

Chapter 5 Medial Node Models

In the preceding two chapters, a method was developed for extracting local properties of shape using core atom clusters. Because the measurement of these local properties has been found to be stable, it can be used to match underlying regions in the image to nodes in a model of anatomical structure. The main subject of this chapter is how to perform this matching using the expected local medial properties and the geometric relationships between them. The chapter also introduces an additional property, *endness*, whose value is high at the cap of a cylinder and the edge of a slab. An expected location and orientation for endness can be provided by a single core atom cluster or by a configuration of core atom clusters. These tasks are organized by means of a *medial node model* (MNM), which defines the expected medial properties at discrete nodes and their geometric relationships in space.

Consider, for example, the simple model of the cardiac left ventricle (LV) introduced in Fig. 1.2. The model consists of 3 nodes: (1) a large dark cylindrical LV, (2) a thin bright slab-like mitral valve at one end of the LV, and (3) a hemispherical cap at the other end of the LV (the ventricular apex). To find these structures in the image, clusters of core atoms are matched to the properties of the LV and mitral valve, described in the first nodes of the MNM. Pairs of clusters are then selected that also match the location, orientation, and scale of the LV and mitral valve relative to each other as predicted by the MNM. From these *pairs* of clusters, an optimum location for the LV and mitral valve in the image is determined, yielding an axis along which to measure *endness* at the apex of the LV. The expected properties of endness at the apex are defined by Node 3 in the MNM. These properties are measured in the image using *core ions*. Core ions are similar to a core atom, but possess only one boundary point.

The details of applying the particular MNM for the LV are described in subsequent chapters. The present chapter concentrates on developing the theory behind matching core atom clusters to an MNM in general and measuring endness based on initial placement of medial nodes.

5A. Homologies Between Core Atom Clusters and Nodes

A medial node model consists of discrete nodes, each node describing the expected properties of core atom (or core ion) populations at a particular section in the anatomical structure being sought. The first step in fitting an MNM to a given image is collecting and clustering core atoms, as developed previously. Chapter 4 described the process of initially collecting core atoms in bins on a regular lattice and then clustering them based on directionally weighted voting between samples in a region. For the purposes of producing clearer illustrations (Figs 4.8 and 4.10), a vote-removal scheme was employed so that only clusters with non-overlapping constituencies would be displayed. When forming clusters to match against a model, however, such a vote-removal scheme is *not* used so that as many clusters can be matched against the model as possible. As a result, neighboring clusters usually represent overlapping populations of core atoms and therefore can be expected to yield multiple matches to a single node in the model. This opens the way for a Hough-like voting strategy to establish homologies between regions in the image and nodes in the model.

The process of matching individual nodes and clusters implies the use of a *single-node* metric, a measure of distance in a parameter space between values expected by a node and those measured by a cluster. The single node metric reflects differences between the model and measured values with respect to medial scale, dimensionality, and other parameters. Some of these parameters are implicit in the overall process of forming the core atom clusters and have already been applied by the time a given cluster is compared to a given node. For example, a node may have an expected scale corresponding to a particular range core atom lengths. If only core atoms with this range of lengths are originally collected, the resulting clusters will already match the expected scale of the node. Likewise, if a certain absolute orientation is expected for a node, then the core atom selection process can be limited to core atoms with the corresponding absolute orientation. For example, a closed mitral valve can be expected to have an absolute orientation roughly parallel to the face of the transducer. Care must be taken when limiting core atom orientation, since it may distort the dimensionality of the resulting clusters. Although pre-selection of core atoms may need to be repeated with different parameters for different nodes in the model, it can still save time, especially when the total number of nodes is small.

Other implicit parameters that effect the single-node metric include the boundariness threshold in selecting candidate boundary points, the face-to-faceness threshold and the polarity employed in selecting core atoms, the σ parameter used to determine the extent of ellipsoidal voting (eq. 4.1), and the spacing of bins for collecting boundary points and core atoms.

If the scale parameter is not already implicit in the pre-selection of core atoms, it can be applied explicitly when matching nodes and clusters by comparing the mean scale of the core atoms in a cluster to the scale expected by the node. The other important parameter of the single-node metric is dimensionality, defined by the two eigenvalues λ_1 and λ_2 within the lambda triangle (Fig. 3.9). The dimensionality parameter can only be applied after clustering.

No specific formulae will be developed here for the single-node metric, since the optimum metric depends on the particular application and even the particular node. The general requirement of the metric is that it be able to select appropriate clusters to match each node, and report some form of goodness-of-fit for each match.

5B. Homologies between Cluster Pairs and Node Pairs

Once candidate clusters have been identified by the single-node metric corresponding to individual nodes of an MNM, pairs of clusters can be tested against pairs of nodes in the MNM to see if they match the model in terms of the expected geometric relationships for each pair. The expected pair-wise relationship between two clusters in an MNM is expressed by a *dual-node* metric, whose parameters describe relative location, orientation, and scale.

In the general theory for the dual-node metric the expected scale of individual clusters and the expected distance between those clusters can be dependent upon each other. For example, assume the aorta is locally a cylinder and the right ventricle is locally a slab. In the average heart, the aorta may be expected to pass at a given distance and with a given relative orientation to the right ventricle. If the heart is larger than normal, the right ventricle and the aorta should both be correspondingly larger and further apart. (In the simple MNM developed subsequently for the LV, expected distance between the LV and the mitral valve is independent from the scale of the individual clusters).

As illustrated in Figs. 5.1-5.3, relating the orientations of clusters in a pair to each other must take into account the orientations of the individual clusters as well as the orientation of the *inter-cluster vector*, which joins the two clusters (see Fig. 5.1). The orientation parameters of the dual-node metric depend upon the dimensionality of the individual clusters, since the orientation of a cluster cannot even be defined without considering its dimensionality. For example, while the orientation of a cylinder is clearly defined as its axis ($\hat{\mathbf{a}}_1$ in Fig. 3.8), that of a slab can only reasonably be defined as its normal ($\hat{\mathbf{a}}_3$ in Fig. 3.8). A sphere has no orientation at all. When describing the relative orientation between clusters, such matters must be considered. A pair of clusters consisting only of slabs or cylinders (see Fig. 5.1) requires 3 parameters to define relative orientation. Two angles are needed between 0° and 90° that may

be thought of as *nose-dive*, how orientation changes moving from each of the clusters to the inter-cluster vector. A third parameter defines the *twist* between the two clusters around the inter-cluster vector, ranging from 0° to 180° . The twist parameter is the only MNM parameter possessing *chirality*, that is, the property of being right or left handed. Thus the twist parameter enables the metric to differentiate an object from its mirror image. When the nose-dive of either cluster in a pair is 0° , the twist cannot be determined. Likewise, a cluster pair containing a sphere as one of its clusters requires no twist and only 1 nose-dive (Fig 5.2). A cluster of pairs consisting of two spheres has no relative orientation parameters at all (Fig. 5.3).

An interesting analogy exists between a core atom and a cluster pair. Each consists of a pair of elements, respectively, of boundary points or clusters. In each case, the elements come from a population originally chosen without regard to orientation or location. The elements are paired using a metric that is sensitive to relative orientation and distance. This general pairwise approach is a central theme of this dissertation.

5C. Endness

Certain medial locations are special in that they represent terminations of medial ridges (Leyton 1992). The property of *endness* has been described by Clary, et al., who used it to identify the apex of the cardiac left ventricle (Clary, Pizer et al. 1997). The concept of endness is illustrated using 2D core atoms in Fig. 5.4. The 1D core within a stripe is blocked by a boundary at the right-hand end of a stripe. A single boundary point associated with the right-most core atom is oriented so that it faces the natural extension of the core -- a medial stop-sign, as it were. Such a boundary point, called a *core ion*, lacks the second boundary point of a core atom and is paired instead with the center of a core atom cluster.

This concept of core ions is extended to three dimensions in Fig. 5.5. Two possible configurations are shown. The cylinder has a 1D axial core that may encounter a core ion in either of two directions along the axis. The slab has a 2D core that may encounter core ions in any direction co-planar with the core. A sphere (not shown) already has core atoms in all directions, and thus will not experience endness.

These observations lead directly to a method for detecting endness. Clusters of core atoms can be used as vantage points from which to look for core ions. For example, once a cylindrical cluster of core atoms has been established, core ions can be sought along the axis of the cylinder in either direction as evidence of a cap on the cylinder. Similarly, once a local slab

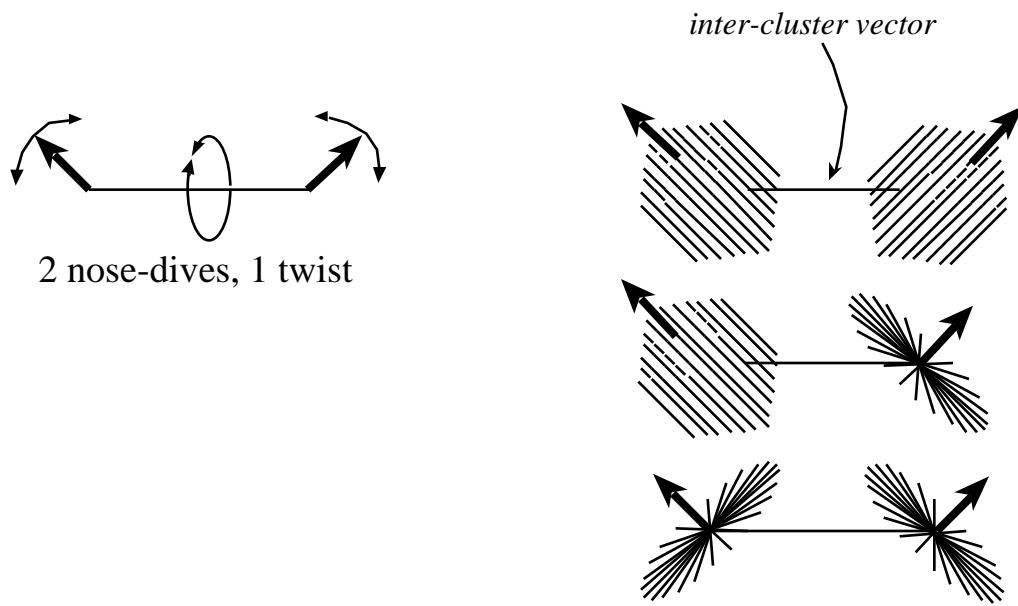


Fig 5.1 Slabs and cylinders have individual orientations relative to the *inter-cluster vector*, yielding 3 relative orientation parameters for the pair.

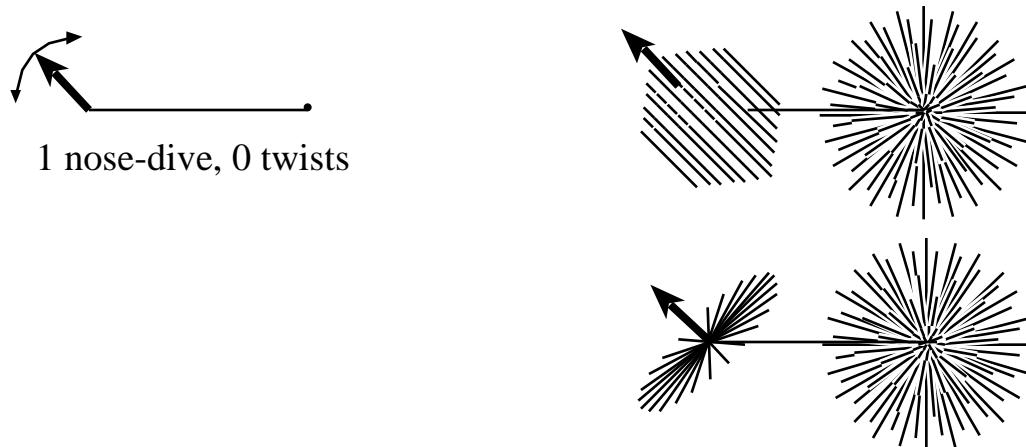


Fig 5.2 Pairs with one sphere have only 1 relative orientation parameter.

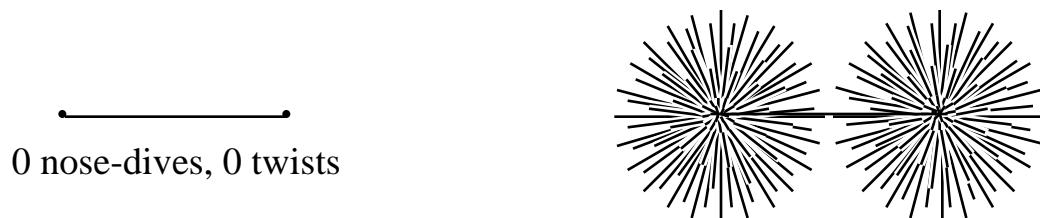


Fig 5.3 A pair with two spheres has no relative orientation parameters.

has been found, boundary points indicating an edge on the slab can be sought in all directions coplanar with the core within the slab. As will be seen in the MNM of the left ventricle, two nodes can cooperate to decrease the uncertainty in endness orientation. Thus the axis between the LV cylinder and the mitral valve slab defines a more accurate orientation along which to find the LV apex than the LV cylinder alone.

A core ion is identified from candidate boundary points by requiring that its relative orientation and location terminate the local core. Once gathered, core ion populations may be analyzed to yield endness parameters, which can then be matched against nodes describing expected endness in a medial node model. The distance from the appropriate boundary point to the center of the core atom cluster should be about half the mean core atom diameter. The boundary point should face a direction of maximum orthogonality to the core atoms in the cluster. The eigenanalysis developed in Chapter 3 can determine if an eigenvector has an eigenvalue near zero (the center of the stripe, axis of the cylinder, or plane of the slab). Such an axis (or axes) defines a general subspace in which to find core ions.

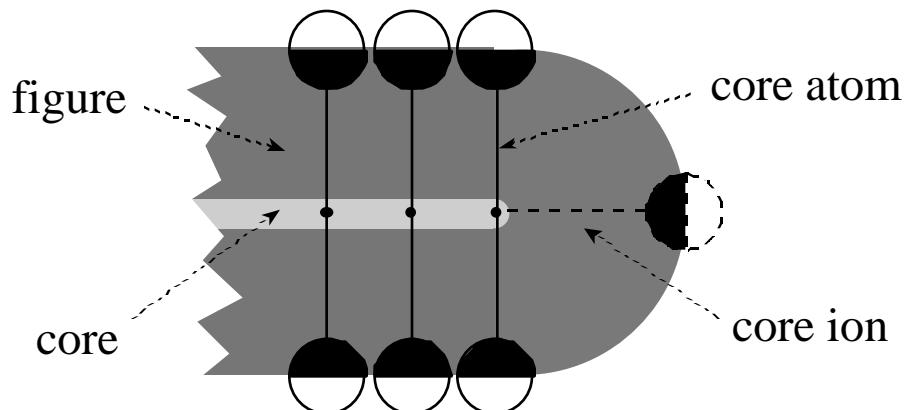


Fig. 5.4 Endness, manifested as the end of a stripe, and detected as *face-to-medialness* by a *core ion*. In 2D, this is the only type of endness.

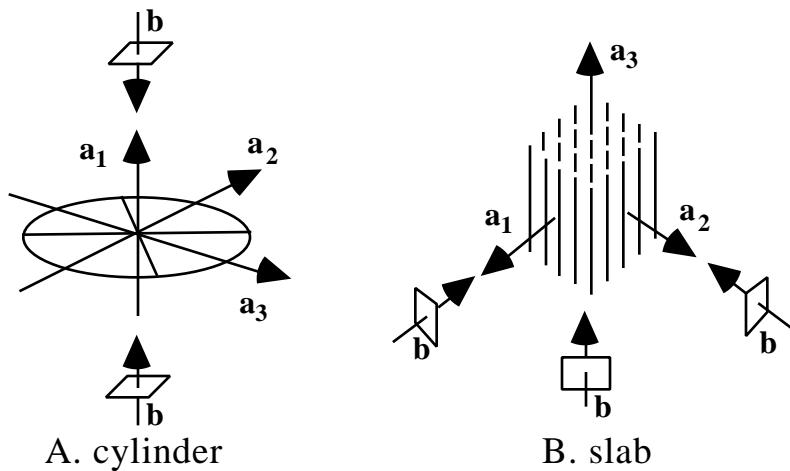


Fig. 5.5 Endness, manifested by a cap on a cylinder (A) and the edge of a slab (B). Boundary points contributing to endness are labeled \mathbf{b} and eigenvectors of core atom cluster labeled \mathbf{a}_1 , \mathbf{a}_2 , and \mathbf{a}_3 as in Fig. 3.8.

The requirements for the formation of a core ion can be stated in three equations analogous to those defined for the core atom in Section 3.1. The mathematics are developed here in general for n dimensions.

(1) For the vector $\bar{\mathbf{r}}$ from the center of the core atom cluster to the boundary point \mathbf{b} , the magnitude

$$r = |\bar{\mathbf{r}}| \quad (5.1)$$

must be within some error ϵ_r of half the mean core atom diameter \bar{c} in the cluster.

$$\left| r - \frac{\bar{c}}{2} \right| < \epsilon_r. \quad (5.2)$$

(2) The face-to-medialness $F(\mathbf{b}, C)$ between boundary point \mathbf{b} and core atom cluster C must be sufficiently high, such that the boundary normal $\hat{\mathbf{n}}$ faces the natural extension of the core across the intervening distance. As with ellipsoidal voting in Chapter 4, eigenanalysis of core atom clusters can be used to determine the natural extension of the core. Face-to-medialness can be defined as

$$F(\mathbf{b}, C) = |f_1 \ f_2| > 1 - , \quad f_1 = \sum_{i=1}^{m-1} (\hat{\mathbf{a}}_i \cdot \hat{\mathbf{r}})^2, \quad f_2 = \hat{\mathbf{n}} \cdot \hat{\mathbf{r}} \quad (5.3)$$

where $\hat{\mathbf{a}}_i$ is one of the $m - 1$ eigenvectors with the smallest eigenvalues λ_i and the coefficients $e_i = 1 - 2\lambda_i$ range from 0 to 1.

(3) As in core atoms, the sign of f_2 is called the *polarity* and is either + or - depending on whether the expected target is lighter or darker than the background..

Multiple core ions can be associated with a single core atom cluster. Analyzing a population of core ions yields parameters describing the orientation and magnitude of the underlying endness. Combining these parameters with the medial dimensionality of the core atom cluster permits differentiation (in 3D) between the cap of a cylinder and the edge of a slab.

Assume a population of n core ions have been collected for a given core atom cluster. A radius vector $\vec{\mathbf{r}}$ has been calculated from the center of the cluster to each core ion. Endness coefficients e_i between $+n$ and $-n$ can then be found corresponding to the first $m - 1$ eigenvectors

$$e_i = \sum_{j=1}^n \hat{\mathbf{a}}_i \cdot \hat{\mathbf{r}}_j, \quad 0 < i < m - 1. \quad (5.4)$$

The sign of e_i is arbitrary, since the unit basis vector from the eigenanalysis of a core atom cluster can equally well be $\pm \hat{\mathbf{a}}_i$. Once assigned, however, the sign of e_i is essential for computing the endness vector $\vec{\mathbf{e}}$ as a linear combination of those $m - 1$ basis vectors.

$$\vec{\mathbf{e}} = \sum_{i=1}^{m-1} e_i \hat{\mathbf{a}}_i \quad (5.5)$$

The endness vector $\vec{\mathbf{e}}$ inhabits the same subspace as the core itself and is oriented in the direction of maximum endness. The magnitude $|\vec{\mathbf{e}}|$ equals the number of core ions and provides a scalar measure of endness that can be incorporated into the single-node metric. Thus the model can use $|\vec{\mathbf{e}}|$ to differentiate the cap of a cylinder from the middle of a cylinder. The orientation $\hat{\mathbf{e}}$, in turn, offers an additional orientation parameter to the dual-node metric. Thus the MNM can specify which end of a cylinder has a cap or in which direction a slab has an edge. If both nodes in a pair of nodes have endness, an additional twist parameter can be specified in the dual node metric (see Fig. 5.6).

Important anatomical locations often include endness in addition to medialness, such as the apex of the left ventricle or the rim of the mitral valve. Medial nodes incorporating endness will be anchored at terminal points on the medial ridge, their location being constrained by one more dimension than the ridge itself. Medial nodes with endness also add an orientation constraint relative to other nodes in the model and an extra orientation constraint (twist) relative to other nodes with endness.

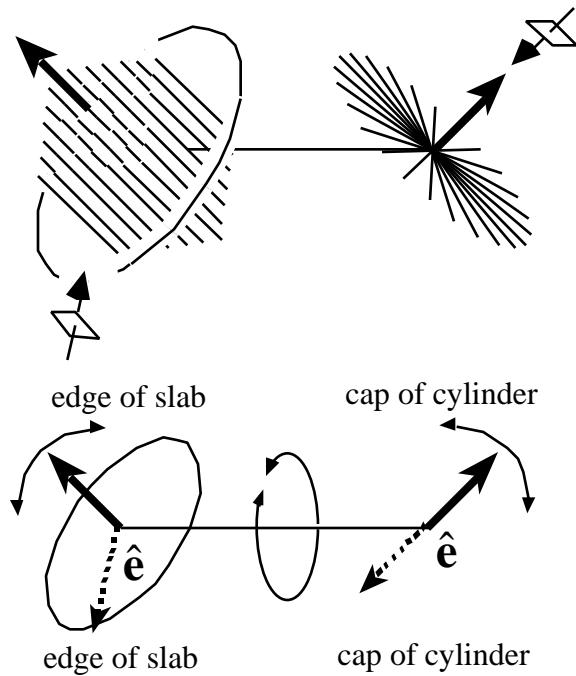


Fig. 5.6 Endness introduces an additional parameter into the dual-node metric for matching node pairs to cluster pairs. The additional parameter can be thought of as a normalized endness vector \hat{e} , which is a linear combination of \hat{a}_1 and \hat{a}_2 for the edge of a slab and $\pm\hat{a}_1$ for the end of a cylinder. If both nodes have significant endness, an additional twist parameter may be specified.