International Journal of Computer Vision 54(1/2/3), 5–11, 2003
© 2003 Kluwer Academic Publishers. Manufactured in The Netherlands.

# Guest Editorial: Computational Vision at Brown

MICHAEL J. BLACK
*Department of Computer Science, Box 1910, Brown University, Providence, RI 02912, USA*
black@cs.brown.edu

BENJAMIN B. KIMIA
*Division of Engineering, Box D, Brown University, Providence, RI 02912, USA*
kimia@lems.brown.edu

From the outside it may not be apparent that Brown University has a large, interdisciplinary, and vibrant computer vision community. Despite being a small school (5,674 undergraduate and 1,343 graduate students), Brown has a tightly knit community of vision researchers in various departments. This can be partly seen in this special issue which has contributions from researchers in the Division of Applied Mathematics, the Department of Computer Science, the Division of Engineering, and the Department of Cognitive and Linguistic Sciences. What we find amazing and wonderful about Brown is the high degree of interaction among researchers from these different disciplines.

We hope the papers in this special issue give some glimpse into the computational vision research at Brown. Any collection of this type is only a snapshot at an instant in time and cannot hope to capture the full diversity of work going on here. While there is a broad range of vision research at Brown, our focus here is specifically on conveying a sense of the breadth and depth of *computational vision* research at Brown. Towards that end, we briefly summarize some of the current research efforts that are not represented by articles in this issue and provide some references for the interested reader.

## The Division of Applied Mathematics

*Elie Bienenstock* (Bienenstock et al., 1997; Bienenstock and Doursat, 1994; Geman et al., 1992; Malsburg and Bienenstock, 1987) studies the "compositional" nature of vision and object recognition. Compositionality refers to the ability of natural-and some artificial—vision systems to represent objects as hierarchies of reusable parts. For example, a face is made of eyes, nose, mouth, and each of these is made of simpler constituents. Constituents can come together in an image only if they obey specific relational rules. The compositionality principle has implications for both computational and biological systems. In the brain for example, there must be a mechanism for rapidly and reversibly binding otherwise uncorrelated spatio-temporal patterns of neural activity.

*Stuart Geman* (Geman and Geman, 1984; Geman et al., 1990, 1993, 2002; Kunsch et al., 1995) also studies compositionality. Compositional hierarchies can be viewed as parse trees derived under formal grammars. This sets up an equivalence between parsing and scene analysis, and it suggests the use of probabilistic grammars for representing the relative likelihoods of different aggregations of parts and objects. Since the well-studied context-free grammars are inadequate, it is important to craft distributions for more general grammars and to specify feasible algorithms for the inference of these grammars. One natural way to implement compositional hierarchies in biological vision systems would be to exploit fine-temporal structure of neuronal spike trains. Statistical methods are being devised to systematically search for fine-temporal structure in multi-unit neural recordings.

*David Mumford* studies vision from a broad range of perspectives (Mumford and Gidas, 2001; Lee et al., 2003; Mumford, 2002, 2003; Lee and Mumford, to appear). His focus has been on constructing probabilistic models for the "variables of vision." This includes capturing the statistics of natural and range images and modeling occlusion and local image structure. His work also explores the concept of *similarity* between two shapes which is fundamental to object recognition. Finally, he is exploring how our understanding of computer vision corresponds to perceptual processes in the human brain.

*Basilis Gidas*'s current research in computer vision includes: (i) Simultaneous tracking and recognition of moving objects on the basis of video images; the work explores hierarchical/syntactic (context-free-grammar type) object representations and Monte Carlo type filters; (ii) design of scale-invariant, non-Gaussian, 2-D stochastic models for image generation and processing. His past work includes estimation of topographic maps ("shape from shading"), tomographic reconstruction, texture modeling and segmentation, multiscale Monte Carlo simulation and optimization algorithms for image processing tasks (Gidas et al., 2000, 2002; Gidas and Mumford, 2001; Gidas and Geman, 1991; Gidas, 1989).

*Ulf Grenander* is the originator of Pattern Theory whose aim is to analyze from a statistical point of view the patterns in images and other signals. This is an area of strength for Brown where numerous other researches actively use this approach. He is a pioneer of the notion of *computational atlases* where the anatomy is digitally represented as patterns observed in 3D medical image datasets, e.g., from CT, MR, and PET (Grenander and Miller, 1998; Bakircioglu et al., 1998; Miller et al., 1993). The goal is to facilitate automatic recognition and component measurement of anatomical structures in these images. Another current interest is Automatic Target Recognition (Grenander and Srivastava, 2001; Grenander et al., 1998) where the research primarily involves developing jump-diffusion algorithms for pattern inference based on knowledge about vehicle dynamics and the physics of sensors. As another example of how computer vision can be studied from the pattern theory perspective, shape can be modeled as templates deformed by Lie group transformations.

*Donald McClure* is interested in problems that benefit from the use of temporal as well as spatial information for the analysis of image sequences. Examples include the detection and removal of defects from motion pictures, resolution conversion, image stabilization and region-based decomposition of image sequences for more effective motion image compression (Geman et al., 1992, 1993; Geman and McClure, 1987; Kutliroff, 2002)

**The Department of Computer Science**

*Michael Black*'s research focuses on motion estimation and motion understanding in video sequences. In particular, his work on optical flow addresses the detection of occlusion boundaries and the estimation of image motion in layers using robust statistical methods and probabilistic models (Black and Fleet, 2000; Black and Anandan 1996). Understanding the motion of the human body is of special interest but is made challenging by the variability of human appearance, the high dimensionality of articulated body models, and the complexity of human motion. His work on this topic exploits learning methods and probabilistic inference techniques to model the motion of the face and body, their appearance, how they change due to motion, and to track deformable face or articulated body models in image sequences (Black and Yacoob, 1997; Black and Jepson, 1998; Black et al., 2000).

*Thomas Hofmann*'s research focuses on the theoretical foundations of learning with applications to computer vision and pattern recognition in addition to natural language learning and data mining (Andrews et al., 2003; Hofmann, 2001; Puzicha et al., 1999; Hofmann et al., 1998). Like many researchers at Brown with an interest in computational vision, his focus is on statistical and information theoretic techniques.

*David Laidlaw* works primarily in the areas of scientific visualization, computational modeling, and computer graphics. His vision-related focus is on developing tools to facilitate scientific understanding of medical imaging data and results of simulations (van Dam et al., 2002). He has worked on geometric model extraction from MR data (Andrews and Laidlaw, 2002), classification of MR data using Bayesian modeling of partial volume mixing (Cooper et al., 2002), and tools for visually interacting with multi-valued volume data. To develop applications for his vision-related tools, he collaborates widely with researchers in many other disciplines (Grimm et al., 2002).

**The Division of Engineering**

*David Cooper* (Cooper et al., 2001; Kang et al., 2001; Tarel and Cooper, 2000; Blane et al., 2000; Barzohar and Cooper, 1996) has been a pioneer of applying the Bayesian approach to computer vision. His recent research focuses on: (i) Automatic extraction of 3D geometrically, algebraically, stochastically modeled structure and semantically meaningful geometry from unordered sets of 3D data points or from images from multiple cameras. (ii) Applications to archaeological site data, which includes: Reconstruction of 3D ceramic pot models by automatically assembling laser scans of the many small pot sherds; Indexing into databases of extracted 3D geometric representations; Semi-automatic site reconstruction for VR viewing. (iii) Using 3D sketching and user-controlled interpolation for high precision, user-intuitive, sculpting in a VR environment. (iv) 3D geometric learning; indexing into image and video databases.

*Benjamin Kimia*'s work focuses on shape representation, object recognition, and perceptual grouping. Robust recognition in the presence of numerous sources of variability requires an explicit topology for shape, constructed by formally studying the transitions of the *shock graph* (with Giblin, this volume), which is a variant of the medial axis (with Tek, this volume). Recognition is based on a metric of dissimilarity which is the cost of the least action deformation path. Computations are made practical by defining equivalence classes and developing an edit-distance approach to shock graph matching. Other research interests include perceptual grouping using shock graphs, exploring the viability of this framework from psychophysics and neurophysiology perspectives, medical imaging, and digital archaeology (Kimia et al., 1992; Sebastian et al., 2003; Siddiqi et al., 2001; Sebastian et al., 2003; Kimia, 2003).

*Joseph Mundy* is working in the area of video processing and analysis with particular emphasis on object recognition and modeling. A current research theme is the concept of active models which combine the activities of modeling and recognition. In this approach, the object model acts as a recognition agent and refines its recognition strategy as more sensor information is accrued. The object representation combines both 3-d geometric and illumination scattering models (Mundy, 1995, 1998; Rothwell et al., 1993; Mundy and Zisserman, 1992)

*Gabriel Taubin*'s main research interests include: Computer Vision, Applied Computational Geometry, Computer Graphics, Geometric Modeling, and 3D Photography (Taubin, 2002; Balmelli et al., 2002; Bernardini et al., 2002; Taubin et al., 1998; Taubin and Rossignac, 1998). For the last few years his main line of research has been on the development of efficient, simple, and mathematically sound algorithms to operate on 3D objects represented as polygonal meshes, with an emphasis on Web-based applications. This work has spanned several areas such as: 3D capturing and surface reconstruction, modeling, compression, progressive transmission, and display of polygonal meshes, and mesh signal processing.

*William Wolovich*'s primary research areas are: (i) Object Modeling/Measurement, where a new conic-line decomposition of algebraic curves has been developed. This decomposition can be used to analyze the boundaries of free-form objects, to define canonical curves for object recognition, and to obtain geometric invariants for object identification and classification. (ii) Geometric Design, where new methods have been developed for blending multiple two-dimensional profile curves to produce three-dimensional sweep surfaces that are easy to model, measure and modify. (iii) Motion Modeling and Control, where motion signature surfaces can be created to analyze databases containing normal and abnormal motion for multiple applications, such as the diagnosis of ergonomic injuries.

### The Department of Cognitive and Linguistic Sciences

*Fulvio Domini* studies how the human visual system can exploit the 2D flow of image features to recover structure and motion. In particular he is concerned with mathematical models of how different cues are combined by the brain to interpret a moving scene. In psychophysical experiments, he has shown (i) the visual system relies on properties of the optic flow that are not necessarily sufficient for deriving the projected object and its 3D motion; (ii) the derivation of the 3D structure and motion is based primarily on a heuristic (rather than a veridical mathematical) analysis of the optic flow (Domini et al., 1997, 1998, 2002; Domini and Caudek, 1999; Domini and Braunstein, 1998).

*Michael Tarr*'s research focuses on behavioral and computational approaches to object recognition and visual perception (Tarr and Cheng, 2003; Tarr and Warren, 2002; Gauthier et al., 2002; Tarr and Gauthier, 2000; Tarr et al., 1998). He is interested in perceptual expertise as a model of plasticity in the visual system (and as an explanation for putatively “face-specific” phenomena), the degree to which surface properties such as color or luminance are intrinsic to object vision, and how different kinds of visual information are recruited for tasks such as object recognition or navigation. Earlier work investigated how observers recognize three-dimensional objects from two-dimensional views.

*Bill Warren*'s research focuses on the visual control of action - in particular, human locomotion and navigation (Duchon et al., 1998; Li and Warren, 2000; Warren et al., 2001; Duchon and Warren, 2002; Kearns et al., 2002). Using virtual reality techniques, his research team investigates problems such as the perception of optic flow, the visual control of steering and obstacle avoidance, and path integration and spatial knowledge used in navigation. This VR technology allows researchers to manipulate visual information during active walking, and to measure their ongoing behavior. The ultimate aim of this research is to understand how adaptive behavior emerges from the dynamic interaction of an organism and its environment.

## The Brain Science Program

While our focus here is primarily on researchers studying computational vision, it should be clear that many of us are inspired by biological vision or directly work on neural models of vision. The departments above are all part of a broader Brain Science Program where we have many other colleagues who study vision from other perspectives. At Brown, research into the cognitive, computational, and neural mechanisms of vision is diverse and interdisciplinary. We cite just a few examples below to give a flavor of this breadth.

*David Sheinberg* in the Department of Neuroscience studies the neural mechanisms of natural vision (Sheinberg and Logothetis, 1997, 2001, 2002). *Michael Paradiso*, also in Neuroscience, studies how neurons code visual information and how the communication and cooperation of these neurons underlies visual perception (Paradiso, 2002; Macevoy and Paradiso, 2001; Rossi and Paradiso, 1999). Both Sheinberg and Paradiso conduct perception experiments during which the activity of brain cells is recorded. *David Berson* studies retinal ganglion cells, how the different types of these cells respond to visual stimuli, and what information they send to the rest of the brain.

In the Department of Psychology *Leslie Welch* studies the human visual psychophysics of motion perception, binocular vision, and spatial vision. Her work also addresses issues of perceptual learning in the visual system (Festa and Welch, 1997; Welch et al., 1997; Matthews and Welch, 1997).

In the Department of Physics *Nathan Intrator* studies biological vision as well as image pattern recognition and classification (Blais et al., 1998; Edelman and Intrator, 2000; Intrator and Edelman, 1997). His work on early vision includes nonlinear feature extraction and dimensionality reduction (Intrator and Edelman, 1997). His work focuses on neural models of vision, receptive field formation in natural environments, learning in early visual cortex, and learning in network models of vision.

## Conclusion

We hope you enjoy this issue and will take the opportunity to explore the work of our colleagues in more depth. If you would like to learn more about computational vision at Brown please see the following web sites:

```
http://www.vision.brown.edu/
http://www.dam.brown.edu/ptg
http://www.cs.brown.edu
http://www.lems.brown.edu/vision
http://www.cog.brown.edu
http://www.brainscience.brown.edu
```

## Acknowledgments

This editorial contains text contributed by many of the above researchers. We thank then for their assistance and accept full responsibility for errors and omissions.

## References

Andrews, S. and Laidlaw, D.H. 2002. Toward a framework for assembling broken pottery vessels. In *Proceedings of the 18th National Conference on Artificial Intelligence*, pp. 945–946.

Andrews, S., Tsochantaridis, I., and Hofmann, T. 2003. Support vector machines for multiple-instance learning. In *Advances in Neural Information Processing Systems (NIPS*15)*, MIT Press (to appear).

Bakircioglu, M., Grenander, U., Khaneja, N., and Miller, M.I. 1998. Curve matching on brain surfaces using induced Frenet distance metrics. *Human Brain Mapping*, 6(5):329–331.

Balmelli, L., Taubin, G., and Bernardini, F. 2002 Space-optimized texture maps. In *Proceedings of Eurographics 2002*, Germany.

Barzohar, M. and Cooper, D.B. 1996. Automatic finding of main roads in aerial images by using geometric-stochastic models and estimation. *IEEE Trans. On PAMI*, 18(7):707–721.

Bernardini, F., Martin, I., Mittleman, J., Rushmeier, H., and Taubin, G. 2002. Building a digital model of Michelangelo's Florentine Pieta. *IEEE Computer Graphics & Applications*, January/February 2002.

Bienenstock, E. and Doursat, R. 1994. A shape-recognition model using dynamical links. *Network: Computation in Neural Systems*, 5:241–258.

Bienenstock, E., Geman, S., and Potter, D. 1997. Compositionality, MDL priors, and object recognition. In *Advances in Neural*

*Information Processing Systems 9*, M.C. Mozer, M.I. Jordan, and T. Petsche (Eds.), MIT Press, pp. 838–844.

Black, M.J. and Anandan, P. 1996. The robust estimation of multiple motions: Parametric and piecewise-smooth flow fields. *Computer Vision and Image Understanding, CVIU*, 63(1):75–104.

Black, M.J. and Fleet, D.J. 2000. Probabilistic detection and tracking of motion discontinuities. *International Journal of Computer Vision*, 38(3):231–245.

Black, M.J., Fleet, D.J., and Yacoob, Y. 2000. Robustly estimating changes in image appearance. *Computer Vision and Image Processing*, 78(1):8–31.

Black, M.J. and Yacoob, Y. 1997. Recognizing facial expressions in image sequences using local parameterized models of image motion. *Int. Journal of Computer Vision*, 25(1):23–48.

Black, M.J. and Jepson, A. 1998. EigenTracking: Robust matching and tracking of articulated objects using a view-based representation. *Int. Journal of Computer Vision*, 26(1):63–84.

Blais, B.S., Intrator, N., Shouval, H., and Cooper, L.N. 1998. Receptive field formation in natural scene environments: Comparison of single cell learning rules. *Neural Computation*. 10(7):1797–1813.

Blane, M., Lei, Z., Civi, H., and Cooper, D.B. 2000. The 3L algorithm for fitting implicit polynomial curves and surfaces to data. *IEEE Trans. On PAMI*, 22(3):298–313.

Cooper, D.B., Willis, A., Andrews, S., Baker, J., Cao, Y., Han, D., Kang, K., Wong, W., Leymarie, F.F., Orriols, X., Velipasalar, S., Vote, E.L., Joukowsky, M.S., Kimia, B.B., Laidlaw, D.H., and Mumford, D. 2001. Assembling virtual pots from 3D measurements of their fragments. In *Proc. Of the Virtual Reality Archaeology and Cultural heritage Symposium (VAST)*, Athens, Greece, Nov. 2001.

Cooper, D.B., Willis, A., Andrews, S., Baker, J., Cao, Y., Han, D., Kang, K., Kong, W., Leymarie, F., Orriols, X, Velipasalar, S., Vote, E., Joukowsky, M., Kimia, B.J., Laidlaw, D.H., and Mumford, D. 2002. Bayesian virtual pot-assembly from fragments as problems in perceptual-grouping and geometric-learning. In *Proceedings of ICPR*, vol. 3, pp. 30297–30302.

Domini, F. and Braunstein, M.L. 1998. Recovery of 3D structure from motion is neither Euclidean nor affine. *Journal of Experimental Psychology: Human Perception and Performance*, 24:1273–1295.

Domini, F. and Caudek, C. 1999. Perceiving surface slant from deformation of optic flow. *Journal of Experimental Psychology: Human Perception and Performance*, 25:426–444.

Domini, F., Caudek, C., and Proffitt, D.R. 1997. Misperceptions of angular velocities influence the perception of rigidity in the kinetic depth effect. *Journal of Experimental Psychology: Human Perception and Performance*, 23:1111–1129.

Domini, F., Caudek, C., and Richman, S. 1998. Distortions of depth-order relations and parallelism in structure from motion. *Perception and Psychophysics*, 60:1164–1174.

Domini, F., Vuong, Q., and Caudek, C. 2002. Temporal integration in structure from motion. *Journal of Experimental Psychology: Human Perception and Performance*, 28:816–838.

Duchon, A.P. and Warren, W.H. 2002. A visual equalization strategy for locomotor control: Of honeybees, humans, and robots. *Psychological Science*, 13:272–278.

Duchon, A., Warren, W.H., and Kaelbling, L.P. 1998. Ecological robotics. *Special issue of Adaptive Behavior*, 6:473–507.

Edelman, S. and Intrator, N. 2000. Coarse coding of shape fragments + retinotopy representation of structure. *Spatial Vision*, 13:255–264.

Festa, E.K. and Welch, L. 1997. Recruitment mechanisms in speed and fine-direction discrimination tasks. *Vision Research*, 37(22):3129–3144.

Gauthier, I., Hayward, W.G., Tarr, M.J., Anderson, A., Skudlarski, P., and Gore, J.C. 2002. BOLD activity during mental rotation and viewpoint-dependent object recognition. *Neuron*, 34(1):161–171.

Geman, S. and Geman, D. 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE-PAMI*, 6:721–741.

Geman, D., Geman, S., Graffigne, C., and Dong, P. 1990. Boundary detection by constrained optimization. *IEEE-PAMI*, 12:609–628.

Geman, D., Geman, S., and McClure, D.E. 1992. A nonlinear filter for film restoration and other problems in image processing. *CVGIP: Graphical Models and Image Processing*, 54:281–289.

Geman, S., Bienenstock, E., and Doursat, R. 1992. Neural networks and the bias/variance dilemma. *Neural Computation*, 4:1–58.

Geman, S., Manbeck, K., and McClure, D.E. 1993. A comprehensive statistical model for single-photon emission tomography. In *Markov Random Fields: Theory and Applications* R. Chellappa and A. Jain (Eds.) Academic Press, Boston, pp. 93–130.

Geman, S., Manbeck K., and McClure, D.E. 1993. A comprehensive statistical model for single photon emission computed tomography. In *Markov Random Fields: Theory and Application*, R. Chellappa and A. Jain (Eds.), Academic Press (Harcourt Brace Jovanovich), pp. 93–130.

Geman S. and McClure, D.E. 1987. Statistical methods for tomographic image reconstruction. *Bulletin of the International Statistical Institute*, 52.

Geman, S., Potter, D.F., and Chi, Z. 2002. Composition systems. *Quarterly of Applied Mathematics*, LX:707–736.

Gidas, B. 1989. A renormalization group approach to image processing problems. *IEEE Transactions PAMI*, 11(2):164–180.

Gidas, B., Almeida, M., and Robertson, C. 2000. Tracking of moving objects in cluttered enviroments via Monte Carlo filter. In *Proceedings of the International Conference on Pattern Recognition (ICPR2000)*, Barcelona, Spain, pp. 175–179.

Gidas, B. and Geman, D. 1991. Image analysis and computer vision. *Spatial Statistics and Image Processing*, National Research Council, National Academy Press, pp. 1–36.

Gidas, B., Gomes, F. C., and Robertson, C. 2002. Model-based tracking of moving objects in cluttered enviroments. *Quarterly of Applied Mathematics*, LX(4):737–771.

Gidas, B. and Mumford, D. 2001. Stochastic models for generic images. *Quarterly of Applied Mathematics*, LIX(1):85–111.

Grenander, U. and Miller, M.I. 1998. Computational anatomy: An emerging discipline. *Quarterly of Applied Mathematics*, LVI(4):617–694.

Grenander, U., Miller, M.I., and Srivastava, A. 1998. Hilbert-Schmidt lower bounds for estimators on matrix Lie groups for ATR. *IEEE Trans. PAMI*, 20(8):790–802.

Grenander, U. and Srivastava, A. 2001. Probably models for clutter in natural images. *IEEE Trans. PAMI*, 23(4):424–429.

Grimm, C., Crisco, J.J., and Laidlaw, D.H. 2002. Fitting manifold surfaces to 3D point clouds. *Journal of Biomechanical Engineering*, 124(1):136–140.

Heller, A. and Mundy, J.L. 1991. Benchmark evaluation of a model-based object recognition system. In *Proc. 3rd International Conference on Computer Vision*, 1990, Reprinted in *Computer Vision: Advances and Applications*, R. Kasturi and R. Jain (Eds.) IEEE Computer Society Press.

Hofmann, T. 2001. Unsupervised learning by probabilistic latent semantic analysis. *Machine Learning Journal*, 42(1):177–196.

Hofmann, T., Puzicha, J., and Buhmann, J.M. 1998. Unsupervised texture segmentation in a deterministic annealing framework. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 20(8):803–818.

Intrator, N. and Edelman, S. 1997. Learning low dimensional representations of visual object with extensive use of prior knowledge. *Network*, 8(3):283–296.

Kang, K., Tarel, J.-P., Fishman, R., and Cooper, D.B. 2001. A linear dual-space approach to 3D surface reconstruction from occluding contours using algebraic surfaces. In *Proc. IEEE Intl. Conf. On Computer Vision, Vancouver*, Canada, pp. 198–204.

Kearns, M., Warren, W.H., Duchon, A., and Tarr, M. 2002. Path integration from optic flow and body senses in a homing task. *Perception*, 31:349–374.

Kimia, B.B. To Appear. On the role of medial geometry in human vision. *Journal of Physiology*, 2003.

Kimia, B.B., Tannenbaum, A.R., and Zucker, S.W. 1992. On the evolution of curves via a function of curvature, I: The classical case. *JMAA*, 163(2):438–458.

Kimia, B.B., Tannenbaum, A.R., and Zucker, S.W. 1995. Shapes, shocks, and deformations, I: The components of shape and the reaction-diffusion space. *IJCV*, 15(3):189–224.

Kunsch, H., Geman, S., and Kehagias, A. 1995. Hidden Markov random fields. *Ananals of Applied Probability*, 5:577–602.

Kutliroff, G. 2002. Approximation in an Adaptive Cosine Basis and Its Application to Image Compression, Ph.D. Thesis, Division of Applied Mathematics, Brown University, May 2002. Available at www.dam.brown.edu/people/gershom.

Laidlaw, D.H., Fleischer, K.W., and Barr, A.H. 1998. Partial-volume Bayesian classification of material mixtures in MR volume data using voxel histograms. *IEEE Transactions on Medical Imaging*, 17(1):74–86.

Lee, T. S. and Mumford, D., to appear. Hierarchical Bayesian inference in the visual cortex. *J. of the Optical Society of America* (submitted).

Lee, A.B., Pedersen, K.S., and Mumford, D. 2003. The nonlinear statistics of high-contrast patches in natural images. *Int. J. Comp. Vision*. 54(1/2/3):83–103.

Li, L. and Warren, W.H. 2000. Perception of heading during rotation: Sufficiency of dense motion parallax and reference objects. *Vision Research*, 40:3873–3894.

Macevoy, S.P. and Paradiso, M.A. 2001. Lightness constancy in primary visual cortex. *Proceedings of the National Academy of Science*, 98:8827–8831.

Matthews, N. and Welch, L. 1997. The effect of inducer polarity and contrast on the perception of illusory figures. *Perception*, 26:1431–1443.

Miller, M., Christensen, G., Amit, Y., and Grenander, U. 1993. Mathematical textbook of deformable neuroanatomies. *Proc. of the National Academy of Sciences*, 90(24):11944–11948.

Mumford, D. 2002. Pattern theory: The mathematics of perception. In *Proceedings of ICM 2002*, vol. 1, Beijing.

Mumford, D. 2003. The shape of objects in two and three dimensions, Gibbs Lecture. *Notices of the Amer. Math. Soc.*, to appear.

Mumford, D. and Gidas, B. 2001. Stochastic models for generic images. *Quarterly Appl. Math.*, 59:85–111.

Mundy, J.L. 1995. Object recognition: The search for representation. Lecture Notes in Computer Science, vol. 994.

Mundy, J.L. 1998. Object recognition based on geometry: Progress over three decades. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences, of the Royal Society*, 356:1213–1231.

Mundy, J.L. and Zisserman, A. 1992. *Applications of Geometric Invariance in Computer Vision*. MIT Press.

Paradiso, M.A. 2002. Neuronal and perceptual correspondence in primary visual cortex. *Current Opinion in Neurobiology*, 12:155–161.

Puzicha, J., Hofmann, T., and Buhmann, J.M. 1999. Histogram clustering for unsupervised segmentation and image retrieval. *Pattern Recognition Letters*, 20:899–909.

Rossi, A.F. and Paradiso, M.A. 1999. Neural correlates of brightness in the responses of neurons in the retina, LGN, and primary visual cortex. *Journal of Neuroscience*, 19:6145–6156.

Rothwell, C.A., Forsyth, D.A., Zisserman, A., and Mundy, J.L. 1993. Extracting projective structure from single perspective views of 3-d point sets. In *Proc. 4th International Joint Conference on Computer Vision*.

Sebastian, T., Crisco, J., and Kimia, B. 2003. Segmentation of carpal bones for accurate measurement of 3D in vivo carpal kinematics. *Medical Image Analysis*, 7(1):21–45.

Sebastian, T., Klein, P., and Kimia, B.B. 2003. Recognition of shapes by editing their shock graphs. *PAMI*. Under minor revision.

Sheinberg, D.L., Logothetis, N.K. 1997. The role of temporal cortical areas in perceptual organization. *Proceedings of the National Academy of Sciences*, 94:3408–3413.

Sheinberg, D.L. and Logothetis, N.K. 2001. Noticing familiar objects in real world scenes: The role of temporal cortical neurons in natural vision. *Journal of Neuroscience*, 21:1340–1350.

Sheinberg, D.L. and Logothetis, N.K. 2002. Perceptual learning and the development of complex visual representations in temporal cortical neurons. In *Perceptual Learning*, M. Fahle and T. Poggio (Eds.), MIT Press, 95–124.

Siddiqi, K., Kimia, B.B., Tannenbaunm, A.R., and Zucker, S.W. 2001. On the psychophysics of the shape triangle. *Vision Research*, 41(9):1153–1178.

Tarel, J.-P. and Cooper, D.B. 2000. The complex representation of algebraic curves and its simple exploitation for pose estimation and invariant recognition. *IEEE Trans. On Pattern. Anal. and Machine Learning*, 22(7):663–674.

Tarr, M.J. and Cheng, Y.D. 2003 Learning to see faces and objects. *TRENDS in Cognitive Sciences*, 7(1):23–30.

Tarr, M.J. and Gauthier, I. 2000. FFA: A Flexible Fusiform Area for subordinate-level visual processing automatized by expertise. *Nature Neuroscience*, 3(8):764–769.

Tarr, M.J. and Warren, W.H. 2002. Virtual reality in behavioral neuroscience and beyond. *Nature Neuroscience Supplement*, 5:1089–1092.

Tarr, M.J., Williams, P., Hayward, W.G., and Gauthier, I. 1998. Three-dimensional object recognition is viewpoint-dependent. *Nature Neuroscience*, 1(4):275–277.

Taubin, G. 2002. BLIC: Bi-level isosurface compression. In *Proceedings of IEEE Visualization 2002*, Boston.

Taubin, G., Gueziec, A., Horn, W., and Lazarus, F. 1998. Progressive forest split compression. *Siggraph'98*, Orlando, FL.

Taubin, G. and Rossignac, J. 1998. Geometric compression through topological surgery. *ACM Transactions on Graphics*, 17(2).

van Dam, A., Laidlaw, D.H., and Simpson, R.M. 2002. Experiments in immersive virtual reality for scientific visualization. *Computers and Graphics*, 26(4):535–555.

von der Malsburg, C. and Bienenstock, E. 1987. A neural network for invariant pattern recognition. *Europhysics Letters*, 4:121–126.

Warren, W.H. In press. Optic flow. In *The Visual Neurosciences*, L. Chalupa and J. Werner (Eds.), Cambridge, MA: MIT Press.

Warren, W.H., Kay, B.A., Duchon, A.P., Zosh, W., and Sahuc, S. 2001. Optic flow is used to control human walking. *Nature Neuroscience*, 4:213–216.

Welch, L., MacLeod, D.I.A., and McKee, S.P. 1997. Motion interference: Perturbing perceived direction. *Vision Research*, 37(19):2725–2736.