Part II.
BASES FOR EVALUATION of the MODEL

There are presently a handful of medical imaging modalities, each of which displays as intensity values on film or a computer screen the differential response characteristics of anatomical structures to the imposed imaging technology. For example, digital radiographs and computed tomography scans image the differential attenuation of x-rays by the structures of the body. Magnetic resonance images depict the different nuclear resonance ("spinning") properties of structures initially placed within a magnet. Parameters dictating the properties of components of each of these systems, such as the initial energy source and collection equipment, or the image processing and display methods, may have an appreciable influence on the observer's proficiency in making clinical detection and interpretation decisions. Optimizing quality thus requires adjusting those parameters that dictate the operation of the system. The next two chapters contain a brief description of the two imaging modalities chosen for the image quality investigations in this research. More importantly, acquisition or processing parameters that might be expected to have an effect on the performance of relevant tasks, and thus on quality of the images, are described.

As emphasized in Chapter 2, inherent in the evaluation of image quality is the determination of the relevant task(s) performed with the images. It may be the case that there are several or many fundamental purposes for which an image is consulted, and optimal acquisition and display might be different for each of those tasks. A mammogram is acquired for the detection of, among other things, breast masses, calcifications, and spiculations,[1] and the optimal digital image processing for improving detection might be different for each of those structures.[2] Two estimation tasks performed with the imaging modalities in this research have been selected. In both cases, the discussion will demonstrate that the tasks are of critical clinical importance, such that one would wish to adjust the operation of the imaging systems to optimize the performance of the tasks. The tasks furthermore require fundamental shape-related visual judgments suitable for computation with the core model. This chapter contains descriptions of 1) the stenosis estimation task frequently performed with angiograms and 2) the treatment beam distance determination undertaken with portal films.

An important component of this research is the collection and generation of images for use in the psychophysical experiments described in Chapters 6. It was necessary, for both the angiography and portal imaging experiments, to create a set of images for which the physical characteristics or processing parameters were systematically varied: the variation in the "appearance" of the images serves as the dependent variable in the assessment of quality. The methods used in generating the experimental images for both investigations are described in Chapters 4 and 5 as well.

Finally, these two chapters describe particular implementations of the core model. For both tasks, the model's estimation procedure consisted of core analysis followed by the extraction of pertinent information from the resulting core. While the mechanisms of core formation seem highly plausible in light of the physiological and psychophysical evidence outlined in the previous chapter, it is difficult to say whether the proposed means of combining and using core information to perform these tasks has any basis in visual interpretation. Nonetheless, these adaptations are arguably fair approximations to fundamental visual operations. The imaging tasks that are detailed here provide a realistic test of the predictability and applicability of this core-based assessment.

4. ANGIOGRAPHY and STENOSIS ESTIMATION

Model and human image quality measurements were obtained in this research for images that simulated those that might be produced from a digital angiography system. Angiography is an imaging modality employed for a host of interpretive tasks related to blood vessels, including assessments about vessel constrictions, or stenoses. In this chapter the acquisition of angiograms and the stenosis estimation task are briefly described in Sections 4.1 and 4.2, respectively. The simulation of angiographic images for use in the experiments in this

---

[1] D.B. Kopans, Breast Imaging (Philadelphia: J.B. Lippincott, 1989).

[2] E.D. Pisano, personal communication.

research is detailed in Section 4.3. Finally, the manner in which the core model was applied to compute the stenosis depth estimates is related in Section 4.4.

4.1 Angiography

Radiographic imaging systems aim a narrow beam of x-ray photons at the anatomy to be imaged and capture whatever photons have passed completely through the body with a screen/film system or electronic collection equipment. This process relies on the differential absorption of x-ray photons by the various structures in the body to produce a picture that is a mapping of the beam attenuation. Many structures in the body, such as blood vessels, absorb x-rays very little and fail to appear in the standard radiograph. Angiography is an x-ray imaging modality in which a radiodense contrast agent is injected into the pertinent regions of the patient's vasculature to improve the radiographic contrast there. For this reason, angiography is consulted for many diagnostic decisions about vessels throughout the body. The image in Figure 4.1 is a digitized angiogram of the carotid arteries.

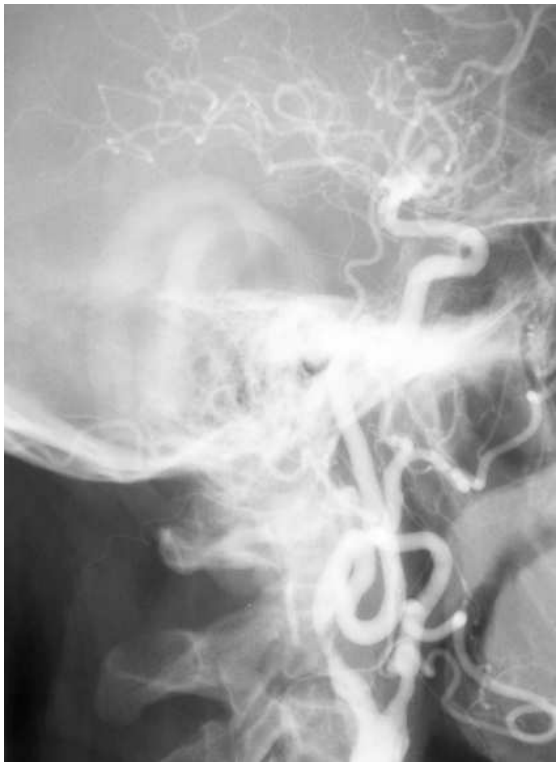

Figure 4.1. Angiogram of the carotid arteries. This digitized angiogram shows the common carotid artery (proximal to the bifurcation), and the internal and external carotid arteries (left and right halves of the bifurcation, respectively).

There are several parameters in a digital angiography acquisition system that have an effect on the quality of the angiogram.† Among those parameters, two related to the observer's ability to perform quantitative tasks regarding the health and integrity of the vessels are the imaging system blur and noise.[3] It is the effect of these two parameters on the model's performance of the stenosis estimation task described in the next section that will serve as the test of the model.

The imaging system blur, a function primarily of the x-ray beam focal spot geometry, results in unsharpness vessel borders (Figure 4.2). The focal spot, the origin of the x-rays, is not a point source but has a finite size that is a function of the x-ray tube geometry (size) as well as the tube operating conditions such as

† While it is the case that a user would probably do little in the way of altering a commercially-available angiography system with respect to physical properties such as noise and blur, at some point during the production of the equipment design decisions must be made that impact the quality of the images that the proposed system will produce. It is at that production stage that image quality assessment is often most important for the adjustment of system parameters.

[3] A.J. Buda, E.J. Delp, Digital Cardiac Imaging (Boston: Martinus Nijhoff, 1985).

voltage and current.[4] The geometry of the finite focal spot causes the edges of the imaged objects to possess a "penumbra," an intensity variation, or fall-off, at the edge of the object that appears as unsharpness or "blur" (Figure 4.3). The penumbra may also be influenced by the positioning of the anatomy: moving the patient as far from the focal spot and as close to the collection equipment as possible minimizes its extent.[5]

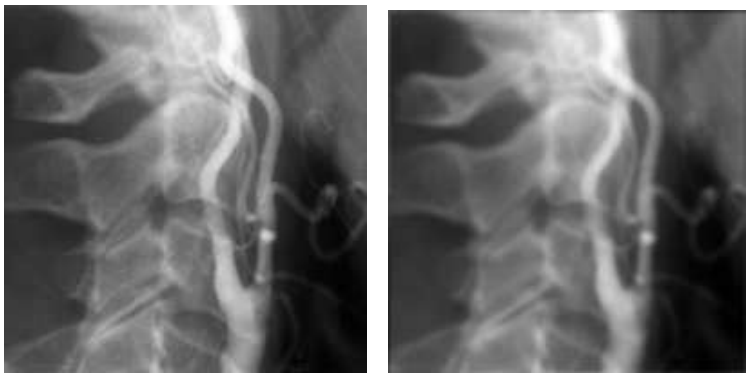

Figure 4.2. Angiography system blur. The angiogram image on the right is a simulation of the unsharpness that might occur in an angiography system. The image is produced by blurring the original image on the left.

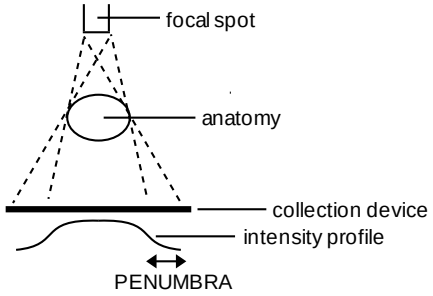

Figure 4.3. Unsharpness caused by focal spot penumbra. The finite size of the focal spot causes the intensity fall-off, or penumbra, at the edges of objects.

A second physical parameter of the angiographic imaging system is noise. Quantum noise, that appears as a fine mottled, "speckled" texture in the image, results from the quantal, or discrete, nature of x-radiation. Random statistical fluctuations in the number of quanta arriving at the x-ray collection device from the source cause quantum mottle, or noise (Figure 4.4). The exposure duration, which controls how many photons are emitted by the x-ray tube and is related to how many photons are ultimately captured by the collection device, thus determines the amount of noise in the image. While longer exposures produce less noise in the image, the exposure is of course in practice limited by considerations of patient safety. Cinematic angiographic sequences, which contain a series of images capturing the flow of the contrast agent through the vessels, are particularly prone to noise. To produce a real-time sequence and at the same time minimize overall radiation exposure, each cine frame is acquired with a brief, low-energy beam pulse. This causes relatively few photons to be detectable per angiographic image.

To a lesser extent, *electronic* noise, statistical variability in the distribution of the electrons constituting an electrical signal, may be present in a digital angiography system. The noise that results from the combination of both quantal and electronic sources causes an uneven intensity distribution within and across anatomical

[4]E.L. Chaney, W.R. Hendee, "Effects of X-Ray Tube Current and Voltage on Effective Focal-Spot Size," Medical Physics 1, no. 3 (1974): 141-147.

[5]T.S. Curry, J.E. Dowdey, R.C. Murry, Christensen's Introduction to the Physics of Diagnostic Radiology (Philadelphia: Lea & Febiger, 1984).

structures. Vessel detection and quantification tasks related to finding vessel abnormalities or determining the position of the vessel edge become perceptually difficult with increased noise levels in the angiogram.[6]

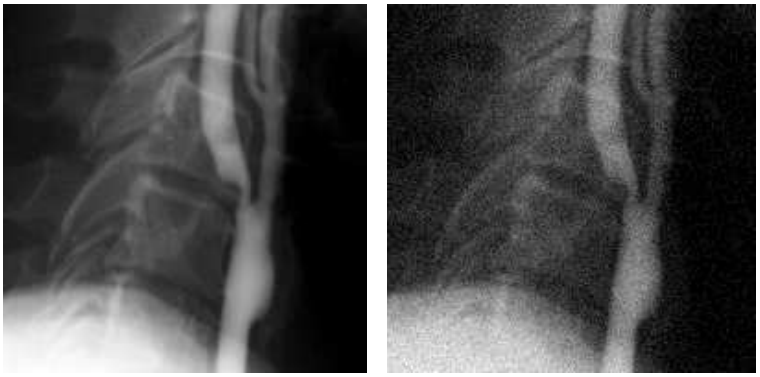

Figure 4.4. Angiography system noise. The simulated noise in the angiogram image on the right results from statistical fluctuations in x-ray photons arriving at the receptor and electronic noise. The image was simulated by adding Poisson noise to the original image on the left.

There exists an inherent blur and noise trade-off in a real angiography system. Lower noise levels are usually accomplished by using a larger x-ray source that produces more photons. However, when the source size is increased, more blur accompanies the improved noise characteristics. This research does not attempt to simulate this trade-off, but instead measures stenosis estimation performance as a function of the independent variation in both noise and blur. A subset of these parameter combinations could conceivably be thought of as defining the one-dimensional blur-noise relationship that would result from changing a single "source size" parameter.

4.2 Stenosis Estimation

A stenosis, or vessel wall plaque that causes localized restriction of blood flow, may induce sudden neurological dysfunction, known as a stroke, that results from the concomitant ischemia. The detection and proper treatment of a stenosis of the internal carotid artery, the chief source of blood flow to the brain (Figure 4.5), is crucial in preserving normal neurological operation. Using angiography, neuroradiologists often estimate the depth of a carotid artery stenosis in choosing among the treatments available (surgery, endarterectomy, or medication, for example) in remedying the critical ischemia.[7] Because it is an important clinical task and one for which the angiography system should be optimized, this stenosis measurement has been chosen as the basis for the image quality investigations in this research.

What is important in the assessment of a vessel stenosis is the *relative* severity of the constriction. Stenoses may be measured in vessels in disparate regions of the body and in different patients with varying vascular maturity or integrity, so any sort of absolute measurement of constriction width is meaningless without an accompanying measure of width along the normal, unconstricted portion of the vessel nearby. Thus, the depth of stenosis may be calculated by forming a ratio of the estimate of the vessel diameter at the point of greatest constriction, S, to the diameter, N, beyond (above or below) the visible constriction at a normal region of the vessel (Figure 4.6). That fraction is subtracted from one to obtain a measure of the relative extent to which the vessel is constricted (Equation 4.1).[8] So a measurement might indicate that a vessel is 50 percent stenosed, for example.

$$\% \text{ stenosis } = 1 - \frac{S}{N} \qquad 4.1$$

[6]K. Ohara, K. Doi, C.E. Metz, M.L. Giger, "Investigation of Basic Imaging Properties in Digital Radiography. 13. Effect of Simple Structured Noise on the Detectability of Simulated Stenotic Lesions," Medical Physics 16, no. 1 (1989): 14-21.

[7]N. Rosenberg, CRC Handbook of Carotid Artery Surgery: Facts and Figures (Boca Raton, FL: CRC Press, 1989).

[8]NASCET Steering Committee, "North American Symptomatic Carotid Endarterectomy Trial: Methods, Patient Characteristics, and Progress," Stroke 22, no. 6 (1991): 711-720.

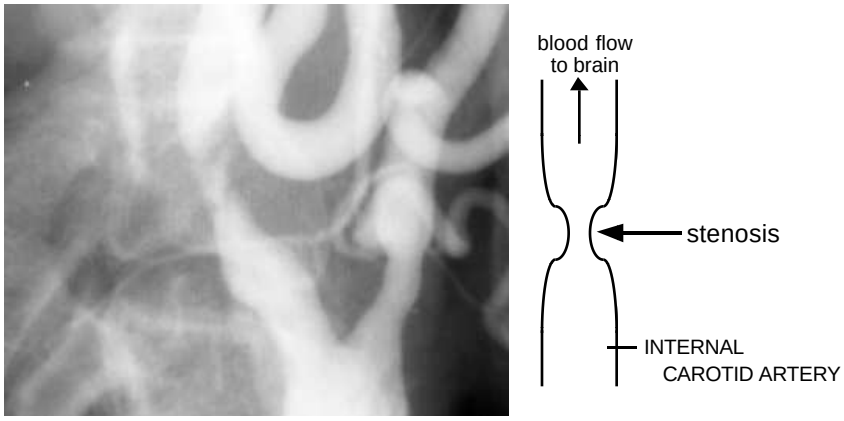

Figure 4.5. Stenosis of the internal carotid artery. The diagram depicts the apparent constriction of the vessel as seen in the angiogram on the left. The stenosis is located in the left branch of the vessel, roughly in the middle of the image.

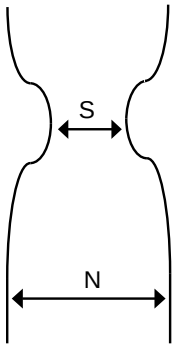

Figure 4.6. Protocol for stenosis estimation. The radiologist expresses the depth of the stenosis as 1 minus the ratio of the width of the vessel at its most constricted point, S, to the width of the normal vessel, N.

4.3 Angiographic Image Simulations

It was necessary to generate a set of digital angiographic images for the experiments in this research in which blur and noise were varied. As it would have been virtually impossible to alter the physical operation of the clinical angiography equipment at this institution, image generation and processing techniques were utilized to simulate digital angiograms with chosen amounts of noise and blur.

If blur and noise are to be added to an image to simulate an angiogram, they must be "superimposed" on anatomy. That anatomy consists of one or more blood vessels filled with contrast agent along with the other surrounding background anatomy. Rather than attempt to develop a model for simulating human anatomy, the following approach was adopted. Digitized angiographic "scout" films, taken during the angiography procedure prior to the appearance of the contrast in the vessels, were manipulated by a process called multiscale nonuniform diffusion to obtain an estimate about the original background radiation distribution. Contrast-filled internal carotid vessels were then simulated and added into this background. Finally, different amounts of Gaussian blur and Poisson noise were added to the composite image, as if it had been captured by the digital collection

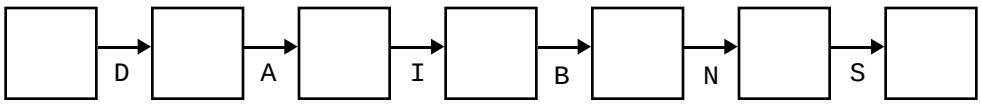

Figure 4.7. Angiographic image simulation process. A digitized scout angiogram was processed with multiscale nonuniform diffusion (step D). A simulated vessel was added (step A) into this image. A greyscale inversion and scaling (step I) was then applied. This image was subjected to selected amounts of Gaussian blur (step B) and Poisson noise (step N). Finally, a reinversion and rescaling (step S) were performed.

equipment of the imaging system. The steps involved in this simulation are diagrammed in Figure 4.7, and are described in turn in the following sections.

4.3.1 Multiscale Nonuniform Diffusion

A multiscale non-uniform diffusion algorithm was used to process the scout angiograms. As these angiograms already contained their own noise and blur, it was necessary to attempt to remove those degradations prior to the artificial imposition of selected amounts of blur and noise that was performed to produce the series of images for the experiment. The diffusion process produces a good guess about the underlying intensity distribution prior to the imaging system's addition of blur and noise. By smoothing away small scale noise, and by sharpening the larger scale edges that were blurred by the focal geometry of the system, this diffusion produces, in a multiscale, principled fashion, a result that is representative of the population of possible true solutions.

Nonuniform diffusion is an image processing method for blurring, or smoothing, an image (L, below) that possesses the property that the diffusion may vary across space. Thus the conductance term, c, in the diffusion equation is established to be a function of spatial properties of the image (Equation 4.2).†

$$\frac{\partial \mathrm{L}}{\partial \mathrm{t}} = \nabla \bullet \mathrm{c}(x, y, t)\nabla \mathrm{L} \qquad 4.2$$

In particular, several authors[9,10] have proposed that the conductance be a monotonically decreasing function, g, of the local gradient magnitude, $\|\nabla \mathrm{L}\|$, or edge "steepness." When the conductance function is chosen properly according to constraints discussed in the references, boundaries of large magnitude may be preserved (or even enhanced), while smaller scale boundaries, that may be of little interest or the result of noise, are blurred away (Equation 4.3).

$$\frac{\partial \mathrm{L}}{\partial \mathrm{t}} = \nabla \bullet \mathrm{c}\nabla \mathrm{L}, \quad \mathrm{c} = \mathrm{g}\left(\|\nabla \mathrm{L}\|\right) \qquad 4.3$$

The gradient measurements calculated in determining the conductance in the previous equations are made at some scale: the scale dictates the local region over which information is integrated in making that measurement. Measurements made at a pixel-level scale are likely inherently unreliable because of the inevitable presence of noise at that scale. An approach to nonuniform diffusion, multiscale nonuniform diffusion,[11,12] suggests that initial, "tentative," gradient measurements be made at some larger scale. Then, under the assumption that the measurements become more reliable as the diffusion process progresses, successively smaller scales may be utilized (Equation 4.4).

$$\frac{\partial \mathrm{L}}{\partial \mathrm{t}} = \nabla \bullet \mathrm{g}\left(\left\|\nabla \mathrm{G}\left(\mathrm{s}(t)\right) * \mathrm{L}(x, y, t)\right\|\right)\nabla \mathrm{L} \qquad 4.4$$

Thus the conductance is dictated by gradient measurements made over a Gaussian-weighted region (G) with scale s(t). An exponentially decreasing shape for the s(t) function may be theoretically advisable.

The relationship between the number of photons that reach an angiographic film and the intensities from the digitized film reflects two transformations that should be considered prior to this diffusion step. The film has a characteristic, "H&D" curve that specifies the relationship between beam exposure and the film's resulting

---

† The ∇ symbol indicates the gradient operator, $\nabla \mathrm{L} = \left\langle \frac{\partial \mathrm{L}}{\partial x}, \frac{\partial \mathrm{L}}{\partial y} \right\rangle$. The $\| \ \|$ symbolizes a vector magnitude, while • is the dot product of vectors.

[9] S. Grossberg, "Neural Dynamics of Brightness Perception: Features, Boundaries, Diffusion, and Resonance," Perception and Psychophysics 36, no. 5 (1984): 428-456.

[10] P. Perona, J. Malik, "Scale-space and Edge Detection Using Anisotropic Diffusion," IEEE Transactions on Pattern Analysis and Machine Intelligence 12 (1990); 429-439.

[11] R.T. Whitaker, "Geometry-limited Diffusion," Ph.D. dissertation, (University of North Carolina-Chapel Hill, 1993).

[12] R.T. Whitaker, S.M. Pizer, "A Multi-scale Approach to Nonuniform Diffusion," Computer Vision, Graphics, and Image Processing: Image Understanding 57, no. 1 (1993): 99-110.

photometric density.[13] Secondly, the film digitizer used to scan the angiograms possesses a transfer function dictating the relationship between film density and digital intensity values. If the original photon distribution was to be ultimately estimated, these transformations should technically have been inverted prior to the diffusion step. However, the digitizer[14] was known to be roughly linear with respect to input optical densities.[15] Second, technical specifications were acquired for the film[16] used in the UNC Department of Cardiology's angiography equipment for the period to which the films used for this dissertation corresponded. It was verified that, for the head and neck regions of several scout films used in the study, optical densities fell along the linear portion of the characteristic curve. Since the relationships were shown to be roughly linear, no attempt was made to invert either the digitizer or film transformation. The result is that the vessel, and later the noise, could be linearly added to the intensities in the digitized, processed image.

To produce the background images, angiographic film sequences of the carotid vessels were obtained from the UNC Hospitals Cardiology division, and a scout view (taken prior to contrast injection) and a representative standard (contrast-filled) view were digitized at 80 micron/pixel resolution. All films were digitized such that anterior anatomy appeared on the right side of images. The scout images were bilinearly interpolated to a size for which the cervical vertebral bodies were roughly 150 pixels (4.4^{o} visual angle at 60 cm) wide.

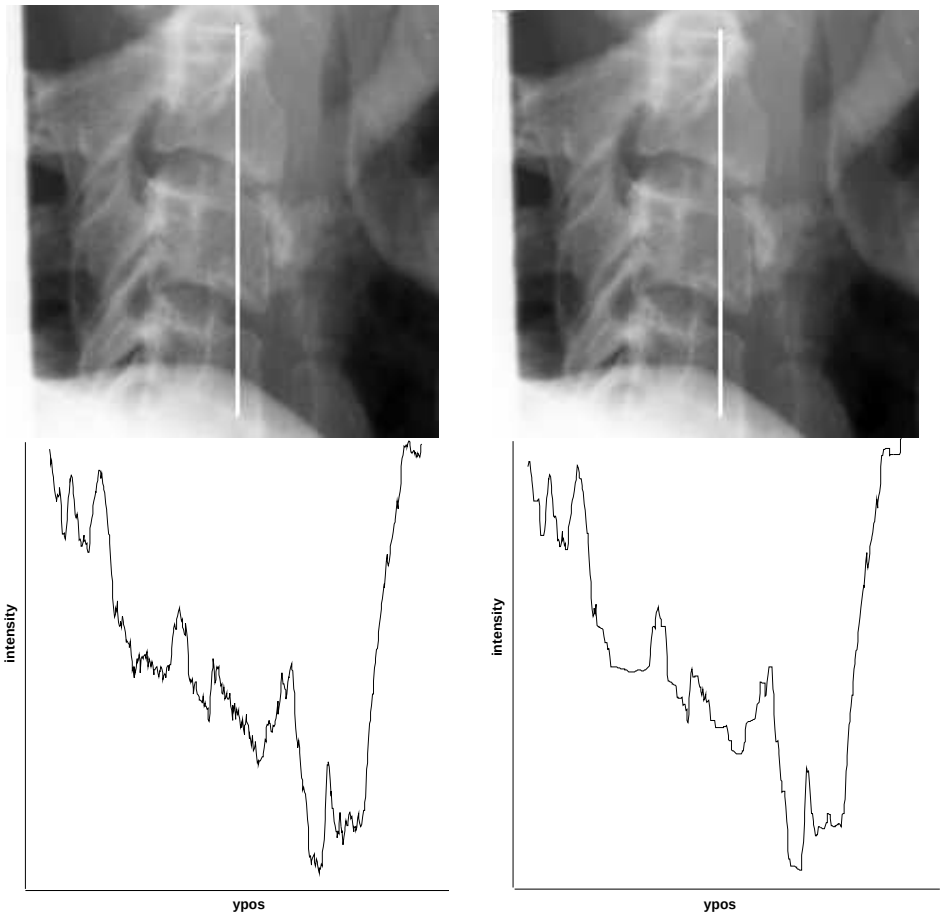

Figure 4.8. Intensity profile for multiscale nonuniform diffusion. The graphs are intensity profiles corresponding to the vertical lines marked in the original (left) and diffused (right) images. The profile of the diffused image demonstrates noise removal and some edge sharpening.

[13] T.S. Curry, J.E. Dowdey, R.C. Murry, Christensen's Introduction to the Physics of Diagnostic Radiology

[14] Lumisys, Inc., Sunnyvale, CA.

[15] A. Dillon, personal communication.

[16] Fuji Super HR-G, Fuji Photo Film Co., Ltd., Tokyo, Japan.

The background images used in the experiments were cropped from the interpolated scout images from the region, containing several of the cervical vertebrae, that corresponded to the area underlying the vessel in the matching image with vessel contrast. The resulting images were 512x512 pixels (13 degrees of visual angle at 60 cm), and were all linearly intensity-mapped to have a range of 0-4095. Multiscale nonuniform diffusion[17] was applied to the images for 100 iterations with an exponential decrease in scale from 4.0 to 0.0. The conductance parameter was set at 0.3 times the mean gradient magnitude of the image. The software defaults were used for all other parameters Empirical analysis (Figure 4.8) suggested these settings were best for reducing the small-scale noise and sharpening edges in these images.

4.3.2 Vessel Simulation

Realistic blood vessels were simulated with software developed for that purpose.[18] The software allows the specification of a random but constrained three-dimensional vessel path. The path was in this case constrained such that it traversed roughly vertically upward in the image with a local curvature and overall orientation similar to that of the internal carotid artery. A tube, centered about the vessel path and with a specified diameter, was finally parallel-projected along the depth axis to generate a two-dimensional image. The vessels in the resulting image, owing to projection of the vessel volume, possess dense lumina and a smooth intensity falloff at the edges that cause them to appear strikingly like the vessels in real angiograms (Figure 4.9).

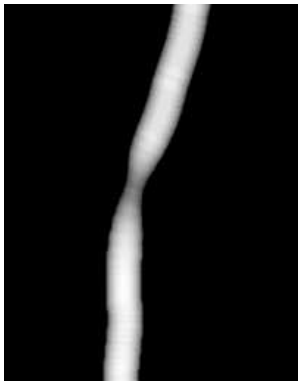

Figure 4.9. Simulated angiographic vessel. A random but constrained vessel volume much like the internal carotid artery is parallel-projected to create this realistic vessel simulation.

A stenosis was positioned along the vessel path. The fall-off of the stenotic width function along the vessel axis was described by a Gaussian profile. The Gaussian has two parameters. The amplitude of the Gaussian, which dictates the severity of the stenosis, was an independent variable in the experiments (Section 6.1). The standard deviation of the Gaussian determining the axial extent of the stenosis was fixed at 30 pixels (0.9^{o} visual angle at 60 cm).[†,19] From an examination of the original angiograms, a vessel diameter (24 pixels, or 0.7^{o}) was chosen to be consistent with the scale of the background anatomy upon which it was superimposed. The greyscale range of the initial vessel image was 0-1023.

The vessel image was added, via simple pixelwise addition, to the result of the diffused scout angiogram to create what might be considered a mapping of the projected attenuation characteristics of the original anatomy. This image possesses large values in attenuating anatomy such as the contrast-injected vessel and bones. Finally, the greyscale mapping for this image was inverted to generate a mapping of the radiation distribution exiting the patient. This distribution is inversely related to the attenuation properties of the anatomy through which the radiation beam passes: fewer photons pass through the radiodense vessel and bones. It is this radiation distribution that is subsequently influenced by the resolution and noise properties of the system.

---

[17] K. Gadepalli, "Pdiffuse," UNC Department of Computer Science Image Processing Software, 1992.

[18] J.P. Rolland, D.T. Puff, "Angiogram Simulation Software Documentation," (University of North Carolina Department of Computer Science Technical Report, TR93-018, 1993).

[†] The axial length of the stenosis may affect blood flow as well. However, it is rarely used as an indicator of disease severity.

[19] R.L. Feldman, W.W. Nichols, C.J. Pepine, C.R. Conti, "Hemodynamic Significance of the Length of a Coronary Arterial Narrowing," American Journal of Cardiology 41 (1978): 865-871.

4.3.3 Blur Simulation

The imaging system's resolution, a property defined by the smallest structure that may be recorded by the system, is negatively influenced, in a radiographic imaging system, by the blur caused by the focal spot geometry discussed in Section 4.1.1. A means of quantifying this characteristic is the point spread function, the intensity distribution recorded by the imaging system for a single point source of light. The point spread function for an imaging system typically imprecisely resembles a Gaussian distribution (Figure 4.10).[20,21]

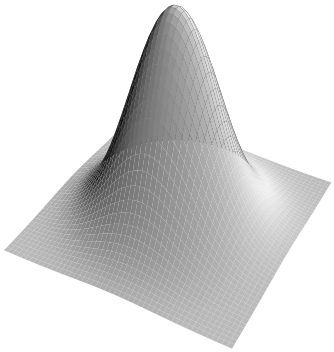

Figure 4.10. Imaging system point spread function. The response of an imaging system to a point source of light often resembles a Gaussian distribution.

For the purposes of the image simulations in this research, the point spread function of the imaging system was modeled as a Gaussian, and its effect upon the photon distribution was simulated by convolution of the photon distribution image with a Gaussian kernel. The standard deviation of the Gaussian was a parameter that was varied in the experiments (Section 6.1). This blur in the image occurs prior to the capture of the photons by the collection device; thus the next step in the simulation process is to superimpose noise.

The standard deviation of the Gaussian that represents the acquisition system's blur function is not sufficient to fully describe the amount of blurring that an object undergoes in the human perception of it. The measurements made by the human visual system impose a perceptual blurring that effectively further blurs the object in question. Furthermore, a hypothesis about the human visual system is that objects are measured with apertures roughly proportional to object width. That is, the scale of perceptual blurring is related to object width, W, by some proportionality constant, 1/k.† Thus the scale of the total composite blurring, σ_T, is the square of the scale describing the convolution of the acquisition and perceptual blurring functions σ_A and W/k, or

$$\sigma_T^2 = \sigma_A^2 + \left(\frac{W}{k}\right)^2 \qquad 4.6$$

A measure of object-relevant blur scale can be gotten by dividing the total blur squared term by the square of the width:

$$\frac{\sigma_T^2}{W^2} = \frac{\sigma_A^2}{W^2} + \left(\frac{1}{k}\right)^2 \qquad 4.7$$

Equation 4.7 is a rearrangement of Equation 4.6 so that k, which might eventually be determined from experimental data, need not be known. The $1/k^2$ term is in this way a constant and may be ignored if all that is of interest are relative levels of blur. An "effective blur scale," or EBS, that appropriately reflects object width then is the square of the acquisition blur scale normalized by object width:

---

[20] H.H. Barrett, W. Swindell, Radiological Imaging: The Theory of Image Formation, Detection, and Processing (New York: Academic Press, 1981), 137.

[21] M.L. Giger, K. Doi, "Investigation of Basic Imaging Properties in Digital Radiography. I. Modulation Transfer Function," Medical Physics 11, no. 3 (1984): 287-295.

† Section 3.2 presents recent experimental evidence that suggests that the visual system is not perfectly zoom invariant in its perception of widths. However, for simplicity this discussion will assume that there is just a single proportionality constant.

$$EBS = \frac{\sigma_A^2}{W^2} \qquad 4.8$$

A result of these principles is that the EBS for judgments about vessel stenoses is proportional to the stenotic width of the vessel. It was the intent of the angiography experiments described in Chapters 6 and 7 to study just three levels of blur. But with later realization of the importance of considering the blur as relative to object width, what really exist in that study are nine different blur conditions, one for each of the system blur and stenosis depth conditions. Chapter 6 contains a list of these EBS values. Furthermore, the statistical analyses in Chapter 7 must take into account this variation in the amount of object-relevant blur and not just the Gaussian standard deviations used to generate the blurred images.

4.3.4 Noise Simulation

The statistical fluctuations that cause the quantum mottle described in Section 4.1.1 may be modeled by the Poisson distribution. This distribution has the property that the variance is equal to the mean. Specifically, the probability, p, of obtaining a photon count of n, when the true mean count is r, is determined by Equation 4.9:[22]

$$p_n = \frac{r_n e^{-r}}{n!} \qquad 4.9$$

This is precisely the behavior observed from x-ray photons: more relative statistical fluctuation occurs with fewer quanta. For example, for 100 photons the percent standard deviation is $\frac{\sqrt{100}}{100} = 0.1$, but for 10000 photons it is $\frac{\sqrt{10000}}{10000} = 0.01$. The electronic noise in the digital equipment that is coupled with the photon capturing materials (such as sodium iodide crystals coupled to photomultipliers, or the charged selenium detectors in an area detector[23]) in a digital imaging system is also Poisson in nature.[24] Noise was added to the images in the simulation process by adding to each pixel from the blurred input image a random number drawn from a Poisson probability distribution with a variance equal to the value at that pixel.[25]

When the greyscale inversion described in Section 4.3.2 was performed, different scalings were applied as well that would dictate the amount of noise applied at this stage. The values of these scalings are listed in the experimental design (Section 6.1). By mapping the intensity values in the image to a higher range, less noise is in effect added (again, the percent standard deviation decreases as the photon count increases). The final step in the simulation process is an inverse greyscale mapping to cause the image to appear as a standard radiograph, with "white" bones and dark soft tissue. Images were remapped to the range 0-255. A typical image is shown in Figure 4.11.

A result of the visual system's strategy of measuring objects roughly in proportion to their width is that larger objects are more perceptible in the presence of noise: larger scale apertures recruited in the measurement of those larger objects average out more noise. Therefore, the influence of some amount of noise on a particular object-based task is very much dependent on the size of the objects involved in the judgment.

The implication from these principles is that the quantification of noise must include a specification of object size. It is not meaningful to refer to a signal-to-noise ratio for an image without some accompanying measurement of the size of the objects that are found and judged in the image. Morse[26] has proposed a "figure-to-noise" ratio (FNR) that is figure intensity divided by the variation in intensity when measured at a scale proportional to the width of the object. This measure logically varies throughout an image and even within a single figure.

---

[22]W.R. Hendee, R. Ritenour, Medical Imaging Physics, 3rd ed. (St. Louis: Mosby-Year Book, Inc., 1992), 286.

[23]R.E. Greene, J. Oestmann, Computed Digital Radiography in Clinical Practice (New York: Thieme Medical Publishers, 1992).

[24]H.H. Barrett, W. Swindell, Radiological Imaging: The Theory of Image Formation, Detection, and Processing, 82.

[25]A.G. Gash, "Noise," UNC Department of Computer Science Image Processing Software, 1992.

[26]B.S. Morse, "Computation of Object Cores from Grey-Level Images," Ph.D. dissertation, (University of North Carolina-Chapel Hill, 1994).

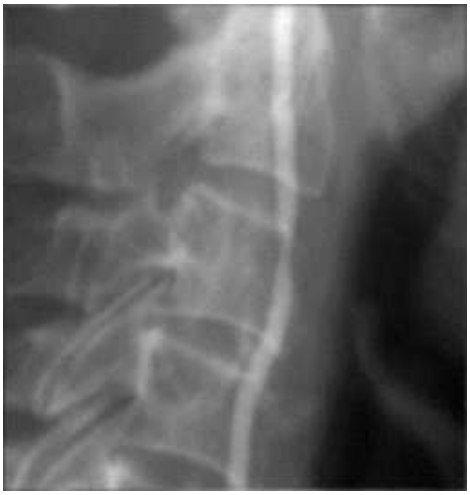

Figure 4.11. Simulated angiogram. This is an image that was used in the angiography experiment described in Chapter 6. It was simulated via the steps described in this chapter, and contains a moderate amount of blur and noise.

In the experiments described in Chapters 6 and 7, it was of interest to study three different noise levels, where the amount of noise was determined by the intensity scalings previously described. When viewing the noise conditions for the experiments in light of these figural concepts, it becomes apparent that noise is better described in units of FNR's than as greyrange scalings. More importantly, because of the different stenosis widths and levels of blur, each noise/depth/blur condition in the experiment possesses not one of three noise amounts but a unique FNR. That is, the three scalings were used to *generate* the noise conditions, but what was *produced*, when measured with the appropriate metric, are many different noise levels.

The statistical methods for analyzing the angiography experiment data in Chapter 7 utilize estimates of FNR's for the different experimental conditions. Those estimates were determined using an FNR defined as

$$\frac{I - B}{\sqrt{I}} * W \qquad 4.10$$

At pixel scale (when object width W is 1), the FNR is just the signal-to-noise ratio, or the figural intensity above background (I-B) divided by the (Poisson) noise variation, $\sqrt{I}$. The FNR then appropriately reflects the increase in perceptibility for a constant signal-to-noise ratio as object width increases. As discussed in the next paragraph, I and B are average signal and background intensities that should be measured with area-normalized apertures proportional to width W.

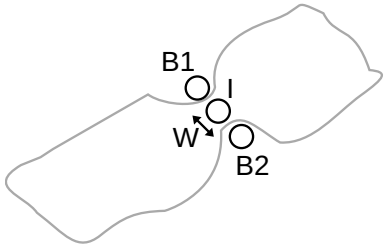

Figure 4.12. Figure-to-noise ratio intensity measurements. For the angiograms in this research, the intensities in Equation 4.6 were measured as normalized Gaussian-weighted averages with standard deviation equal to the stenotic width, W. Vessel intensity I was measured at the center of the stenosis, while B was taken as the mean of the measurements B1 and B2 made at distance $\pm$W from the stenosis center.

The actual FNR results are presented in the experimental design description (Section 6.1). Those measurements were computed for the angiographic images via Equation 4.10 using the following information (Figure 4.12). Intensity I was determined by a normalized Gaussian-weighted average at the known center of the vessel constriction using a standard deviation that was the stenotic width of the vessel. Background intensity B

was taken as the average of two measurements on either side of the stenosis. Specifically, normalized Gaussian-weighted averages, with standard deviations that were the stenotic width of the vessel, were computed at positions plus and minus one stenotic width from the center of the stenosis in the direction perpendicular to the vessel path. W was simply the known stenotic width of the vessel.

4.4 Model Computation of Stenosis Depth

This section describes specifically how the core image analysis methods were utilized in the computation of estimates for stenosis depth and how core formation and the stenosis estimates from cores might be influenced by blur and noise. Recall that it was the hypothesis that the model implementation was consistent with the human performance of this task. There is an abundance of literature relating proposed image processing methods for automatic vessel quantitation from angiograms.[27,28,29] Some of those methods are quite accurate and thus are useful diagnostic tools for computing consistent[30] estimates. It may be the case that core image analysis is itself a viable means of representing and making measurements from blood vessels.[31] However, what is desired from the core model in its incarnation in this research as a model for visual processing is that it produces estimates that are in fact systematically variable: its performance ought to be influenced in the same way that the human would as physical characteristics of the angiogram change.

Core formation was initiated and utilized as follows. An initial guess, an indication of roughly where in position and scale to search for the ridge that is the core, was provided as the known position and width of the vessel at its most constricted portion. The single core obtained from this initial guess was used to the make the stenosis estimate. A stenotic region encompassing the stenosis was designated for each vessel that was as wide as the normal width of the vessel and extending on both sides from the center of the stenosis outward to the normal portions of the vessel. Two measurements were needed from the core to make the stenosis estimate. The estimate for the constricted width, S, of the vessel was taken to be the mean of the three minimum core widths within the stenotic region. To determine the normal width, N, a "sliding" window that contained ten core points was moved along the longer of the two core segments outside the stenotic region to find that subsection of the core with minimum variance. The mean width in that minimum-variance window was used as the normal width estimate. The normal and constricted estimates were then combined to produce a percent stenosis estimate via Equation 4.1. Figure 4.13 shows a typical vessel core, while 4.14 diagrams the estimation process.

The philosophy behind these decisions about how to extract the relevant information from the core was that the human would presumably behave similarly. First, the human must first establish the position of the stenosis and then seek its most constricted point in order to make a constricted width estimate. Secondly, he/she must decide where the stenosis ends and must visually move beyond that point to make a judgment about normal. For that second measurement, the observer will likely use some region along the vessel that is most believable, or salient, or consistent.

A core was deemed sufficient to use in the estimation protocol when it spanned the entire stenotic region and possessed at least ten core points on either side outside the region. In the event that a core that met those requirements was not produced from the initial guess, two additional cores were computed with initial guesses from known positions at normal portions of the vessel above and below the stenosis. If both those cores spanned the stenotic region, their stenosis estimates were averaged. If only one of the two additional cores was successful, then it alone was used as the estimate. While occasionally, for some vessel path and background image combinations, core calculations of stenosis depth would entirely fail by this criterion (all three cores generated from the three initial guess positions failed to capture the stenosis and some portion of the normal vessel), the images used in the experiment were all ones for which a model stenosis estimate was successfully determined.

---

[27] J.R. Spears, T. Sandor, A.V. Als, M. Malagold, J.E. Markis, W. Grossman, J.R. Serur, S. Paulin, "Computerized Image Analysis for Quantitative Measurement of Vessel Diameter from Cineangiograms," Circulation 68, no. 2 (1983): 453-461.

[28] T. Sugahara, H. Maeda, Y. Yanagihara, "Automatic Detection Method of Stenotic Lesions in Coronary Cineangiograms," International Journal of Cardiac Imaging 5 (1989): 17-23.

[29] J.H.C. Reiber, P.M.J. van der Zwet, G. Koning, C.D. von Land, B. van Meurs, J.J. Gerbrands, B. Buis, A.E. van Voorthuisen, "Accuracy and Precision of Quantitative Digital Coronary Arteriography: Observer-, Short-, and Medium-Term Variabilities," Catheterization and Cardiovascular Diagnosis 28 (1993): 187-198.

[30] L.M. Zir, S.W. Miller, R.E. Dinsmore, J.P. Gilbert, J.W. Harthorne, "Interobserver Variability in Coronary Angiography," Circulation 53, no. 4 (1976): 627-632.

[31] D.S. Fritsch, D. Eberly, S.M. Pizer, M.J. McAuliffe, "Stimulated Cores and Their Applications in Medical Imaging," submitted to Information Processing in Medical Imaging '95.

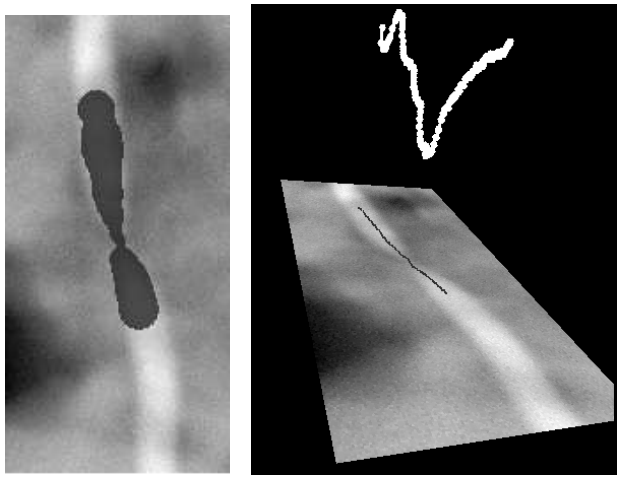

Figure 4.13. Typical vessel core. The two pictures display by different means a core calculated for a simulated vessel. On the left, the information conveyed by the core is projected onto the vessel as a series of discs of radius and position specified by the core. The three-dimensional view on the right depicts the core hovering in scale space; its position is projected onto the vessel as a thin dark line, and its height is proportional to the width of the vessel. In both cases note how the core scale reflects the constriction in the vessel.

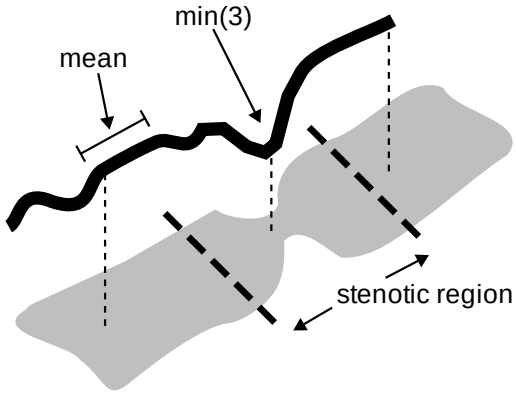

Figure 4.14. Model stenosis estimation. Model estimates for stenotic width were computed as the mean of the three minimum core widths within the designated stenotic region. The mean core width in a minimum-variance window along the longer of the two core segments outside the stenotic region was used as the normal width.

Blur and noise influence the core representation of the vessel at the two stages of core construction, medialness formation and ridge tracking. Figure 4.15 depicts medialness, as computed with a Laplacian of a Gaussian, for a one-dimensional figure profile, a step function. For that plot, medialness was measured at a scale equal to the width of the figure. Appropriately, medialness is maximal at the position of the object center. Figure 4.16 (left) shows medialness for a figure profile that is the step function from Figure 4.15 convolved with a Gaussian of standard deviation 4.0. The effect of the figure blurring on medialness is to broaden its distribution in the region of the figure and decrease its amplitude at the figure center. Both of these phenomena occur in proportion to the scale of the blurring. Figure 4.16 (right) plots the medialness for a noisy version of the step function from Figure 4.15. Poisson noise was added with a standard deviation that was 12.5 percent of the intensity of the figure center above background. Medialness for a linear convolution kernel such as the Laplacian can be separated into medialness for the figure profile and medialness for the noise. The medialness for the noise is a smooth, quasi-"periodic" distribution whose amplitude and frequency are related to the noise amplitude and frequency and the scale of the medialness measurement. The effect is that medialness maximum for the noisy figure profile can shift from the true figure center in position and scale by an amount related to those amplitude, frequency, and kernel scale characteristics.

The ridge tracking that is performed on the medialness for blurred or noisy images may suffer as a result of these degradations. When figure contrast is low and/or when figures are closely adjacent, blurring may flatten the medialness to such an extent that the eigensystem solutions used in computing ridges may produce eigenvalues of zero. Ridge tracking will thus not be initiated or will be terminated. This is equivalent in the human visual system to being unable to detect or form a representation of a figure at a particular position and scale of interest. The shift in position and scale of maximum medialness caused by noise will result in a less accurate representation of the figure. Also, the perturbations in the medialness can cause the formation of additional cores that in turn cause core termination or tracking along the "wrong" figure.

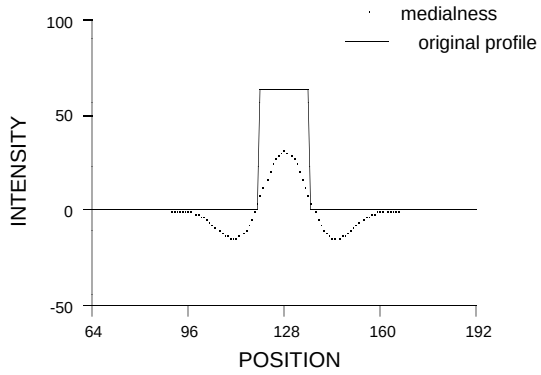

Figure 4.15. Medialness for step function. Laplacian medialness at a scale equal to the width of the figure is shown superimposed on the figure profile.

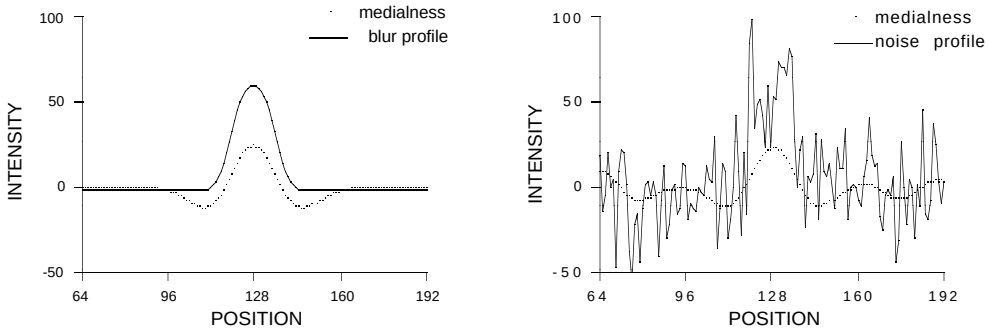

Figure 4.16. Medialness for blurred (left) and noisy (right) step function. At left, the step function from Figure 4.15 was convolved with a Gaussian of standard deviation 4.0. At right, Poisson noise with a standard deviation of 12.5% of the figure intensity above background was added to the step function. In both cases, Laplacian medialness is superimposed on the figure profiles.

This discussion of the effects of blur and noise on the model is important for several reasons. First, the angiography experiment described later in this research studied the correspondence between the model and human stenosis estimates as a function of these physical properties. It is thus necessary to understand how the model is affected by these manipulations, and it will be useful in the discussion of the experimental results to speculate whether the way that the model captures or is perturbed by blur and noise is sufficiently similar to human perception to make the model potentially predictive under those conditions. Second, sometimes the blur and noise degradations caused failures in stenosis estimation during the image generation and model computation process. Several situations under which that occurred are discussed next.

There are a number of reasons why, even with three different initial guesses, a core might not have been produced that would allow an estimate of stenosis by the previous criterion. Figure 4.17 shows several possibilities.

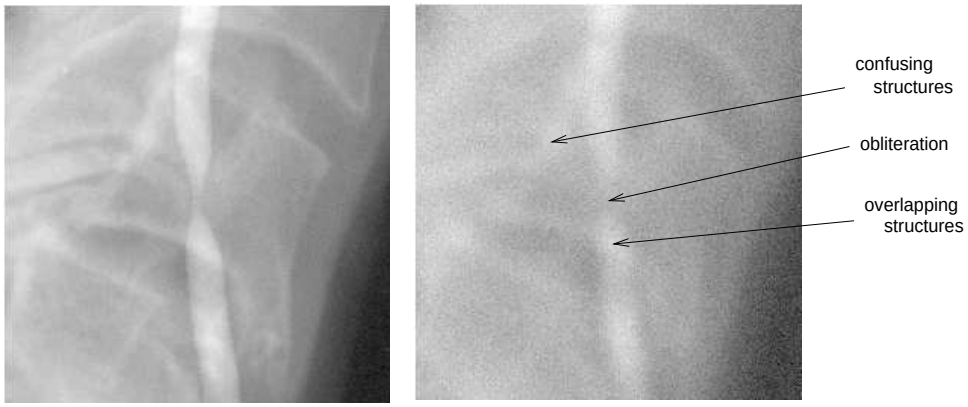

Figure 4.17. Model stenosis estimation failures. Blur and noise were added to the image on the left to create the image on the right. Obliteration of the vessel at the severe constriction, and confusing overlapping or adjacent structures, may divert or terminate core tracking.

Projection radiographs represent an integral of photon attenuation along the trajectory of the beam and thus superimpose overlying structures. All sorts of features, particularly the sharp vertebral edges in the neck region through which the internal carotid arteries pass, can appear as lying "under" the vessel and in some cases cause the termination of core tracking. Second, some of the blur and noise conditions were extreme enough to cause the vessel to virtually disappear in highly-stenosed regions. Without sufficient boundary strength to produce focused medialness in turn, the ridge conditions cannot be satisfied and core tracking fails to initiate or proceed. Finally, sometimes it is the case that adjacent cores may lie very close to the core that is being traversed. It was thus common in the low quality images for a core to get "sidetracked" and follow other prominent neighboring structures. Section 6.1 discusses how the angiographic images for the experiment were generated so as to produce a set with valid core estimates for each. This dissertation will not attempt to characterize further how and when core formation failed to come about during the image generation process. There were indeed many failures to produce cores lengthy enough to make the estimates by the described method, but most occurred in the degraded or constricted vessel images. While a characterization of when and how the failures occurred would say much about the extensibility of this approach, this research has as its scope only an initial test of its potential. Finally, improvements to the core analysis are continually developed; the stability of the methods used at the time of this research may not be entirely indicative of the capability of current implementations.

The step of providing known positions and widths as initial guesses only serves to initiate the core-finding process. The core ultimately determined will not likely pass precisely through these guess points, and therefore this step does not aid or enhance the accuracy of the model. Even if that knowledge did somehow aid core accuracy, it would presumably do so in the same manner for each image. That is, if estimates were made for the same vessel in the presence of many different physical characteristics, then the model ought to produce different estimates even though the core tracking would have been initiated in the same position each time. The purpose of this research was to measure how core estimates were influenced by variations in physical characteristics.

More sophisticated methods might have been employed to determine from the core where to take the estimate for normal width that would not have required user designation of the stenotic region. For instance, variance calculations in a window sliding along the core might have been used to determine where the width estimates "leveled off." However, this information that was provided to the model about the location of the stenotic region is again not relevant to the test of the model. The issue for this research is whether the width information derived from the core construction is correlated with human judgments, regardless of how the core might have been initiated or what other information might have been used to determine important positions in the figures that cores represent.

4.5 Summary

Angiography is an example of a medical imaging system that possesses parameters dictating its performance that a developer or user might wish to adjust in a manner that optimizes the human observer's utilization of the system. It is a common modality with a demand for image quality considerations. The accuracy of stenosis estimation, among the many tasks that might be conducted with an angiogram, is important for directing clinical intervention. Those vessel estimations are a natural application for the shape descriptions of the visual model tested in this research. For these reasons together, it was precisely this system and task that served as one basis for the model and human comparisons in the experimental investigation.

While the object representation mechanisms posited by the core model are likely to be consistent with those possessed by the human, the processes for extracting and combining that shape information to form an estimate like the stenosis judgment studied here are much less understood. The decisions that were made regarding how to generate a normal or stenotic width from the core structure were based purely on common sense. The comparisons with human performance in this research are made not as much with the core model itself as with the particular adaptation and extension of the model that was developed to calculate these clinical estimates.

The image simulations were an extensive step in this research, and they were crucial for several reasons. As was described previously, it would not have been feasible to adjust an angiography machine to produce images with different blur and noise characteristics. Simulation was the only way to achieve this variation. Likewise, simulation allows the manipulation of the physical characteristics of a single image for which all other factors, such as background anatomy and vessel position, remain constant. This aspect is important in producing images that are controlled stimuli in a psychophysical experiment. Finally, the simulated vessels provide the "truth" that could be used, together with the proof of the equivalence of the model with the human, to enable the model alone to choose optimal parameter settings. The model's accuracy (the difference between its estimate and truth) would dictate how best to set the parameters for human interpretation. This truth about the actual state of the anatomy, the size of the stenosis, would be impossible to fully determine from a standard image. The depth of a stenosis in a vessel can not ever be known, but only inferred, from viewing or even measuring an angiogram.

This chapter has described novel means of quantifying blur and noise in medical images that appropriately account for the width of the objects in question in a particular task. The measures, the EBS and FNR, were presented here in part to emphasize that they are principled perceptual descriptions of physical properties and to motivate their continued use for that purpose. It is also the case that these measures are the way that the amounts of blur and noise in each condition in the angiography experimental design will be determined.

Angiography is just one of many imaging modalities that are available for the vast array of visual diagnostic needs in medicine. The next chapter describes a second important imaging procedure and accompanying task that was examined in this research. Model performance was measured for images acquired from two imaging modalities in an attempt to demonstrate another property of model-based approaches: image quality may be measured for any imaging system for which a visual model exists for predicting performance for a relevant task.