angle of 2γ which, when halved, produces what we call the *axis vector* \mathbf{a} (whose phase angle is γ). Halving the phase angle is accomplished using the trigonometric identities:

$$\sin^2 \gamma = \frac{1 - \cos 2\gamma}{2} \quad \cos^2 \gamma = \frac{1 + \cos 2\gamma}{2} \quad (10)$$

$$a_y = \pm |\mathbf{a}'| \sqrt{\frac{1}{2} - \frac{a'_x}{2|\mathbf{a}'|}} \quad \begin{array}{ll} a_y \geq 0 & \text{if } a'_y \geq 0 \\ a_y < 0 & \text{if } a'_y < 0 \end{array} \quad a_x = |\mathbf{a}'| \sqrt{\frac{1}{2} + \frac{a'_x}{2|\mathbf{a}'|}} \quad (11)$$

Attention is paid to signs of the square roots used to calculate a_x and a_y to confine \mathbf{a} to the two right-hand quadrants. Orientation is only meaningful over 180 degrees, since the boundary points in a core atom are order-independent.

The modulus of the *axis* vector \mathbf{a} represents *eccentricity*, which lies between 0 and 1. Eccentricity provides a measure of whether or not the underlying core at that scale is directional, i.e., forms a 1D ridge. During the calculation of \mathbf{a} , angle doubling causes the contributions from perpendicular core atoms to cancel. A uniform distribution of orientations therefore has 0 eccentricity, since each core atom, on average, has another one perpendicular to it. This is shown in figure 3A, where core atoms are evenly distributed around a circular figure. The corresponding experimental data (see figure 4A) shows the eccentricity close to 0, especially at scales where the number of core atoms is greatest. The dip toward zero in the eccentricity spectrum corresponds to the maximum number of core atoms, a relationship predictable from the \sqrt{n} property of noise.




Figure 3. The two extremes of eccentricity. **A.** Core atoms with evenly distributed orientations. **B.** Core atoms aligned along a 1D ridge.

At the other extreme, an eccentricity of 1 is shown in figure 3B. Here the core atoms have the same orientation, resulting in an axis vector \mathbf{a} with a modulus of 1. Since \mathbf{a} has a non-zero length, we may also report its direction. The underlying direction of the ridge will be perpendicular to the orientation of the core atoms. Doing this at each scale results in a third spectrum, that of *orientation*. Although the orientation spectrum shifts vertically when the figure is rotated as shown in figure 5, the difference between orientations as a function of scale remains fixed, constituting a useful form of rotational invariance.

Figure 4B shows an ellipse. The magnitude spectrum is now split into two distinct humps. Within each hump, the eccentricity is near 1, but the orientations are 90° apart. We recognize these humps to 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. The drop in eccentricity at scales between the humps results from perpendicular populations of core atoms at scales between those of the major and minor axes of the ellipse canceling each other.

Note that the display for the orientation spectrum is duplicated modulo- 180° and that it wraps around from -180° to $+180^\circ$. The orientations are printed lightly (as in figure 4A) at scales where insufficient eccentricity exists. Likewise, eccentricity is printed lightly at scales where insufficient magnitude exists. Also, note the large number of core atoms involved, e.g., 4016 for the circle, which is close to the total number of pixels in the image! This is due to the fact that many boundary points form core atoms with multiple partner points across the figure.

5. GRAPHICAL TEST FIGURES

In addition to the circle and the ellipse already described, other graphical test objects have been subjected to core atom formation and their spectra included for review.

Figure 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 everywhere near 1, since at each given scale the sub-population is uniformly oriented across the paisley. Larger scales have been excluded to avoid the formation of longer core atoms that would obscure the illustration.

Figures 4D, 4E, and 4F show the effect of varying a figure's scale. A vertical rectangle is shown in 4D resulting in two distinct core atom populations, each with eccentricities essentially of 1. The first population consists of small scale core atoms across the narrow axis of the rectangle from left to right with 0° orientation, resulting in a vertical core (remember, the white dots show core atom centers). The second consists of larger scale core atoms traversing the long axis of the rectangle from top to bottom with 90° orientation, resulting in a horizontal core. The two cores cross in the center of the figure, although they miss each other when scale is added as a third dimension, in so-called *scale space*. The direction of the local core is approximately perpendicular to the average orientation of the core atoms. Subsequent reduction in the rectangle's size in figures 4E and 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 to the left would occur.

Figures 5A-F show the effect of rotation on the rectangle figure. The magnitude and eccentricity spectra are essentially unaffected. The orientation spectrum shifts upward as the figure is rotated in a counter-clockwise direction. Note that the relationship between orientations at different scales remains unaffected.

Figure 6A shows the effect of Gaussian noise added to the rectangle figure. As before, white dots represent core atom centers and black dots, core atom boundaries. Core atoms are dispersed throughout the image as a result of the added noise. Comparing the spectra of 6A to those of the same rectangle without noise in figure 6D, it can be seen that most of the noise occurs at small scales. At larger scales the aperture size filters out the noise. Figures 6B and 6C show the results of a simple segmentation routine which extracts clusters of core atoms with the following criteria: (1) the maximum distance between the centers of any two core atoms and between the scales of those two atoms are below certain thresholds, and (2) the cluster contains at least a certain number of core atoms. The two resulting clusters are shown in 6B and 6C, corresponding to the two cores described in figure 4D. The spectra of 6B and 6C combine to approximately equal that of 6D, the rectangle without noise. (Figures 4D and 6D differ in the settings for *gradient* and *face-to-faceness* thresholds). A small amount of noise remains connected to the cluster in the lower corner of the rectangle in figure 6B.

Figures 6E and 6F demonstrate the ability of the spectra to differentiate between a heart and a spade, which are virtually identical except for the stem of the spade. At small scales for the spade in figure 6F, a cluster of core atoms across the stem results in an eccentricity near 1 and an orientation of 0° .

Figures 7A-F show the effect of rotating the spade. The magnitude and eccentricity spectra remain essentially unchanged, while the orientation spectra progress upward. The strongest signal in the orientation spectra is at small scale, where the stem accumulates a cluster with high eccentricity.

6. CONCLUSION

Our approach is novel in several respects. Rather than filtering directly for locations with high medialness, we gain efficiency by finding individual boundary points and then linking appropriate pairs. We do not find ridges, but simply identify clusters. Unlike methods which find medial axes by following ridges, we apply non-stringent criteria to core atoms, searching throughout the image in an initial sweep and gathering enough to perform significant statistics, without user assistance. Core atom populations may be segmented by center location and scale to reduce noise and isolate individual figures or cores. We introduce the concept of scale spectra as an invariant feature of the underlying figure.

Our current research aims to generalize these methods to 3D and to use them to identify and measure cardiac chambers within real time 3D ultrasound data, which requires complete automation and relative efficiency to be clinically useful.

ACKNOWLEDGMENTS

Supported through a Whitaker Foundation Biomedical Engineering grant, NSF grant CDR8622201, and NIH grants 1K08-HL03220 and P01-CA47982.

REFERENCES

1. H. Blum, R. N. Nagel, "Shape description using weighted symmetric axis features," *Pattern Recognition* **10**, 167-180 (1978).
2. L. R. Nackman, "Curvature relations in 3D symmetric axes," *CVGIP* **20**, 43-57 (1982).
3. L. R. Nackman, S. M. Pizer, "Three-Dimensional shape description using the symmetric axis transform I: Theory," *IEEE Trans. PAMI* **2**, 187-202 (1985).
4. C. A. Burbeck, S. M. Pizer, "Object representation by cores: Identifying and representing primitive spatial regions," *Vision Research* **35**, 1917-1930 (1995).
5. S. M. Pizer, D. H. Eberly, B. S. Morse, D. S. Fritsch, "Zoom invariant vision of figural shape: the mathematics of cores.," *Technical Report TR96-004, University of North Carolina, 1996*. To appear in *Computer Vision and Image Understanding*, (1996).

Stetten, Landesman, and Pizer

SPIE Medical Imaging '97

page 8

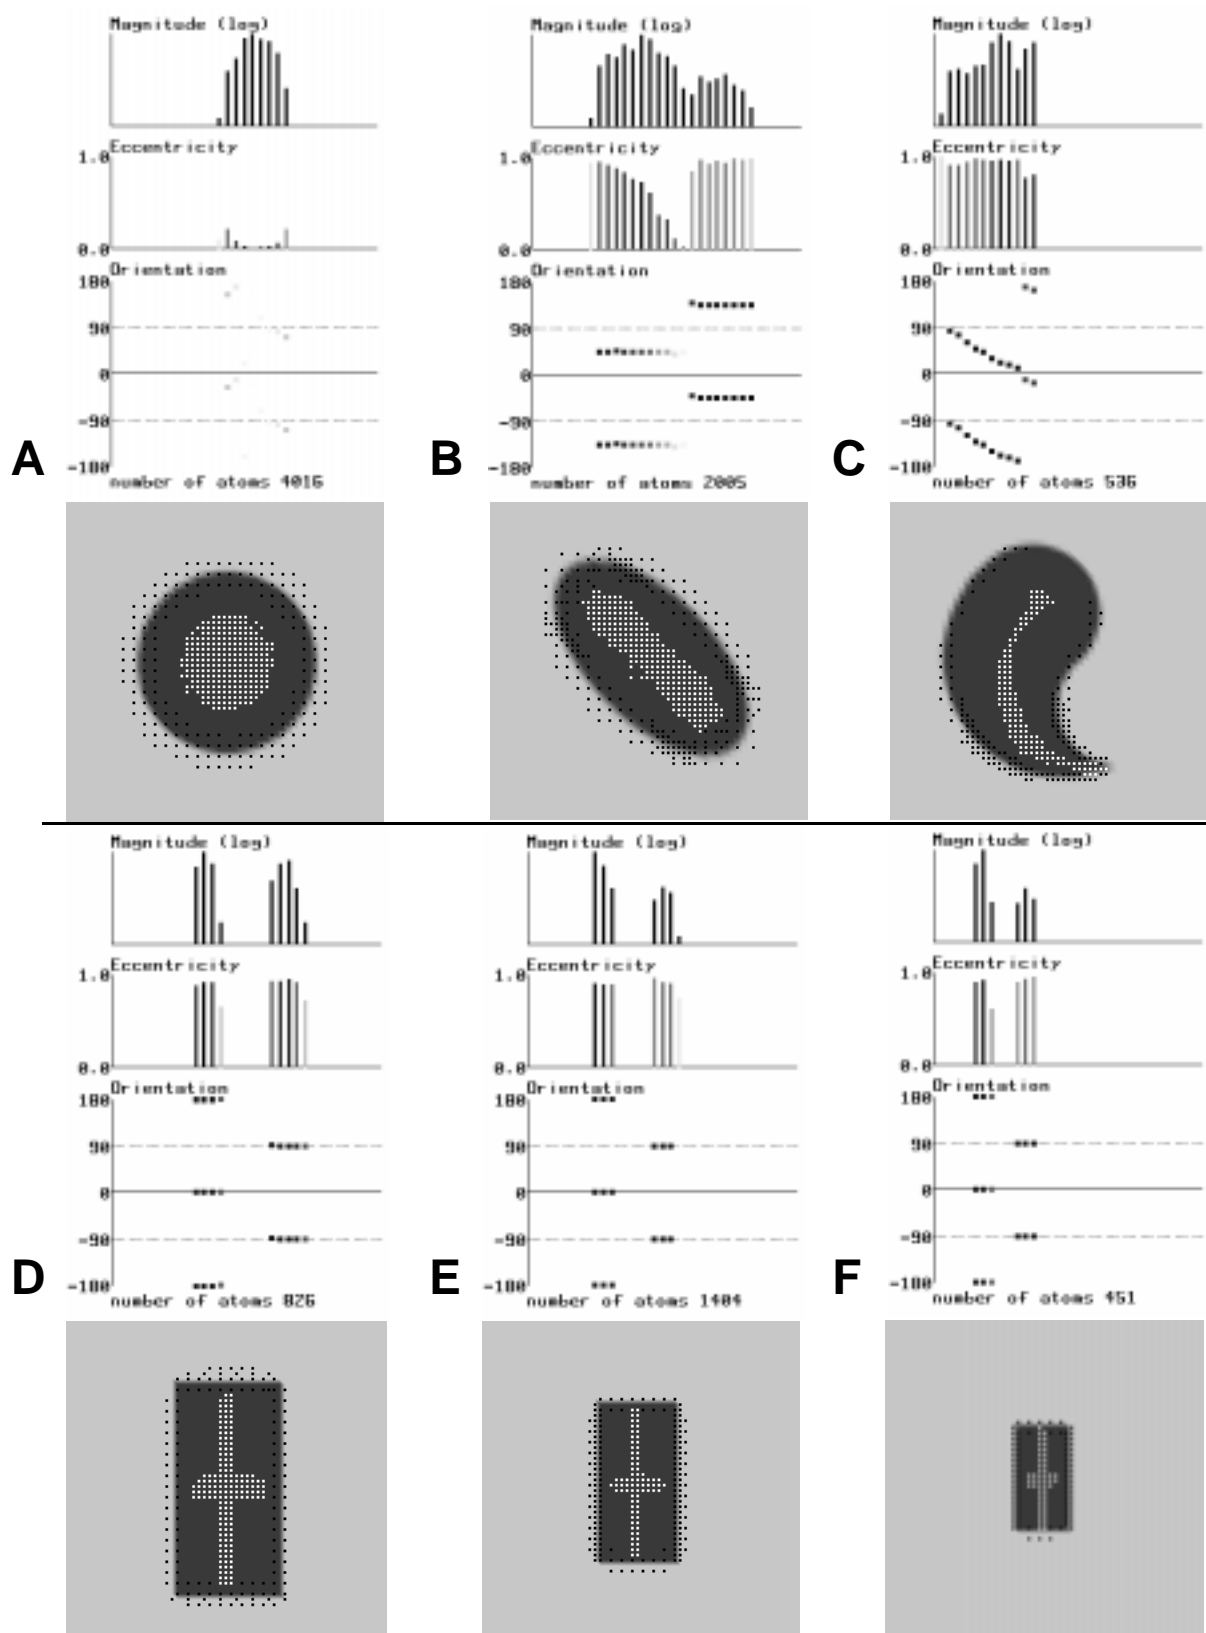


Figure 4

Stetten, Landesman, and Pizer

SPIE Medical Imaging '97

page 9

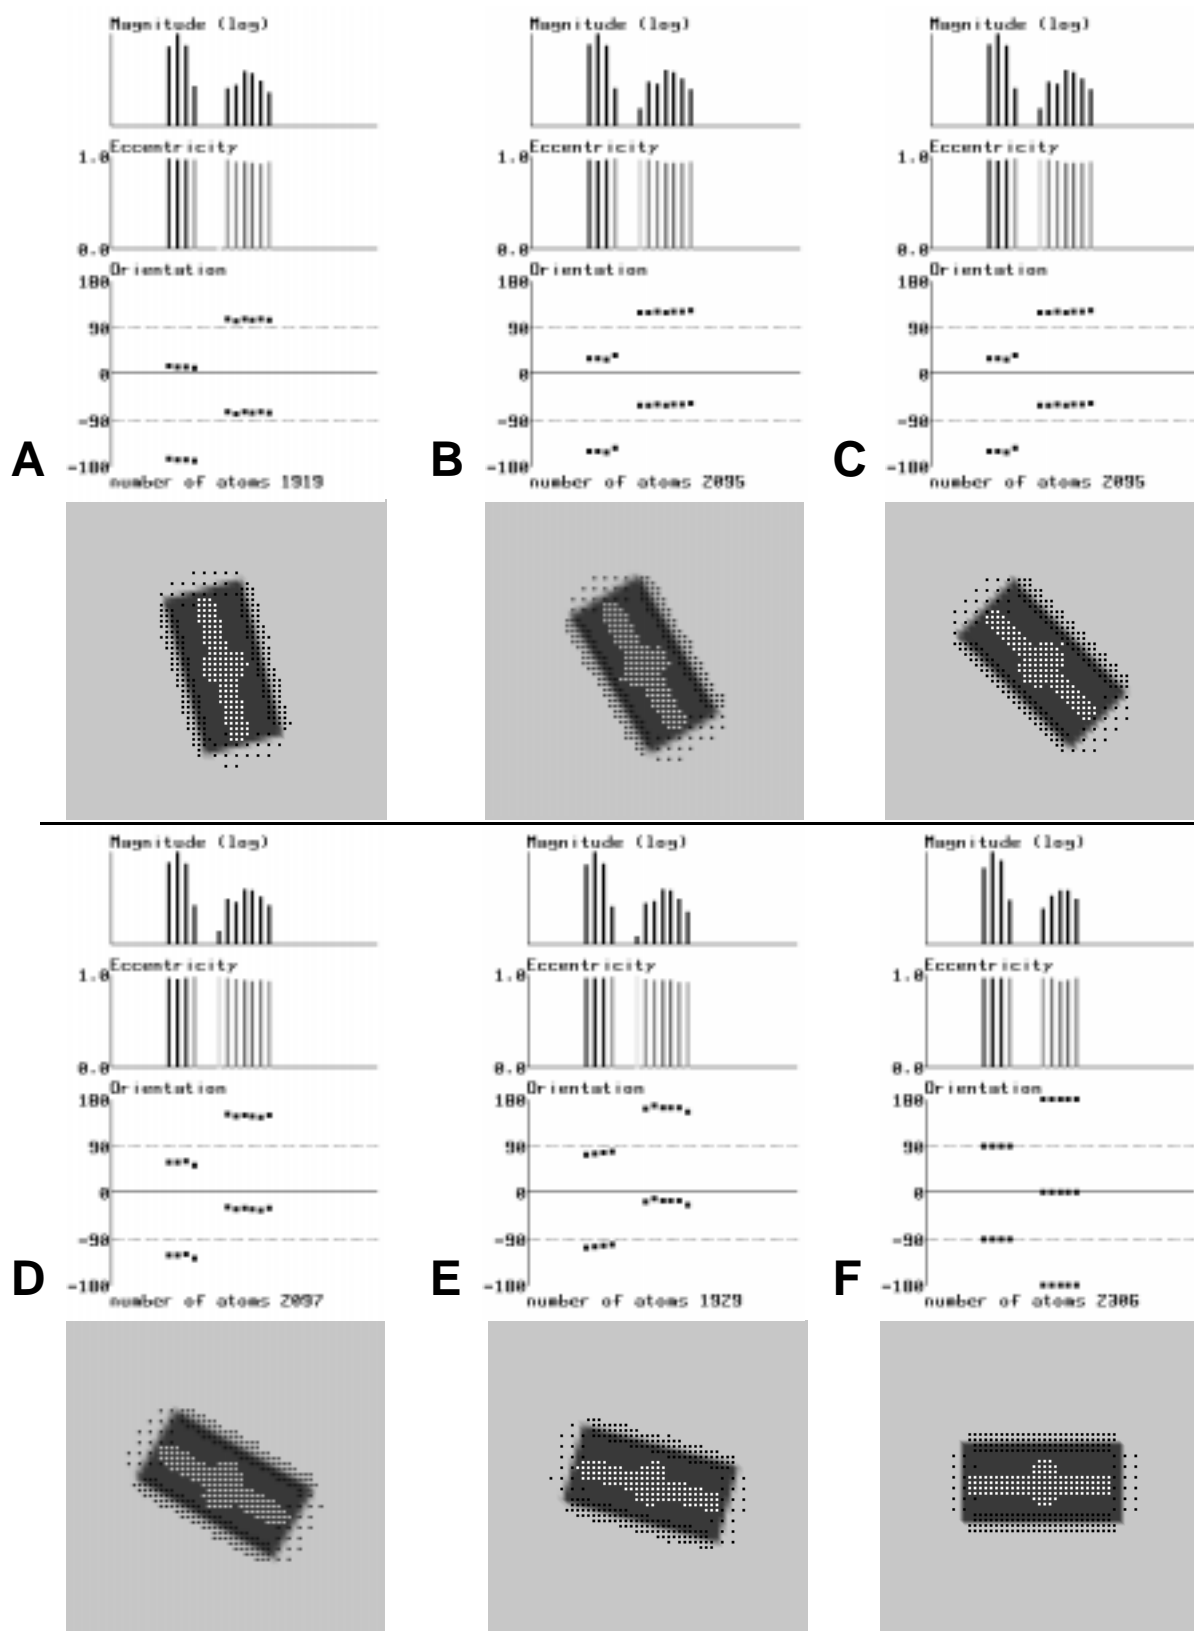


Figure 5

Stetten, Landesman, and Pizer

SPIE Medical Imaging '97

page 10

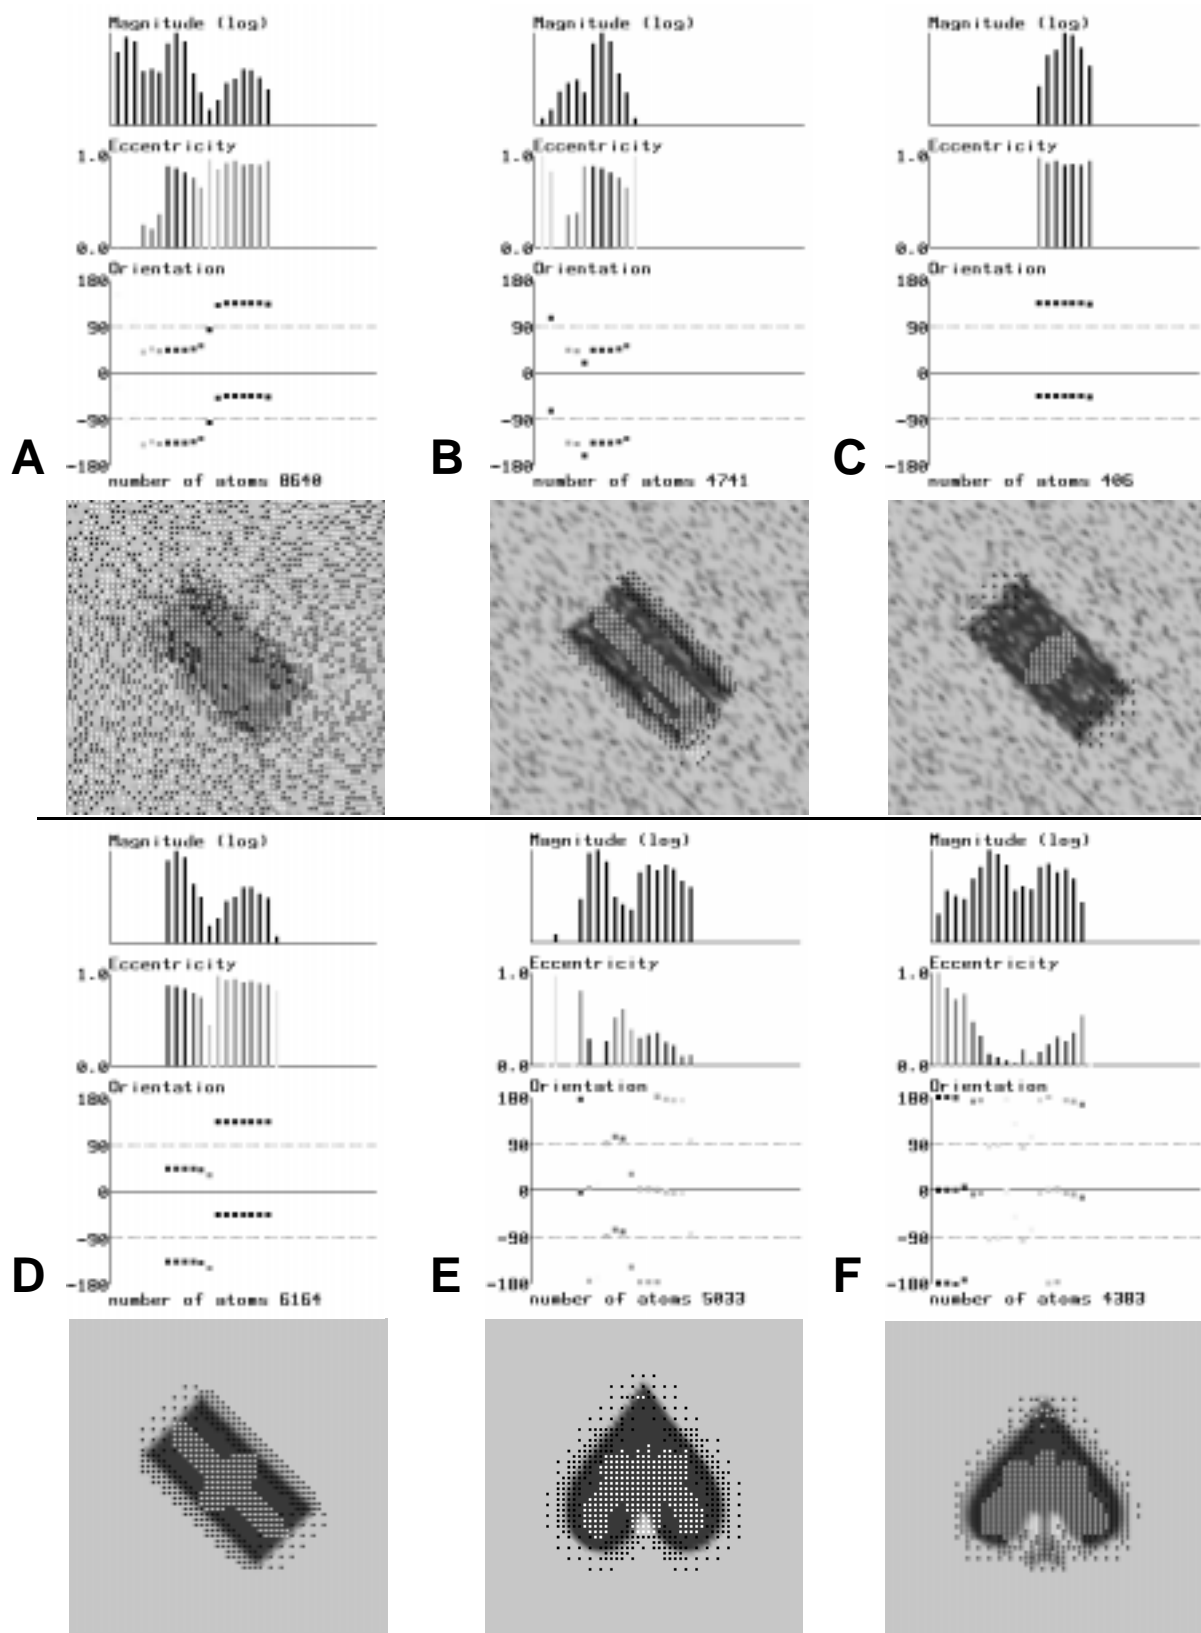


Figure 6

Stetten, Landesman, and Pizer

SPIE Medical Imaging '97

page 11

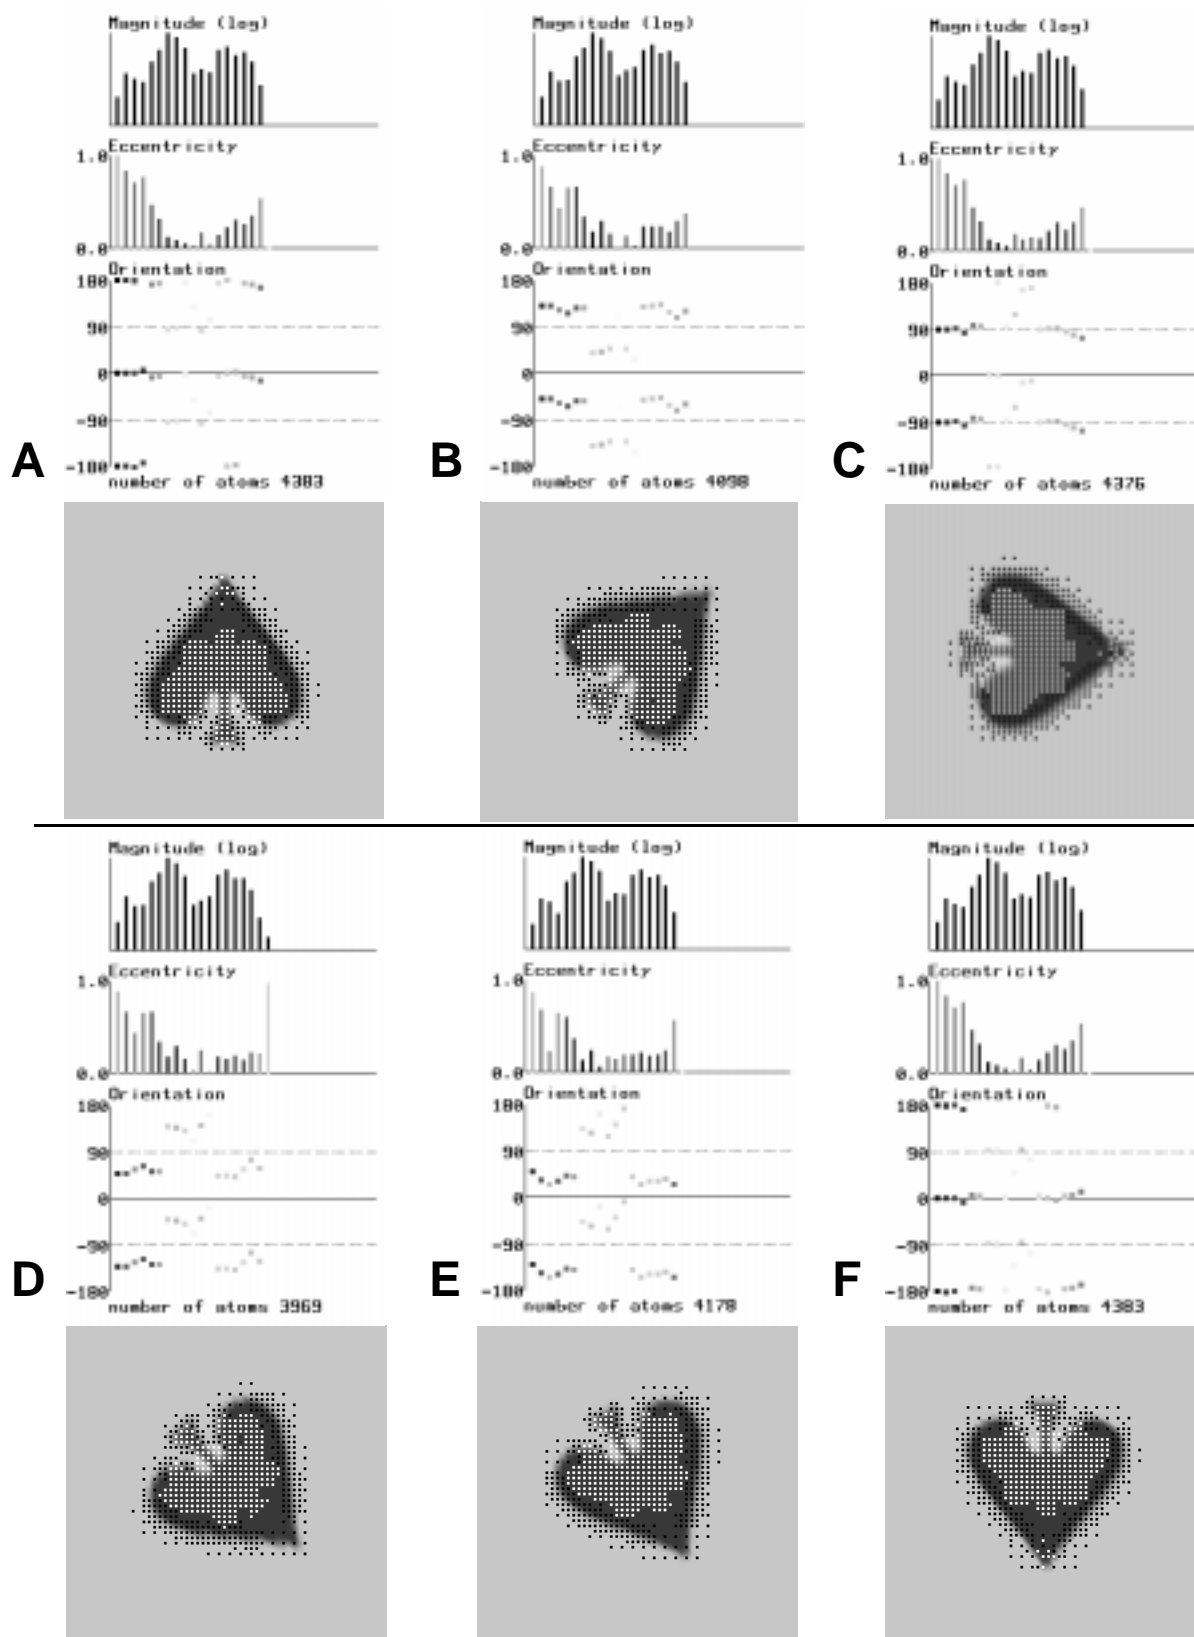


Figure 7