

5. PORTAL IMAGE CONTRAST ENHANCEMENT and TREATMENT VERIFICATION

A second imaging process for which the image quality model was tested is the contrast enhancement of portal images. Performance for the task of treatment field positioning was studied as parameters of the enhancement algorithm SHAHE were varied. Sections 5.1 through 5.3 describe respectively the portal imaging process and the SHAHE algorithm, the treatment field distance estimation task, and the production of images for the experiments in this research. Finally, the implementation of the core model that was developed for the distance estimation task is documented at the end of the chapter in Section 5.4.

5.1 Portal Imaging and SHAHE

In the delivery of a radiation dose for treatment of malignant tumors, the paramount interpretive consideration is the accuracy with which the radiation beam is delivered to the tumor while avoiding adjacent healthy or radiation-sensitive tissue. The geometric accuracy of the treatment beam, or field, is verified by the acquisition and consultation of a radiographic film, the "portal film," exposed by the exiting treatment beam (Figure 5.1).¹ Portal films are acquired continually throughout the treatment regimen to ensure consistent and safe delivery of radiation dosage.

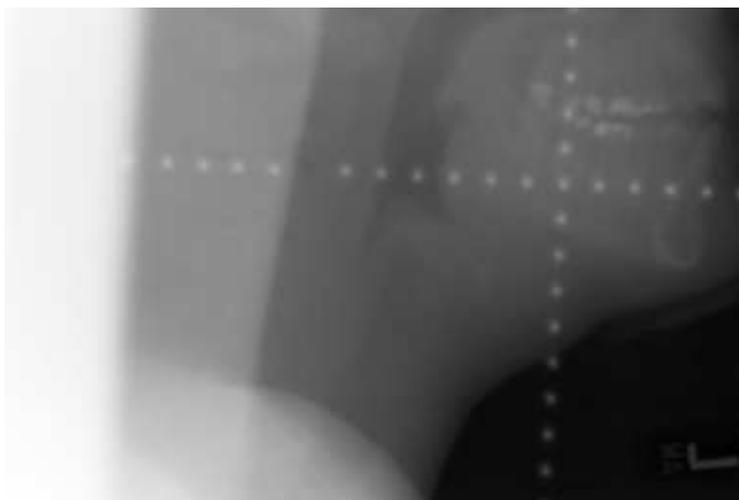


Figure 5.1. A portal image. This image is cropped from a digitized portal film of the head and neck region. The prominent vertical edge in the middle of the image is the edge of the treatment beam; the interior of the beam lies to the right of the edge in the darker, more exposed, portion of the image. The shoulder (lower left) and jaw (upper right) may also be recognized in the image.

The radiation beam used in tumor treatment is characterized by very high energy photons--it is the high energies that are needed to actually destroy diseased tissue. However, standard radiographic imaging relies on the differential attenuation, or "stoppage," of photons in a lower energy beam to form an image with good contrast. Unfortunately the portal film resulting from exposure by the treatment beam therefore possesses little contrast. Several researchers have noted the benefit of image processing algorithms for enhancing the contrast of the digitized portal film.^{2,3} In particular, SHAHE ("sharpened" AHE),⁴ an image processing algorithm that utilizes a

¹F.M. Khan, *The Physics of Radiation Therapy* (Baltimore: Williams & Wilkins, 1984).

²G.W. Sherouse, J. Rosenman, H.L. McMurry, S.M. Pizer, E.L. Chaney, "Automatic Digital Contrast Enhancement of Radiotherapy Films," *International Journal of Radiation Oncology Biology and Physics* 13 (1987): 801-806.

³K.W. Leszczynski, S. Shalev, N.S. Cosby, "The Enhancement of Radiotherapy Verification Images by an Automated Edge Detection Technique," *Medical Physics* 19, no. 3 (1992): 611-621.

⁴R. Cromartie, S.M. Pizer, "Edge-affected Context for Adaptive Contrast Enhancement," *Information Processing in Medical Imaging XII*, A.C.F. Colchester, D.J. Hawkes, eds., Lecture Notes in Computer Science 511 (1991): 474-485.

form of digital unsharp masking⁵ as a preprocess to contrast-limited adaptive histogram equalization (CLAHE),⁶ has been shown to significantly benefit the task of portal and simulator film alignment.⁷

CLAHE is a locally-adaptive histogram equalization method that attempts to optimally transfer the greyscale information in the image to the observer and in effect enhance the contrast of anatomical structures. The algorithm of histogram equalization by itself is a remapping by which the histogram of the image is flattened so that the pixel intensities in the image are uniformly distributed among the grey levels available for display. This remapping, accomplished by assigning a new intensity value that is the rank of the original intensity value in the cumulative histogram for the entire image, devotes proportionately large numbers of greylevels to intensity values with high frequencies of occurrence. Regions of the image with small yet frequent intensity variations are rendered with greater contrast. Adaptive histogram equalization performs histogram equalization within a contextual region surrounding the pixel under consideration for remapping, and thus enhances contrast with respect to local image structure. The contextual region size may be adjusted as dictated by *a priori* input about the size of the objects to be enhanced. Finally, contrast-limited adaptive histogram equalization allows limitation of the extent of contrast enhancement achieved. The amount of contrast enhancement is proportional to the slope of the input-to-output mapping: a slope of 1 generates the original input with no enhancement, and greater slopes produce increasingly greater contrast. The way that this slope is controlled to limit the amount of contrast enhancement is by “clipping” the histogram, or limiting the number of pixels allowed at any given histogram bin.

The enhanced contrast provided by CLAHE appears to be quite beneficial in interpreting the low contrast portal films. However, standard CLAHE exhibits an artifact that causes high contrast edges, such as the treatment field edge found in every portal film, to be blurred, “smeared,” or “shadowed” (Figure 5.2). The artifact becomes increasingly objectionable as the contrast is limited less and as the contextual region size is made smaller. Since almost all the decisions made with the portal film are made with respect to the treatment beam edge, the blurring of that important feature by CLAHE is especially unwanted. The particular form of unsharp masking utilized by SHAHE as a preprocess to CLAHE attempts to counteract this artifact.



Figure 5.2. AHE edge artifact. The broad, roughly vertical, band running through the middle of the image is a “smearing” artifact caused by AHE at the treatment beam edge.

Standard unsharp masking produces an “edge-enhanced” image. The result is obtained by the addition of a blurred, or low-pass filtered, version of the original image with a detail, or high-pass filtered, version of the image. The detail image, gotten by subtracting the blurred image from the original, may be multiplied by a gain factor that dictates how much the detail in the image is enhanced (Equation 5.1).

$$L'(\vec{x}) = \tilde{L}(\vec{x}) + \gamma(L(\vec{x}) - \tilde{L}(\vec{x})), \quad \tilde{L} \text{ blurred image, } \gamma \text{ gain factor} \quad 5.1$$

⁵J.A. Sorensen, L.T. Niklason, J.A. Nelson, “Photographic Unsharp Masking in Chest Radiography,” Investigative Radiology 16 (1981): 281-288.

⁶S.M. Pizer, E.P. Amburn, J.D. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. ter Haar Romeny, J.B. Zimmerman, K. Zuiderveld, “Adaptive Histogram Equalization and Its Variations,” Computer Vision, Graphics, and Image Processing 39 (1987): 355-368.

⁷J. Rosenman, C.A. Roe, R. Cromartie, K.E. Muller, S.M. Pizer, “Portal Film Enhancement: Technique and Clinical Utility,” International Journal of Radiation Oncology Biology and Physics 25 (1993): 333-338.

SHAHE employs a blurring process, nonuniform diffusion, for use in unsharp masking that takes into account the structure in the image. To summarize the discussion from Section 4.3.1, the diffusion equation contains a conductance term that dictates the rate of diffusion. This conductance term may be made a function of local properties, such as the “steepness” of edges, so that the diffusion progresses in an anisotropic manner as influenced by the structure in the image.

By allowing edge strengths to dictate the diffusion, nonuniform diffusion may produce a blurred image in which small edge detail and noise are blurred away, but in which higher contrast edges are maintained or even sharpened further. When a blurring with this behavior is utilized in unsharp masking, small scale edges are enhanced relative to larger ones, and that result need not be enhanced with CLAHE to such extremes. That is, the CLAHE contrast may be limited more, and contextual region size may be enlarged. The result of the concatenation of these processes is an image in which several crucial aspects of object detection and visualization are improved. Not only is the overall contrast of objects improved by virtue of histogram equalization, but their structural details, edge characteristics and textures, are enhanced by the nonuniform unsharp masking. This combination of effects seems to be most beneficial to many interpretive tasks, and in particular those, such as the one related to the treatment beam, that involve mixtures of high and low contrast objects and large and small scale edges.

SHAHE, and many of the other enhancement algorithms that might potentially be applied to the portal film in an attempt to enhance the contrast, have a number of parameters that dictate the characteristics and extent of the enhancement. SHAHE has several parameters for the nonuniform diffusion, and two CLAHE parameters determine the contrast limitation and the contextual region size. Two from among all of SHAHE’s parameters were chosen for study in this research: the unsharp masking gain factor and CLAHE contrast limitation parameter.

The way that the unsharp masking gain is set turns out to be important. The reduction in the CLAHE shadowing enabled by nonuniform diffusion unsharp masking is important in determining the position of the treatment field edge. Furthermore, it is advantageous in interpreting the position of other structures in the portal image to enhance smaller scale edge detail. However, choosing a very large gain factor can cause the image to appear very noisy; after all, the noise in the original image is contained within the small scale portion of the image that is in turn multiplied by the gain factor. So an optimal choice of the gain parameter is of great interest in computing SHAHE for portal films (Figure 5.3).

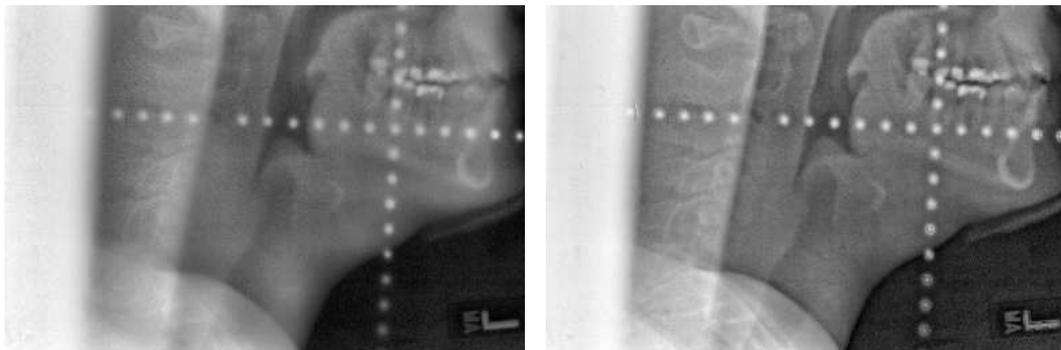


Figure 5.3. SHAHE gain parameter. The images are digitized portal films enhanced with different values for the gain parameter that multiplies the detail information in the SHAHE unsharp masking process.

The second SHAHE parameter manipulated in this research is the contrast-limitation factor for CLAHE. This parameter is a value that places an upper limit on the slope of the cumulative histogram to in effect control the “rate” or extent of the contrast enhancement. The parameter is implemented as a value that multiplies the average frequency in the region histogram to produce a value that in turn limits the frequency allowable at any particular bin in the histogram. Variation of this parameter has a big effect on the appearance of the image, with larger values causing the image to look more noisy and less like a standard radiograph. Yet higher contrast may enable better detection and interpretation of structures (Figure 5.4). This dilemma illustrates the necessity for a task-based image quality approach: only a performance study can reveal the true clinical benefit of a processing method that may not appear to be pleasing or useful.⁸

⁸J.B. Zimmerman, S.B. Cousins, K.M. Hartzell, M.E. Frisse, M.G. Kahn, “A Psychophysical Comparison of Two Methods for Adaptive Histogram Equalization,” *Journal of Digital Imaging* 2, no. 2 (1989): 82-91.

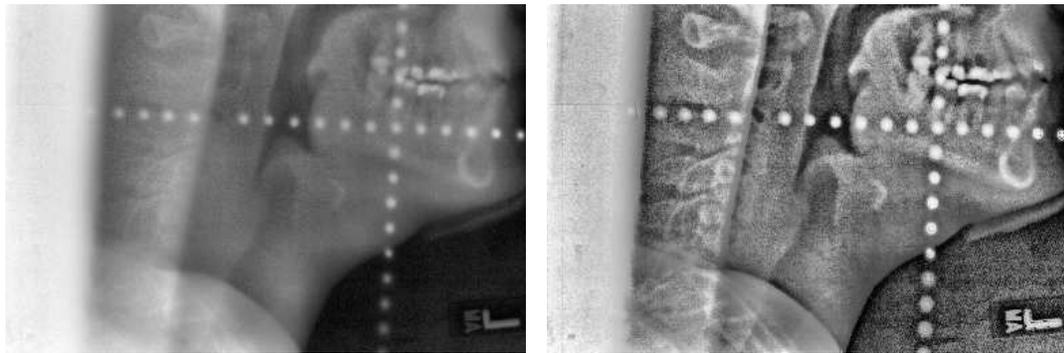


Figure 5.4. SHAHE contrast-limitation parameter. The images were enhanced with different values for the contrast-limitation parameter that controls the degree of contrast enhancement achieved by the CLAHE step in SHAHE.

5.2 Radiation Treatment Field Placement

In treating tumors in the lower head and neck region, the radiation oncologist will when possible construct the treatment field such that it avoids passage through the radiation-sensitive tissue of the spinal cord. The position of the spinal cord, soft-tissue anatomy that does not normally attenuate radiation sufficiently to appear distinctly in even a high-contrast diagnostic radiograph, is inferred from the position of the vertebrae that surround it. In particular, in the lateral view of the vertebrae, it is the posterior edges of the vertebral bodies that delineate the path of the anterior edge of the spinal cord. At the left in Figure 5.5, the vertebral body is the rectangular structure forming the right portion of the vertebra. The posterior edge of the vertebral body separates it from the spinous processes (left). The posterior edges are shown in a radiograph in Figure 5.6.

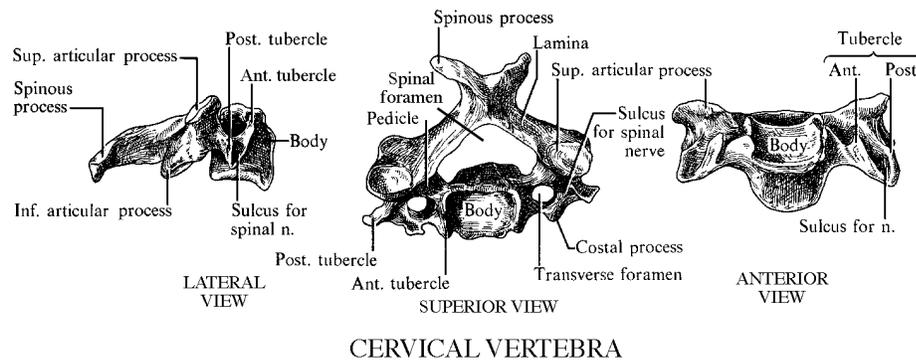


Figure 5.5. A cervical vertebra.⁹ There are seven cervical vertebrae, at the top of the spine, that form the neck. Note in the superior view the spinal foramen, which encases the spinal cord, and in the lateral view the square appearance of the vertebral body and its posterior edge.

An important judgment then in interpreting the portal film, and the one investigated in this research, is to determine the *clearance* of the treatment field, or assess just how close the edge of the field is to the posterior edges of the vertebral bodies. The portal film in Figure 5.1, which was produced for this purpose, faintly shows these edges, and the images in Figures 5.3 and 5.4 show the edges slightly better. Figure 5.7 diagrams the essential structures.

⁹taken from B. Pansky, E.L. House, Review of Gross Anatomy, 3rd ed. (New York: Macmillan Publishing Co., 1975), 147.

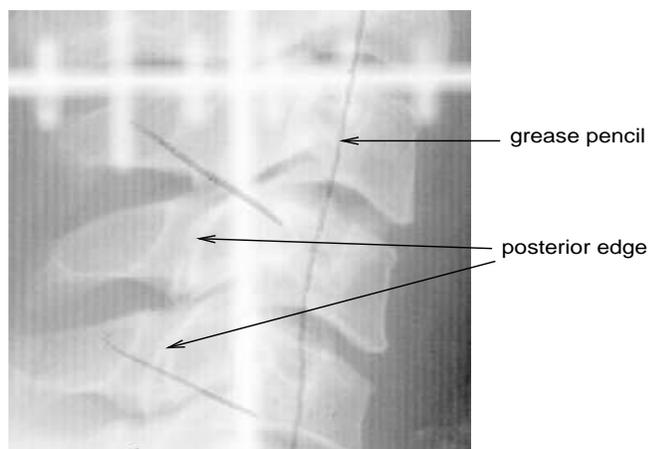


Figure 5.6. Radiograph of the cervical vertebrae. This image is cropped from a digitized radiograph of the cervical vertebrae acquired with a treatment simulator, an imaging device that participates in the planning of treatment field positioning and that produces diagnostic quality radiographs with additional grid lines. The film was digitized after the tentative field shape had been drawn with a grease pencil by the radiation oncologist. The posterior edges of two vertebral bodies are indicated.

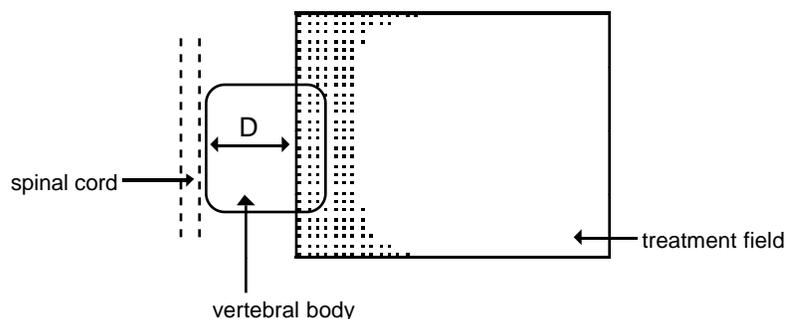


Figure 5.7. Treatment field positioning. The radiation oncologist's task is to verify that the treatment field lies sufficiently anterior to (toward the front of the body from) the spinal cord by measuring the clearance, distance D , from the treatment field edge to the posterior edge of the vertebral body. The position of the spinal cord, which is not imaged in a radiograph or portal film, is inferred from the edge of the vertebral body.

5.3 Portal Image Production

The images for the portal imaging experiments described in Chapter 6 were digitized portal films acquired from the UNC Department of Radiation Oncology. All of the films had been utilized in the treatment of tumors in the head and neck region and contained treatment beams whose posterior edge resided anterior to the posterior edge of the vertebral bodies. The films were positioned so that the anterior anatomy appeared on the right side of the digitized images. The digitized portal films were bilinearly interpolated such that the vertebral bodies contained in them were roughly 200 pixels long in the horizontal direction (5.9° visual angle at 60 cm). The images used in the experiments were 512x512 images cropped from the interpolated images to contain the beam edge and vertebral bodies. The images were chosen and cropped in such a way as to provide a prominent, or at least detectable, vertebral body edge at the vertical level in the image where the distance estimates would be made. SHAHE¹⁰ was applied to the images using the range of gain parameter and contrast limitation values specified in Chapter 6. Other parameter settings for the multiscale nonuniform diffusion in SHAHE included 100 iterations of the algorithm and a conductance value of 2.0 times the mean gradient magnitude of the image. A CLAHE region size of 64 was used. Figure 5.8 shows a typical image.

¹⁰D.T. Puff, "Ashahe," UNC Department of Radiology Image Processing Software, 1992.

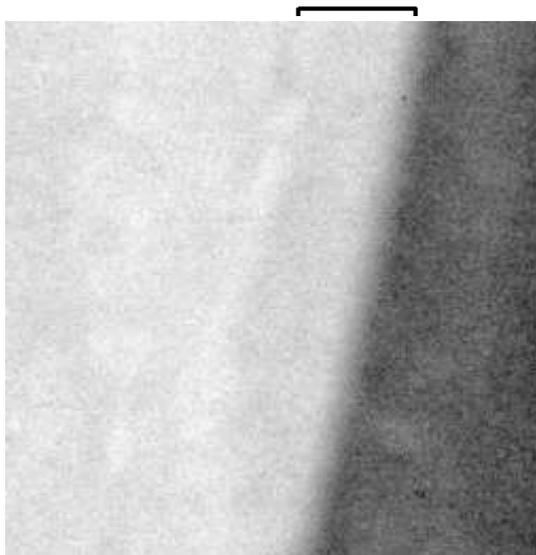


Figure 5.8. SHAHE-processed experiment image. This image is a subimage, focusing on the treatment beam and cervical vertebrae (bracketed above), of a digitized portal image processed with SHAHE. Images such as this one were used in the portal imaging experiment in Chapter 6.

5.4 Model Computation of Treatment Field Clearance

The way that the core model was used to produce an estimate of treatment field clearance, or the distance between the treatment beam and the posterior edge of the vertebral body, was to determine a core for the “gap-object” composed by the edges of those structures, and to use the width information from that core as the clearance estimate.

A gap-object is one that is formed from the edges of two or more other objects. Evidence for the existence of gap objects as visual entities comes from the work of Westheimer,¹¹ who showed that the ability to detect increments in spatial intervals was not significantly influenced by the polarity or the breadth of the sharp edges demarcating the interval. The discrimination thresholds for several combinations of bright and dark lines and edges, separated by 3 minutes of visual angle, were all similar. The objects and gap objects from that experiment were quite a bit smaller than those here, but nonetheless there doesn't seem to be any lack of visual primacy or preferential linking of edges that might not be considered to belong to the same object.

Medialness for the gap object was computed using the operator described in Section 3.3.3. As mentioned, the operator was designed for the specific situation present in the portal images in the experiments. The dark field interior was positioned at the right in each of the images, and the sharp, bright posterior edge of the vertebral body was always found to the left of the field edge. Two examples are now shown to help explain the medialness filter and resulting medialness distribution. First, Figure 5.9 diagrams the essential situation, a bright, thin edge and an adjacent broader edge. Figure 5.10 shows a local view of a typical enhanced portal image and an accompanying slice through its medialness scale space. In each case, a maximum of medialness in the region of the center of the gap object can be found and used to describe the middle and the width of the object.

As there is no known position for the center of the gap object formed by these two edges, a means of determining an initial guess for core calculation was needed. The position and scale for the guess were determined from the midpoint between the ridges computed for the two important edges.

First, the two edges used in the task were estimated as ridges of the magnitude of the first derivative with respect to the x direction, $|L_x|$, for the image in question. An initial guess, the author's judgment of the edge position, was supplied for initiation of ridge traversal per original, preprocessed, patient background. Those same two initial guesses were used for all of the images that resulted from different amounts of SHAHE processing on the original background. Initial guesses at horizontal neighbors successively farther away (up to 20 pixels in either direction of the initial guess) were attempted until an actual ridge was found. There were several cases where the initial guesses from the preprocessed image resulted in flow to edges that were not those that were to be used in the estimation task. In those cases, another guess that was more appropriate for the processed image was

¹¹G. Westheimer, S.P. McKee, “Spatial Configurations for Visual Hyperacuity,” *Vision Research* 17 (1977): 941-947.

provided by the author, and that resulting ridge was used. Using this procedure, there was never a failure to find ridges of the derivative for any of the images generated for the experiment.

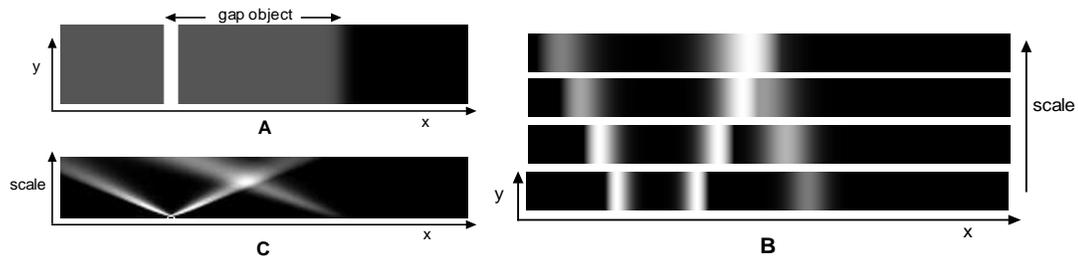


Figure 5.9. Medialness for simulated portal image. Figure A contains the essential structures from the portal image in Figure 5.10. Figure B shows the medialness distribution at four, increasingly larger, scales. Medialness at a position corresponds to the existence of properly-oriented edges at a distance away that is proportional to the scale. Medialness is maximal in the middle of the third scale slice, where the scale of the medialness kernel is such that it engages both edges. Figure C is a slice through scale along the x dimension and depicts the maximal medialness at a scale where the responses from the two boundaries cohere.

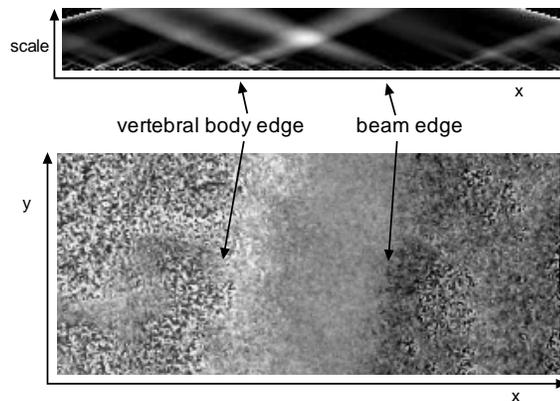


Figure 5.10. Medialness for portal image. The top image shows a slice along the x dimension through the medialness scale space computed for the portal image below. The position and scale of the maximal medialness correspond respectively to the position and width of the gap object formed by the two important edges.

The derivative measurements were made with a scale that was expected to be similar to that used in core construction. That is, the relationship between boundariness scale, σ , and object half-width, r , was chosen for the core computations (described next) to be $\sigma = \frac{1}{8}r$. Thus a rough estimate of the half-width of the gap object, computed from the user-supplied initial guesses for the derivative ridges, was entered into that equation to determine a scale for the derivative measurements.

Along each of the two resulting ridges, a maximum derivative within a 20-pixel vertical distance of the initial guess was found. The midpoint and distance between those two maxima were used as the position and radius, respectively, of the initial guess for core-finding. Core flow and traversal were initiated from that position via the methods described in Section 3.3. The beam-to-vertebral body distance was then taken as the mean width of the core generated from that initial guess (Figure 5.11). For cores that extended beyond the central region of the image in which the estimate was to be made, core width measurements contributing to the mean were only those from the core subsegment within a 20-pixel vertical distance of the midpoint initial guess.

Core formation according to this initiation procedure occurred with good success (Figure 5.12). However, for 28 percent of the images generated for the experiment, the distance estimate was computed via an alternative but related method. Those cases were mainly a result of core formation failures at the initial guess. Alternatively, occasionally cores were determined via the automated procedure that corresponded to objects, or more frequently gap objects, that were clearly not made up of the edges which were intended to be used in making the distance judgment. The cores representing the intended objects might well have been obtained with a different initial guess. However, for both core initiation failures and cores for the wrong objects, a gap object width estimate was taken to be the object radius proportional to the scale at which maximal medialness at the single

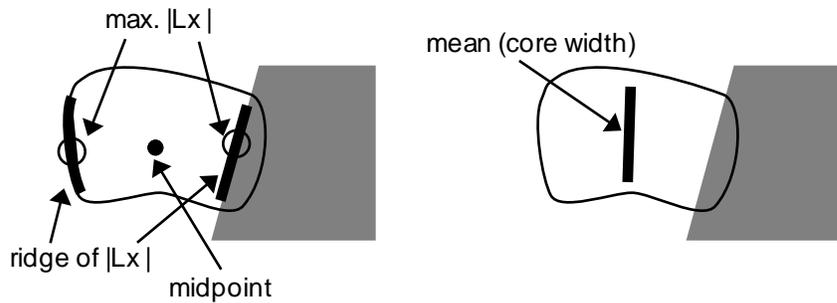


Figure 5.11. Initial guess for core estimation of treatment field distance. Ridges of the derivative with respect to the x direction were computed for the vertebral body and treatment field edges. The midpoint between the two maxima along the ridges was used as the guess to initiate core finding. The mean width of the resulting gap-object core was used as the estimate of the treatment field distance.

guess position occurred. Specifically, the medialness kernel used for this application was convolved with the image at the guess position for a range of scales. At that scale at which the output of the operator was locally maximal, the corresponding object radius (a factor of 8 times the scale) was used as the half-width estimate. As the same medialness operator was used for this procedure as was used in the core analysis software, the resulting estimate was expected to be very similar to that produced by the core. In fact, when estimates obtained via this maximum-through-scale method were attempted with images which did have successful core estimates, the agreement was excellent. For a random sample of ten cases that were tried, the mean difference between the estimate from the scale maximum at a single position and the estimate derived from the core was -1.02 pixels for a mean core-estimated distance of 112.42 pixels.

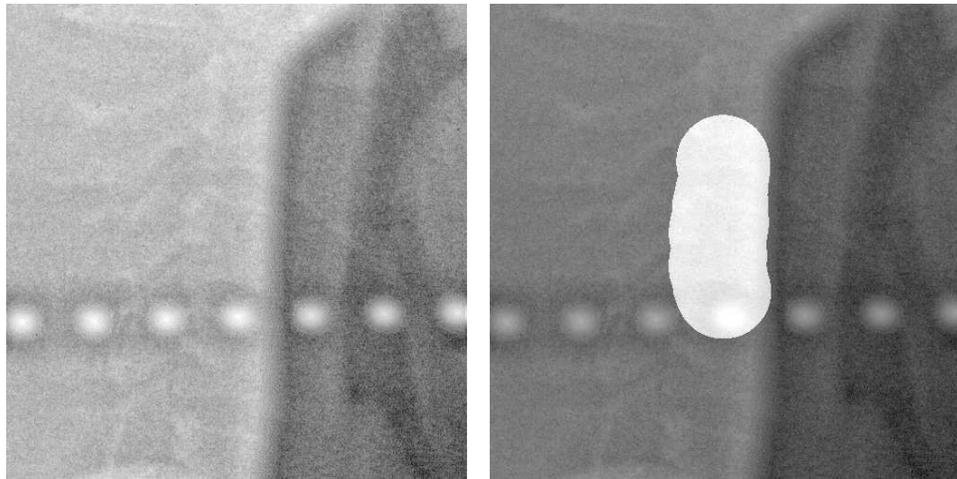


Figure 5.12. A gap object core. The core, drawn as discs of appropriate radius about the core middle, is shown superimposed in the image on the right.

The seemingly ad-hoc methodology for choosing a position at which to initiate core finding and traversal was an attempt at simulating what the human observer might do in his/her recognition and perception of the gap object. The observer would presumably localize the two important edges, and in a vertical region adjacent to the measuring tool designate a prominent position on each edge to use in the judgment. This approach also avoids the potential biases in having a user specify the guess for core finding.

5.5 Summary

One of the attractive properties of a computed measure of image quality is its ability to relatively painlessly compute quality assessments for the many parameter settings of a system or processing method. The potential combinations of those parameters constitute a parameter space. The optimal acquisition of an image or application of a processing algorithm for a given task results from the use of some parameter combination from within the space. As the parameter space for multiparameter methods may be enormous, a computed image quality measure might be used to determine the quality of the images throughout the parameter space and thus

eliminate the need for the many observer trials required in providing that assessment. Cromartie¹² has demonstrated the feasibility of an image quality model similar to the one advocated here for computing performance estimates in the massive SHAHE parameter space.

SHAHE is a complex, multistep process dictated by several parameters. As the examples show, altering the parameters does have an effect on the appearance of the image. The images clearly look different; the question is, which is “better.” With so many parameters contributing to the resulting image, it is difficult to characterize the effect of the variation of a single parameter. The result of jointly varying multiple parameters is often even less understandable or predictable. The only objective means for determining how to choose SHAHE parameter settings in a way that optimizes performance with the processed images is to actually measure performance. Thus the SHAHE optimization problem illustrates a fundamental difficulty for image quality research: objective performance assessments must be made at a multitude of parameter settings. An approach such as is proposed in this research might allow the selection from within the parameter space of those setting that were best for human use.

The images from the portal experiment were not simulated but were real digitized portal films processed with SHAHE. The variation in the physical characteristics of the images was achieved by manipulation of SHAHE’s parameters. In this case it was not possible to know the true distances between the treatment beams and vertebral body edges. Several months of research were invested exploring the possibilities for portal image simulation. It is possible to use treatment planning software from the Department of Radiation Oncology to design treatment beams and produce simulated portal radiographs containing those beams. This would have allowed fairly accurate determination of both edges for “truth” purposes. There were however a number of artifacts that arose in producing the simulated radiographs. Their appearance was not entirely like that of real portal images. Secondly, some of the positioning materials used to place patients consistently in the treatment equipment would appear in the simulated radiographs. These materials, which do not appear in real portal films, created veiling, confusing edges in the simulated images. For these reasons, at the expense of knowledge of the truth, real digitized portal films were utilized in the experiments.

As it turns out, it was of interest to designate a rough estimate of truth in order to perform the analyses in Chapter 7. First, the appropriate measure of performance for the distance estimation task was a relative error. It is known that the accuracy for a visual distance judgment is approximately inversely proportional to the distance itself. In order to make estimation differences between the model and human commensurable at different distances, the absolute difference between the model and human distance estimates was normalized by the “true” distance that was judged. Second, it was desirable to measure human and model accuracy as a function of the SHAHE parameters, and those individual accuracy outcomes necessarily required a value for truth. These accuracy measurements made it possible to further understand separately the human and model behavior that participated in the ultimate comparison of the two. In both cases, truth was provided via a combination of the author’s judgment and image analysis methods (see Section 7.4).

An objective determination of clinical truth is a contentious undertaking for any image quality methodology.¹³ The reliability of the source of the truth necessarily has an impact of the validity of any experimental results. Furthermore, designations of truth are prone to unforeseen biases that accompany any human or even computed judgment. Nevertheless, there are accepted means for arriving at estimates of truth: the decision of a panel of experts or data from pathology inquiries are common examples. So the attention given to the knowledge or determination of truth in this research is not uncommon. Diagnostic or clinical performance with respect to the true physical or physiological status is simply how medical image quality is quantified.

The final part of this dissertation describes the design and analysis of the experiments that were conducted, measuring performance for the tasks just described, to compare the core model’s results with human performance.

¹²R. Cromartie, “Structure-Sensitive Contrast Enhancement: Development and Evaluation,” Ph.D. dissertation, (University of North Carolina-Chapel Hill, 1995).

¹³C.E. Metz, “ROC Methodology in Radiologic Imaging,” *Investigative Radiology* 21 (1986): 720-733.