International Journal of Computer Vision 35(1), 5–12 (1999)

Computational Vision at Yale

PETER N. BELHUMEUR
*Center for Computational Vision and Control, Departments of Electrical Engineering and Computer Science, Yale University,
New Haven CT 06520, USA*

JAMES S. DUNCAN
Imaging Processing and Analysis Group, Departments of Diagnostic Radiology and Electrical Engineering, Yale University, New Haven CT 06520, USA

GREGORY D. HAGER
Center for Computational Vision and Control, Departments of Computer Science and Electrical Engineering, Yale University, New Haven CT 06520, USA

DREW V. McDERMOTT
*Center for Computational Vision and Control, Department of Computer Science, Yale University,
New Haven CT 06520, USA*

A. STEPHEN MORSE
*Center for Computational Vision and Control, Department of Electrical Engineering, Yale University,
New Haven CT 06520, USA*

STEVEN W. ZUCKER
Center for Computational Vision and Control, Departments of Computer Science and Electrical Engineering, Yale University, New Haven CT 06520, USA

Received June 29, 1999; Revised June 29, 1999

Abstract. We present a brief introduction to the five articles that make up this special issue: *Shock Graphs and Shape Matching, The Bas-Relief Ambiguity, Incremental Focus of Attention for Robust Vision Based Tracking, What Tasks can be Performed with an Uncalibrated Stereo Vision System*, and *Volumetric Deformation Analysis Using Mechanics-Based Data Fusion: Application in Cardiac Motion Recovery*. This introduction and accompanying articles provide a by no means exhaustive, but hopefully representative sampling of the computational vision research at Yale University.

1. Introduction

Research in computational vision is inherently interdisciplinary; at Yale University it spans seven departments: electrical engineering, computer science, diagnostic radiology, neuroscience, mathematics, psychology, and statistics. The collection of articles presented in this special issue were chosen because they represent some of the core research in the areas most relevant to this journal, and because they reflect the interdisciplinary nature of the work. The papers address questions ranging from vision to robotics; from theory to application; and from mathematics to systems implementation.

These axes are only a low-dimensional projection from a much richer web of interactions, and our goal in this preface is to provide a glimpse of this richer structure "from the papers out." Each section below introduces one of the papers, and uses this introduction to sketch a roadmap of the research that we are building from it. *Shock Graphs and Shape Matching* develops a theory of generic shape recognition. *The Bas-Relief Ambiguity* considers ambiguities inherent in reconstruction and object recognition under variable illumination. *Incremental Focus of Attention for Robust Vision-Based Tracking* offers an architecture for coordinating multiple visual tracking algorithms. *What Tasks Can Be Performed with an Uncalibrated Stereo Vision System* stratifies the effects of camera calibration uncertainty on the performance of vision-based control tasks. *Volumetric Deformation Analysis Using Mechanics-Based Data Fusion: Applications in Cardiac Motion Recovery* considers the use of realistic mechanical models for recovering deformable object motion from multiple image sources, including structural information and noisy velocity estimates.

In the Section 7, we revisit the question of interactions and attempt to outline several of the meta-level questions.

2. Generic Shape Recognition and the Tangent Hypothesis

While computer vision systems can be very effective at recognizing a particular individual from a database of individuals (Georghiades et al., 1998; and next section), humans exhibit a complementary ability to recognize objects at what cognitive psychologists refer to as a "basic level." This organization is abstract with respect to particulars, and plays a role in both generalization over types of images (e.g., how can cartoons and photographs of the same person be identified?) and in organizing our memory for visual shapes (e.g., by what criteria should memory be organized for rapid indexing and retrieval?). It has even been suggested that such generic descriptions support motor tasks, such as gripping (see Section 5). In the first contribution on recognition, we explore the application of the shock-based theory of shape description to formally developing such basic level recognition strategies. In brief: given the outline of a (2-D) object, allow it to evolve according to a hyperbolic p.d.e. (Kimia et al., 1995). The singularities of this evolution provide a discretization of the space of continuous shapes into equivalence classes; these singularities are called shocks. In the paper by Siddiqi et al. (1999), it is shown that these shocks enjoy a rich syntactic structure, and that this structure can support recognition.

Taking the (projected) outline of objects as given presupposes a solution to the problem of boundary detection, to which we now turn. While this is a classical problem, the classical approaches have not produced satisfying solutions. As above we employ a mathematical abstraction, this time from differential geometry, in which the goal of local "edge detection" is to return discrete tangent approximations to the equivalence class of curves that could pass through a (discrete) image point with a (quantized) orientation. The difference from the classical approach is that, rather than approaching edge detection as a linear filter followed by a selective non-linearity, we design non-linear filters that enforce continuity (from the limiting process inherent in defining the tangent) and photometric conditions (Iverson and Zucker, 1995). Consistency between such local measurements is enforced by co-circularity in a relaxation labeling process (Zucker et al., 1989), where tangents are implicitly transported along the osculating circle. In effect, the local tangents are consistent up to curvature almost everywhere, and discontinuities are represented as multiple tangents at the same location.

Such a scheme is radical in several respects. First, it suggests operating at a mathematical level abstract with respect to images, and places the tangent bundle in a central position. Of course, this bundle must be extended to handle the discontinuities. Second, a biological substrate is suggested (identify an orientation hypercolumn with orientation fibre in the above bundle, and the endstopped property of neurons with curvature representation (Dobbins et al., 1987). The relaxation process can then take place within cliques of pyramidal cells (Miller and Zucker, 1999), and particular neural computations can be formulated by relating relaxation labeling to polymatrix games (Miller and Zucker, 1992). Such models may be interesting to computational theorists not concerned with neurobiology, because they suggest how reliability can be achieved in a way that utilizes redundancy to increase accuracy; for related uses of game theory, see Section 6.

One measure of the strength of a theory lies in its predictions; we mention recent psychophysical tests of boundary (Elder and Zucker, 1998) and shape

(Siddiqi et al., 1999a, 1999b) features. These latter experiments are remarkable, in that the shock abstraction reveals aspects of human shape perception. This motivated a search for more biologically-plausible graph-isomorphism mechanisms to support matching (Pelillo et al., 1998), and for dynamical systems other than curve evolution that can solve the p.d.e.'s (Siddiqi et al., 1999). In a technical sense these dynamical systems connect to the tangent information supplied above, while, in the larger context of curve evolution, the Hamiltonian re-formulation may have much wider application. But in a sense what is most exciting are those connections that are developing in intermediate-levels of vision between the tangent abstraction and the shape abstraction (Dubuc and Zucker, 1995; August et al., 1999a; August et al., 1999b). This suggests a much richer view of intermediate-level vision, and the principles driving perceptual organization. With this support indexing experiments now can be performed (Shokoufandeh et al., 1999), and we are developing extension to the full 3-D problem.

3. Generative Models for Object Detection, Tracking, and Recognition

The complex issues associated with developing a generic description of shape (as described above) are quite different from those arising in the design of a system for the recognition of rigid objects from a pre-described database. These differ in the sense that in the former, one considers the problem of generalization and related issues of categorization, while in the latter, one restricts the problem to recognition within a given category or collection of objects. The problem of generalizing to categories of objects from a collection of exemplars and recognizing new instances of a category seems to require extracting suitable abstractions from the images. In contrast, more concrete representations of an object's appearance may be needed to recognize specific objects within a category.

In either case, a significant obstacle in the development of recognition systems is the variability of an object's appearance from one image to the next. In particular, with slight changes in lighting conditions and viewpoint often come large changes in the object's appearance. In an effort to tame the variability in the images of the same object seen under different conditions, researchers in computer vision and robotics have long searched for representations that are invariant to lighting configurations and viewing direction of the camera.

Only recently have "appearance-based" approaches been developed in an effort to use intensity information to model or learn a representation that captures a large set of the possible images of an object under pose and/or illumination variation (Murase and Nayar, 1995; Pentland et al., 1994; Poggio and Sung, 1994; Sirovitch and Kirby, 1987; Turk and Pentland, 1991; Tagare et al., 1998). These methods have gone a long way in demonstrating the advantages of using much richer descriptions than simply sparse features like edges and corners for recognition. Still, a drawback of these approaches is that in order to recognize an object seen from a particular pose and under a particular illumination, the object must have been previously imaged under the same conditions.

This raises the question: Is there some underlying "generative" structure to the set of images of an object under varying illumination, pose, and shape such that to create the set, the object does not have to be viewed under all possible conditions? We have focused research on modeling the set of images of an object under all sources of variability from a small number of examplars. In the process we have considered a number of questions: What is the set of images of an object under all possible changes in illumination (Belhumeur and Kriegman, 1998), pose (Hager and Belhumeur, 1998), and shape? What is the connection between the image changes due to illumination and image changes due to pose (Hager and Belhumeur, 1998; Belhumeur et al., 1997)? And, finally, how can we design computer vision algorithms for object recognition (Georghiades et al., 1998), 3-D reconstruction (Georghiades et al., 1999), visual tracking (Hager and Belhumeur, 1998; Belhumeur and Hager, 1999), and human/computer interfaces (Belhumeur and Hager, 1999) that take advantage of the answers to these questions? We have applied these solutions largely in the context of face recognition (Belhumeur et al., 1997; Georghiades et al., 1998) and face tracking (Belhumeur and Hager, 1999), though the resulting techniques can be readily applied in other application domains.

This section introduces a paper that arose out of our characterization of the set of images of an object under varying illumination conditions. In Belhumeur and Kriegman (1998) were able to show that for a fixed pose the set of images of an object under all possible illumination conditions is a convex cone in the image space. For objects with Lambertian reflectance, we also showed that this illumination cone can be generated from as few as three properly illuminated images of the object. Since the illumination cone is the complete set of images of an object under all possible lighting conditions, it is natural to ask: When do two objects of differing shape possess the same illumination cone and, thus, are indistinguishable from a particular vantage point?

The paper presented here—*The Bas-Relief Ambiguity*—considers this question showing that there is a three parameter family of transformations (generalized bas-relief transformations) on shape and a corresponding family of transformations on albedo that, when viewed orthographically, preserve the illumination cone of an object with Lambertian reflectance. For each image of the object illuminated by an arbitrary number of distant light sources, there exists an identical image of the transformed object illuminated by similarly transformed light sources. This result holds both for the illuminated regions of the object as well as those in cast and attached shadows. (Shadow ambiguities and perspective projection are further explored in (Kriegman and Belhumeur, 1998)). Thus from a fixed viewpoint with varying unknown lighting, one can neither distinguish objects differing by generalized bas-relief transformations, nor determine an object's precise Euclidean shape.

4. Visual Tracking

Given these attempts to formulate object characterizations from static imagery, we now turn to developing models and algorithms for time-varying imagery. Although these are traditionally rooted in vision problems such as estimating optical flow and/or ego-motion, many techniques have evolved to be closer to the theory of dynamical systems. This is pleasing given that observing cameras and the observed world are governed by a set of well-understood physical principles, and that these principles already emerged in the previous section characterizing the variation of images due to changing illumination.

More generally, we feel the core in visual tracking is the development of effective models for the variation of images over time—models that include the effect of illumination (Hager and Belhumeur, 1998) and pose (Belhumeur and Hager, 1999). However, the ultimate goal of such modeling is to develop efficient and robust algorithms by explicitly taking advantage of the time varying properties of images. For example, we have shown that the visual tracking problem can be posed as one of (dynamically) stabilizing image variability over a class of allowed changes in images over time (Hager and Belhumeur, 1998). The result is an efficient class of tracking algorithms with provable performance properties.

In addition to basic questions in modeling and algorithm design, we are interested in developing effective software systems for visual tracking: systems that are portable, practical, and well-targeted toward this problem domain. The XVision system (Hager and Toyama, 1998) was our first attempt in this direction (Hager and Toyama, 1998). XVision abstracts the notion of a "tracked object" as a state-based system, and allows objects to be combined in a natural manner using inter-object constraints, much as in current animation systems. The Frob system (Peterson et al., 1999; Reid et al., 1999) is an example of work on a new generation of languages for visual tracking and control, in which new application-specific functional programming techniques are utilized.

Although visual tracking problems often admit clean and clear solutions in the abstract, in the real world robust visual tracking is elusive. The article by Toyama and Hager (this volume) is one approach to developing robust tracking methods. The ideas that motivate this approach are simple: first, algorithms that compute more information about an object tend to be more sensitive to imaging conditions; and second, collections of tracking algorithms with different failure modes and sensitivities can often be combined using a switching logic to produce something more robust than any of the pieces individually. The result is a layered architecture in which simple, robust tracking algorithms are used to "prime" more complex ones. It is encouraging that, although simple, this idea leads to an intuitive algorithm design strategy that enjoys a principled means of mathematical evaluation.

There are many other approaches to robust visual tracking. For example, combining multiple cues within an estimation framework can greatly improve the performance of tracking (Rasmussen and Hager, 1998). Likewise, the extra information available when embedding vision within an application can be used to improve tracking performance. For example, in vision-based control (see Section 5), we can often use the fact that we know the control inputs to the controlled system as feed-forward to the tracking algorthm.

All of these ideas ultimately reflect our initial thesis: visual tracking is ultimately a problem in the design and analysis of time-varying systems for estimation and control. The research at Yale is indicative of a larger movement in this direction (Kriegman et al., 1998).

5. Vision-Based Control

While attempts to link computer vision with actuated systems go back to the roots of the field, many of the early approaches separated vision from control: vision was used as a means to reconstruct a geometrically accurate model of the environment, and this reconstruction was then used as a basis for robot control.

While conceptually simple, this approach has the drawback that it relies on extremely general, robust, and accurate solutions to the reconstruction problem. An alternative that has become far more popular over the last decade is to integrate vision with control, and to develop feedback laws that work directly from image information with no intervening 3D representation. Developing practical vision-based feedback algorithms of this sort that have provable performance properties poses a number of problems not encountered in the design of more conventional sensor-based systems. In particular, because video signals are inherently a nonlinear, i.e., derive from a two-dimensional projection of a three-dimensional world, many conventional control paradigms (e.g., set-point regulation via error feedback) do not readily apply to vision-based feedback systems.

Our research is concentrated on several fronts. The first, represented by the article in this volume, is to develop a formal, structured basis for vision-based control. In particular, the reason to use vision in a feedback loop is to eliminate positioning errors that would occur in an open-loop system. However, as it turns out, many approaches to vision-based control (for example, those employing reconstruction) are themselves subject to positioning errors if the underlying camera system and/or actuator is not well-modeled. The article in this issue addresses exactly this issue—namely *what* tasks a given image-based servoing framework can perform with precision, even if the underlying camera and/or robotic system are mis-calibrated.

One interesting outgrowth of this characterization is a structured representation for visual serving tasks. As a result it is possible to develop an algebra of "performable" tasks which, in turn, forms the natural basis for a vision-based programming environment (Dodds et al., 1999; Toyama et al., 1996). A similar line of research is being followed for vision-based mobile systems (Rasmussen and Hager, 1996; Hager et al., 1998). Here, the problem is not so much the design of a class of control algorithms, as it is the design of robust algorithms than can handle the field of view limitations that occur as a mobile system moves through the world.

Finally, we note that a key limiting factor in vision-based control can be *robust vision* (Vincze and Hager, 1999). As a result, many of the ideas behind the layered vision architectures described above find application in vision-based control (Dodds et al., 1999). It is interesting to speculate whether these layered architectures have counterparts with those in biological vision systems. And more abstractly, it is interesting to question whether generic visual descriptions (Section 2) are relevant to manipulation strategies; i.e., should all parts be "grippable?" We expect these and similar developments at other institutions to be fundamental to the development of effective, practical, yet well-founded vision-based control systems.

6. Model-Based Medical Image Analysis

Medical image analysis can no longer be considered simply an application area for computer vision. The distinctive nature of the problems encountered have led to the development of a significant body of work addressing such issues as fully three-dimensional data, nonrigid models for motion, deformation and comparison, and the statistical variation of normal and abnormal structure. Of particular note has been the development and use of physical models both for enforcing continuity and smoothness conditions, as well as truly modeling the underlying tissue properties. While this area of research derives its relevance from the clinical applications, which must be well understood, the methodologies encompass an array of techniques that have advanced computer vision independently of the application.

A broad range of work in this area has been pursued at Yale. A variety of methodologies for the segmentation of anatomic structure, such as the heart and the brain, using deformable models (Staib and Duncan, 1992; Wang and Staib, 1998; Zeng et al., 1998) have been developed. A key theme has been the incorporation of prior information of shape as constraints or biases and the integration of multiple information sources (Chakraborty et al., 1996; Chakraborty and Duncan, 1999). Such information is key to the solution of medical image segmentation problems in the face of noise, ambiguity and structural complexity. It is interesting to note the interplay between these particular segmentation solutions and those developed for more generic tasks (Section 2). Finally, by using the strategy of deforming templates along orthogonal curves, it is possible to obtain globally optimal fits of doformable templates to data (Tagare, 1977a).

The tracking and modeling of the non-rigid motion of the heart has been a major effort here at Yale (e.g. Duncan et al., 1998; Tagare, 1997, 1999). An important facet of this work has been the development of realistic physical models. The article by Shi et al. (this volume) approaches this problem using a continuum mechanical model of the heart integrating shape-based estimates of displacement with velocity estimates measured using phase velocity magnetic resonance imaging. The problem is solved in a finite element framework. The integration of shape with direct, dense measurements of three-dimensional velocity provides a powerful set of constraints on the interpretation of motion from image data. We are applying similar ideas in soft tissue modeling to the computation of brain deformation during neurosurgery (Škrinjar et al., 1998) and well as further developing the heart model to include muscle fiber directions (Papademetris et al., 1999). Finally, we are using a point matching approach (Rangarajan et al., 1997) incorporating a spline-based deformation to perform non-rigid mapping of anatomic brain volumes (Rangarajan et al., 1999; Chui et al., 1999).

Throughout our work, we have endeavored to take advantage of all available information both from prior knowledge of physical properties, geometric constraints or statistical variation as well as imaging data from various modalities. Medical imaging problems typically lie in a well-defined domain. The challenge is to take best advantage of the given domain to solve the often complex and subtle problems posed.

7. Summary

The selection of articles in this special issue provides an instantaneous cut through five research projects at Yale, and indicates some of the dimensions along which inter-relationships occur. The first two papers, on object representation and recognition, raise differences between generality and specificity, and exemplify the very different mathematical techniques implied. The next two papers raise questions about vision and control, and highlight the importance of image sequences and feedback. The final paper, in biomedical image analysis, exemplifies the important differences between 2D, 3D, and 4D data. There is no doubt that recent computer and imaging hardware advances have made the experiments and analysis in many of these papers feasible.

Any overview of research such as this provides a forum for raising more questions than are answered. What representations are appropriate for object recognition? How are the dynamical systems that arise in relaxation labeling related to those that arise in feedback control theory? Even though a matching problem involves discrete representations, are continuous techniques appropriate? Are techniques that are effective on 2D image data appropriate for 3D or 4D data gathered using different imaging modalities? What is the role of biology in stimulating computational theory, or of applications in stimulating mathematical questions? The list is endless.

Perhaps the final inference that one can draw from an overview of research such as this relates to the style and personality of an institution. We hope we have provided a suitable picture of research in computational vision at Yale, and we invite all readers to visit our web pages (`http://cvc.yale.edu`) for a constant update and paper archive.

References

August, J., Siddiqi, K., and Zucker, S.W. 1999a. Ligature instabilities in the perceptual organization of shape. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition*.

August, J., Tannenbaum, A., and Zucker, S.W. 1999b. On the evolution of the skeleton. In *Int. Conf. on Computer Vision*.

Belhumeur, P.N. and Hager, G.D. 1999. Tracking in 3D: Image variability decomposition for recovering object pose and illumination. *Pattern Analysis and Applications*, 2(1):82–91.

Belhumeur, P.N., Hespanha, J.P., and Kriegman, D.J. 1997. Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection. *IEEE Trans. Pattern Anal. Mach. Intelligence*, vol. 19, pp. 711–720.

Belhumeur, P.N. and Kriegman, D.J. 1998. What is the set of images of an object under all possible lighting conditions. *Int. J. Computer Vision*, 28:245–260.

Belhumeur, P.N., Kriegman, D.J., and Yuille, A.L. 1997. The bas-relief ambiguity. In *IEEE Proc. Conf. Computer Vision and Pattern Recognition*.

Chakraborty, A. and Duncan, J. 1999. Game theoretic integration for image segmentation. *IEEE Trans. Pattern Anal. Machine Intell.*, 21(1):12–30.

Chakraborty, A., Staib, L.H., and Duncan, J.S. 1996. Deformable boundary finding in medical images by integrating gradient and region information. *IEEE Trans. Med. Imaging*, 15(6):859–870.

Chui, H., Rambo, J., Duncan, J., Schultz, R., and Rangarajan, A. 1999. Registration of Cortical Anatomical Structures via Robust 3D Point Matching. Information Processing in Medical Imaging, LNCS Vol. 1613, pp. 168–181.

Dobbins, A., Zucker, S.W., and Cynader, M. 1987. Endstopped neurons in the visual cortex as a substrate for calculating curvature. *Nature*, 329:438–441.

Dodds, Z., Hager, G., Morse, A., and Hespanha, J. 1999. Task specification and monitoring for uncalibrated hand/eye coordination. In *Proc. IEEE Int. Conference Rob. Automat.*, pp. 1607–1613.

Dodds, Z., Jagersand, M., and Hager, G.D. 1999. A hierarchical architecture for vision-based robotic manipulation tasks. In *First Int. Conf. on Computer Vision Systems*, H.I. Christensen (Ed.), vol. 1542 of *Lecture Notes in Computer Science*, pp. 312–330.

Dubuc, B. and Zucker, S.W. 1995. Indexing visual representations through the complexity map. In *Fifth International Conference on Computer Vision*.

Duncan, J., Shi, P., Constable, T., and Sinusas, A. 1998. Physical and geometrical modeling for image-based recover of left venticular deformation. *Progress in Biophysics and Molecular Biology*, 69:333–351.

Elder, J. and Zucker, S.W. 1998. Evidence for boundary-specific grouping in human vision. *Vision Research*, 38:143–152.

Georghiades, A., Belhumeur, P.N., and Kriegman, D.J. 1999. Illumination-based image synthesis: Creating novel images of human faces under differing pose and lighting. In *MVIEW'99 Workshop at IEEE Conf. CVPR*.

Georghiades, A., Kriegman, D.J., and Belhumeur, P.N. 1998. Illumination cones for recognition under variable lighting: Faces. In *IEEE Conf. CVPR*.

Hager, G.D. and Belhumeur, P.N. 1998. Efficient region tracking with parametric models of geometry and illumination. *IEEE Trans. Pattern Anal. Mach. Intelligence*.

Hager, G., Kriegman, D., Georghiades, A., and Ben-Shahar, O. 1998. Toward domain-independent navigation: Dynamic vision and control. In *Proc. IEEE Int. Conf. on Dec. and Control*.

Hager, G.D. and Toyama, K. 1998. The “XVision” system: A general purpose substrate for real-time vision applications. *Computer Vision and Image Understanding*, 69(1):23–27.

Iverson, L.A. and Zucker, S.W. 1995. Logical/linear operators for image curves. *IEEE PAMI*.

Kimia, B.B., Tannenbaum, A., and Zucker, S.W. 1995. Shape, shocks, and deformations I: The components of two-dimensional shape and the reaction-diffusion space. *Int. J. Computer Vision*, 15:189–224.

Kriegman, D.J. and Belhumeur, P.N. 1998. What shadows reveal about object structure. In *European Conf. on Computer Vision*.

Kriegman, D., Hager, G., and Morse, A. (Eds.) 1998. *The Confluence of Vision and Control*. Number 237 in Lecture Notes in Control and Information Sciences. Springer-Verlag.

Miller, D. and Zucker, S.W. 1992. Efficent simplex-like methods for equilibria of nonsymmetric analog networks. *Neural Computation*, 4:167–190.

Miller, D. and Zucker, S.W. 1999. Computing with self-excitatory cliques: A model and an application to hyperacuity-scale computation in visual cortex. *Neural Computation*, 11:21–66.

Murase, H. and Nayar, S. 1995. Visual learning and recognition of 3-D objects from appearence. *Int. Journal of Computer Vision*, 14:5–24.

Papademetris, X., Shi, P., Dione, D., Sinusas, A., Constable, R., and Duncan, J. 1999. Recovery of soft tissue object deformation using biomechanical models. In *Information Proc. Med. Imaging*.

Pelillo, M., Siddiqi, K., and Zucker, S.W. 1998. Matching hierarchical structures using association graphs. In *Proc. European Conf. on Computer Vision*.

Pentland, A., Moghaddam, B., and Starner. 1994. View-based and modular eigenspaces for face recognition. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 84–91.

Peterson, J., Hager, G., and Hudak, P. 1999. A language for declarative robotic programming. In *Proc. IEEE Int. Conference Rob. Automat.*, pp. 1144–1151.

Poggio, T. and Sung, K. 1994. Example-based learning for view-based human face detection. In *Proc. Image Understanding Workshop*, vol. II, pp. 843–850.

Rangarajan, A., Chui, H., and Bookstain, F. 1997. The Soft Assign Procrustes Matching Algorithm Information Processing in Medical Imaging, LNCS Vol. 1230, pp. 29–42.

Rangarajan, A., Chui, H., and Injolsness, E. 1999. A New Distance Measure for Non-Rigid Image Matching. In Energy Minimization Methods in Computer Vision and Pattern Recognition, LNCS Vol. 1654, pp. 237–252.

Rasmussen, C. and Hager, G.D. 1996. Robot navigation using image sequences. In *Proc. Am. Assoc. Art. Intell.*, pp. 938–943.

Rasmussen, C. and Hager, G. 1998. Joint probabilistic techniques for tracking multi-part objects. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pp. 16–21.

Reid, A., Peterson, J., Hager, G., and Hudak, P. 1999. Prototyping real-time vision systems: An experiment in DSL design. In *Proc. 21st Int. Conference Software Engineering*, pp. 484–493.

Shokoufandeh, A., Dickinson, S., Siddiqi, K., and Zucker, S.W. 1999. Indexing using a spectral encoding of topological structure. In *CVPR'99 (IEEE Computer Society Conference on Computer Vision and Pattern Recognition)*.

Siddiqi, K., Bouix, S., Tannenbaum, A., and Zucker, S.W. 1999. The hamilton-jacobi skeleton. In *International Conference on Computer Vision*.

Siddiqi, K., Kimia, B., Tannenbaum, A., and Zucker, S.W. 1999a. Categorical features in shape perception. In *Investigative Opth. and Vis. Science*.

Siddiqi, K., Kimia, B., Tannenbaum, A., and Zucker, S.W. 1999b. Shapes, shocks, and wiggles. *Image and Vision Computing*, 17: 365–373.

Siddiqi, K., Shokoufandeh, A., Dickinson, S., and Zucker, S.W. 1999. Shock graphs and shape matching. *International Journal of Computer Vision*.

Sirovitch, L. and Kirby, M. 1987. Low-dimensional procedure for the characterization of human faces. *J. Optical Soc. of America A*, 2:519–524.

Škrinjar, O., Spencer, D., and Duncan, J.S. 1998. Brain shift modeling for use in neurosurgery. In *Medical Image Computing and Computer-Assisted Intervention (LNCS 1496)*, Berlin, Springer, pp. 641–649.

Staib, L.H. and Duncan, J.S. 1992. Boundary finding with parametrically deformable models. *IEEE Trans. Pattern Anal. Machine Intell.*, 14(11):1061–1075.

Tagare, H.D. 1997a. Deformable 2D template matching using orthogonal curves. *IEEE Trans. on Medical Imaging*, 16(1):108–117.

Tagare, H.D. 1997b. Non-Rigid Curve Correspondence for Estimating Heart Motion, Information Processing in Medical Imaging, LNCS Vol. 1230, pp. 489–494.

Tagare, H.D., McDermott, D.V., and Xiao, H. 1998. Visual place recognition for autonomous robots. In *Int. Conf. on Robotics and Automation*, pp. 2530–2535.

Tagare, H.D. 1999. Shape-based nonrigid correspondence with application to heart motion analysis. *IEEE Transactions on Medical Imaging*, 18(7):570–579.

Toyama, K., Wang, J., and Hager, G.D. 1996. SERVOMATIC: A modular system for robust positioning using stereo visual servoing. In *Proc. IEEE Int. Conference Rob. Automat.*, pp. 2636–2643.

Turk, M. and Pentland, A. 1991. Eigenfaces for recognition. *J. of Cognitive Neuroscience*, 3(1).

Vincze, M. and Hager, G. (Eds.) 1999. *Robust Vision for Vision-Based Control of Motion*. IEEE Press.

Wang, Y. and Staib, L.H. 1998. Boundary finding with correspondence using statistical shape models. In *Computer Vision and Pattern Recognition*, pp. 338–345.

Zeng, X., Staib, L.H., Schultz, R.T., and Duncan, J.S. 1998. Volumetric layer segmentation using coupled surfaces propagation. In *Computer Vision and Pattern Recognition*, pp. 708–715.

Zucker, S.W., Dobbins, A., and Iverson, L. 1989. Two stages of curve detection suggest two styles of visual computation. *Neural Computation*, 1:68–81.