Introduction to the Special Issue on Learning in Computer Vision and Pattern Recognition

THE GOAL OF computer vision and pattern recognition (CVPR) research is to provide computers with human-like perception capabilities so that they can sense the environment, understand the sensed data, identify patterns, take appropriate actions, and learn from this experience in order to enhance future performance. The field has evolved from the application of classical pattern recognition and image processing techniques to advanced applications of image understanding, dynamic scene analysis, model-based vision, knowledge-based vision and systems that exhibit learning capability.

Over the years, there has been an increased demand for CVPR systems to address "real-world" applications, such as autonomous navigation, target recognition, manufacturing, photointerpretation, remote sensing, situation awareness, image/video database management, etc. This requires that the vision techniques be robust and flexible to optimize performance in diverse scenarios encountered in a given application.

Past research in applying learning techniques to CVPR problems has been limited [1], [2]. Some of the reasons for this were the lack of understanding and availability of tools for low-level image analysis, interdisciplinary nature of learning in CVPR research and the lack of machine learning [14]–[16] and statistical learning tools [12], [13]. However, in the last decade, some progress has been achieved toward these problems. Solving the signal-to-symbol transition problem remains one of the key challenges in the application of symbolic learning to vision. Learning requires a lot of data and speed for its practical use.

The field of machine and statistical learning is driven by the idea that computer algorithms and systems can improve their own performance with time. Vision provides interesting and challenging problems and a rich environment to advance the state-of-the art in learning. There is a strong potential in learning technology to contribute to the development of flexible and robust vision algorithms, thus improving the performance of vision systems for practical use. Learning-based vision systems are expected to provide a higher level of competence and greater generality. Learning may allow using the experience gained in creating a vision system for one application domain to a vision system for another domain by acquiring and maintaining knowledge.

Thus, an innovative integration of learning and CVPR techniques has the promise of advancing the field which will contribute to better understanding of complex images of real-world dynamic scenes. There is another benefit of incorporating a learning paradigm in the computational CVPR framework. To mature the laboratory-grown vision systems into real-world working systems, it is necessary to evaluate the performance characteristics of these systems using a variety of real, calibrated data. Learning offers this evaluation tool, since no learning can take place without appropriate supervised or unsupervised mode of evaluation of the results.

Since the days of the first NSF/DARPA Workshop [1], the first AAAI workshop on Learning in Computer Vision [2], and the first special section on Learning in Computer Vision published in the IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE [3] we have come a long way. There have been many developments in this area. A number of books related with learning in CVPR [4]–[11], [17], [18] and other popular books in learning [12]–[16] have been published. Several workshops and conferences or special sessions are being held in major conferences in CVPR and learning fields [19]–[26]. This shows that there is a growing interest in the research and application communities on this topic.

Theoretical and practical advances are being made in the field of computer vision and pattern recognition by new techniques and processes of learning, representation, and adaptation. Examples are new regression and classification techniques (e.g., boosting, bagging, margin classifiers), multiclassifier systems, Bayesian learning, interactive, continuous and active learning, reinforcement learning, graphical models, etc. Learning offers effective methods for automating the model/concept acquisition and updating processes, adapting task parameters and representations, and using experience for generating, verifying, and modifying hypotheses. Expanding this list of CVPR problems, we find the current topics of research interest and research on multilevel vision systems for segmentation, feature extraction and object recognition; new learning paradigms including concept learning, learning by examples, task-level learning; learning rules, relations, features, discriminant functions, and evaluation strategies; learning and refining visual models; indexing and recognition strategies; integration of vision modules; learning shape representation and surface reconstruction strategies; self-organizing algorithms for clustering and pattern learning; biologically motivated modeling of vision systems that learn; parameter adaptation and self-calibration of vision systems; evolutionary approaches to learning and synthesis of recognition systems; neural network learning for solving significant vision problems and applications to real-world problems. As an eventual goal, machine learning may provide the necessary tools for synthesizing vision algorithms. Some initial attempts have been made in that direction [17], [18].

The effective usage of learning technology in real-world computer vision problems requires understanding the domain of application, abstraction of a learning problem (the what problem) from a given CVPR task and the selection of appropriate representations for the learnable (input) and learned (internal) entities of the system. Since learning in vision is a new area of

Digital Object Identifier 10.1109/TSMCB.2005.847940

1083-4419/$20.00 © 2005 IEEE

research there are many unexplored issues and there are potentially many different ways in which learning can be applied to solve vision problems and to optimize the resources needed by a vision system. However, learning needs to be carefully applied to selected problems where it makes sense.

A learning system has to clearly demonstrate and answer the questions like what is being learned, from what is being learned, how it is learned, what data is used to learn, how to represent what has been learned, how well and how efficient is the learning, how the learning improves the performance of the vision system over the usual static (or baseline) counterparts and what are the evaluation criteria for the task at hand. Experimental details are essential for learning experiments. They include scientific experimental design methodology for training/testing, parametric studies, and measures of learning and performance improvement with experience. Experiments demonstrating scalability of learning-based vision systems are also very important.

Papers in This Special Issue

The goal of this special issue is to provide the reader with samples of recent developments that use adaptive and learning approaches in computer vision and pattern recognition. The special issue has 11 regular papers and seven correspondence papers. These papers use a variety of techniques for adaptation and learning. All the papers have gone through several rounds of reviewing according to the guidelines and standards of IEEE Transactions. The papers cover a broad area of theory and applications of computer vision and pattern recognition. We hope that the special issue will give a boost to the research in learning in computer vision and pattern recognition and their applications.

Regular Papers

The first paper, by Makris and Ellis, is on learning semantic scene models from observing activities over long periods of time in video streams of data. It addresses the problem that in visual surveillance it is desired to not only detect and track objects, but also understand the activity in the scene. The activity-based semantics learned in this paper aims at generic activities. The input to the system is motion-tracking data. These trajectories are used to identify zones in the camera field of view where objects appear and disappear, paths they follow, junctions where they might change their routes, and stop areas. The two levels of representations that are used include a topographical representation and a topological representation. They focus on spatial and probabilistic nature of the model, respectively. A Gaussian mixture model is used to encode spatial and probabilistic characteristics of various zones. The learning is performed at multiple steps using the Expectation-Maximization (EM) algorithm. The results are shown using real-world data consisting of a large number of trajectories.

The next three papers (second, third, and fourth papers in this special issue) take inspiration from biological evolution and two of them use evolutionary computation for object recognition. The second paper, by Krawiec and Bhanu, presents an evolutionary approach for visual learning. It uses the idea of cooperative coevolution to handle the complexity of the object recognition task. It uses linear genetic programming, a hybrid of genetic algorithms (GAs) and genetic programming (GP), as an evolutionary approach to represent feature extraction procedures. The procedures are learned using elementary image processing operations and filters, and commonly used features in pattern recognition. A variety of techniques and ideas are used for binary classification problems, multiple class recognition problems, scaling the number of classes, and recognizing object variants. Synthetic aperture radar (SAR) imagery is used to demonstrate the results.

The third paper, by Schneider *et al.*, presents a biologically inspired hierarchical vision model for 3-D object recognition. Object recognition has been a very challenging task for neural networks since the networks are large and the search space for computing the network parameters is of very high dimensionality. Evolutionary optimization techniques (direct coding and indirect coding) are used to determine optimal higher-order complex features and nonlinearities of the visual hierarchy. The authors propose an efficient and indirect coding of a neural vision network and propose local unsupervised learning rules based on sparse coding concepts. The optimized networks are applied to a classification task of real-world images that are available in the public domain. Various experiments are performed to demonstrate the capabilities of the system.

Learning by imitations has been important in the evolution of species. The fourth paper, by Lopes and Santos-Victor, is on visual learning by imitation with motor representations. It presents a biologically inspired architecture for arm-hand gesture imitation and recognition. Two different forms of imitations (action level and program level) are distinguished. Action-level (mimicking) imitation consists of replicating the gestures of a demonstrator. Program-level (gesture) imitation involves recognizing the performed gesture so that the learner can produce its own interpretation of the gesture or action effect. A Bayesian framework is used for program level imitation. The robotic system consists of an anthropomorphic arm/hand that is equipped with a single camera. The robot arm/hand is simulated but the real data from a camera is used in extensive experiments.

Concept learning in content-based retrieval systems is a difficult task. The fifth paper, on active concept learning in image databases, by Dong and Bhanu, describes a mixture model approach to concept learning. It addresses two critical issues facing image databases: 1) changing (image insertion/removal) nature of a database and 2) user queries. To achieve concept learning, a novel user directed, semi-supervised expectation-maximization algorithm for parameter estimation is proposed. The model consistency is determined through Bayesian analysis. The proposed concept learning algorithm is able to deal with image insertion as well as any inconsistency in relevance feedback. Experiments on Corel database images are provided to show the efficacy of the proposed technique.

The next three papers deal with human faces. The sixth paper, by Waring and Liu, presents a face detection method that uses the spectral histogram representation and the support vector machines as a margin classifier. Spectral histogram, computed in each image window, is a feature vector that consists of the histogram of filtered images. Thirty three filters used in the paper are: four gradient filters, five Laplacian of Gaussian filters and

24 Gabor filters. It is shown that this representation groups face images together. The results are compared with other recent methods on two commonly used datasets for face detection.

The seventh paper, by Guo and Dyer, presents a linear programming-based technique that handles the problem of feature selection and classifier training simultaneously, especially for the case when the size of the training set is small. For multiclass recognition, the feature selection is performed pairwise. The results are shown for face expression recognition, which has seven classes (neutral, happy, sad, surprise, anger, disgust, and fear) and each class has only two to four images for each facial expression. The features are obtained using Gabor wavelets and are selected by hand. The results are compared with other classifiers, including support vector machine and Adaboost.

Subspace analysis is an important approach to learning low-dimensional representations for classification. The eighth paper, on kernel pooled local subspaces for classification, by Zhang *et al.*, investigates a dimensionality reduction technique that pools locally discriminant information and extend it to the nonlinear case using the kernel trick. They compare their technique against competing subspace methods: kernel principal component analysis (KPCA) and Generalized Fisher discriminant analysis (GDA) in a number of classification problems. Their evaluation is based on the classification performance of the nearest neighbor rule with each subspace representation. The experimental results demonstrate the effectiveness and performance superiority of the proposed kernel pooled subspace method over the competing methods such as KPCA and GDA in some classification problems.

The next two papers use self-organizing maps as a learning technique. The self-organizing maps (SOMs) implement two important operations: topology-preserving mapping and vector quantization. In most of the current work, based on edit distance concept in pattern recognition, the edit costs are derived manually in a task specific manner. The ninth paper, by Neuhaus and Bunke, addresses the problem of learning graph edit distance cost functions for numerically labeled graphs from a given sample set of graphs. The learning approach is based on self-organizing maps (SOM) which models the distribution of the node and edge labels occurring in a population of graphs. The actual edit costs are obtained from a distance measure for labels that is derived with respect to the distribution encoded in the SOM. The results are shown for two different applications involving line drawing graphs of alphabetic letters dataset and graphs representing diatoms.

The tenth paper, by Xu *et al.*, proposes an on-line hierarchical SOMs (H-SOMs) with faster learning. A new learning rule is proposed that delivers the efficiency and topology preservation. The learning rule is superior to other structures of SOMs. The computational complexity of H-SOMs is O(logN) rather than O(N) as in base SOMs. Experimental results reported in the paper demonstrate that the reconstruction performance of H-SOMs is comparable to full-search SOMs, while its computation is much faster. In addition, H-SOMs generate a hierarchical mapping of code vectors and support progressive transmission and decoding property.

The last regular paper, by Joshi *et al.*, addresses the problem of how to create a super-resolution image of a scene from images of the scene taken at different camera zooms. It attempts to build an image of the entire scene at a resolution corresponding to the most zoomed image that is observed. A learning-based approach is proposed since it will be difficult to capture the richness of real world images analytically. Under the assumption that the entire scene is statistically homogeneous, the parameters of the super-resolved image are learned from the most zoomed observation and they are used to estimate the super-resolution image for the least zoomed entire scene. Homogeneous Markov random field model and simultaneous autoregressive models are used. Results are shown on both synthetic and real textured images.

Correspondence Papers

The first two papers deal with the computational aspects of evolutionary learning in computer vision and pattern recognition. The first paper, by Lin and Bhanu, develops a genetic programming (GP) based approach to synthesize composite operators and features from a combination of primitive ones for object detection. While human experts tend to focus on conventional mixtures of primitive operations and features, GP is capable of exploiting unconventional combinations that in some cases yield exceptionally good results. To make GP efficient, while not being too restrictive in search, Lin and Bhanu propose smart crossovers and mutations as well as a public library to identify and keep the effective components of composite operators. In addition, a new fitness function based on minimum description length is proposed to deal with the well-known code bloat problem of GP. Experiments on real images demonstrate that the proposed GP algorithm is quite effective.

The second paper, by Cagnoni *et al.*, presents an approach based on two parallel evolutionary techniques (cellular programming and submachine-code genetic programming) for the design of an ensemble of binary classifiers for recognition. Cellular programming is an evolutionary computational technique in which nonuniform cellular automata (different local computations for every cell) are evolved. Submachine code genetic programming exploits the intrinsic parallelism of bitwise instructions of sequential CPUs and can be effectively run on traditional computer architectures. Different fitness functions are proposed which emphasize the computational accuracy or the efficiency. The experimental results are presented for low-resolution (13×8 pixels) digit recognition application using a set of binary classifiers.

Kernel biased discriminant analysis (KBDA) is a learning method for content-based image retrieval. Crucial to the success of KBDA is the selection of kernel parameters. In the third paper, on optimizing learning performance of KBDA, Wang *et al.*, propose a novel criterion for optimal parameter selection to maximize retrieval performance. Under this criterion, the optimal kernel parameters are the ones that put positive images in a tight cluster and drive negative images away from the positive ones as far as possible in a kernel-induced space. Experimental results show that the proposed criterion can effectively optimize kernel parameters, thus the learning performance of KBDA.

The next paper, by Luo *et al.*, presents an approach to improve semantic scene classification performance by using transformations in the image space to increase the number of exemplars.

The approach, called image-transform bootstrapping, uses geometric transformations to minimize the impact of undesirable foreground objects and color transformations to minimize the effects of color variability. The transforms are based on variations in the nature (lighting changes) and the picture-taking process (zoom, pan). Three schemes (mirroring, cropping and color shifting) are used to increase the number of training and testing images. Multiple transformed versions of either training images or testing images or both are used in conjunction with a single classifier to achieve bootstrapping. The technique has been applied for sunset detection, outdoor scene classification and image orientation detection.

The fifth paper, by Draper *et al.*, addresses the problem associated with the EM algorithm in high dimensions, especially when the number of data samples is smaller than the feature dimensions. As compared to the conventional EM algorithm, the paper combines EM with principal component analysis (PCA) and fits the data at a dimension as high as possible for more precise estimation. It uses eigenvalues and eigenvectors from PCA decompositions of weighted sample covariance matrices to represent high dimensional Gaussian distributions. It proposes ways to get soft clustering, defines Gaussian probability density function and estimates parameters. It shows experimental results on both synthetic data and real images.

The sixth paper, by Chen and Meer, presents a technique for robust fusion of uncertain information where the measurements and their uncertainty described by the covariance matrices are given. The characteristics of an unknown number of information sources are obtained by minimizing the sum of Mahalanobis distances from the measurements to those particular sources. Multivariate kernels are used to compute the sample point density estimate of the measurements. Adaptive mean shift locates the modes of the multivariate density distributions that are related to the information sources. Examples from two computer vision applications (range image segmentation and multiple affine transforms estimation) are given to demonstrate the effectiveness of the technique.

The last correspondence paper, by Yu *et al.*, presents a recursive two-step approach based on Kalman filtering to recover structure and motion from image sequences. The first step estimates the pose of an object using an extended Kalman filter and the second step uses a set of extended Kalaman filters for the refinement of positions of the model features in 3-D space. The performance of the algorithm is compared on both synthetic and real data. It is also compared with other recent approaches.

The Future

The field of machine learning in computer vision and pattern recognition (CVPR) is just emerging. The number of CVPR systems that incorporate some learning component is expected to increase as vision and pattern recognition researchers embrace the need for adaptation and learning in their systems. There will be significant advances as the field evolves and there is increased interaction between machine learning and computer vision, and pattern recognition researchers. Some of the questions of interest from the artificial intelligence (AI) perspective include the role of learning and high-level knowledge in Computer Vision, foundations for flexible/robust intelligent vision systems for real-world dynamic scene understanding involving a network of imaging/nonimaging sensors. The statement that was made in 1994 [3] is still true in a strong sense that it is probably fair to claim that learning represents the next challenging frontier for computer vision and pattern recognition research.

ACKNOWLEDGMENT

We thank the authors who submitted and contributed their work to this special issue and worked under a tight schedule. We thank the referees who spent their valuable time in reviewing the manuscripts. We greatly appreciate the promptness and useful suggestions that were provided, by Larry Hall, Editor-in-Chief (EIC) of SMCB, during various stages of the special issue. Thank you very much, Larry. We would like to thank Diane Cook, Co-EIC of SMCB for her help in handling all the final papers and getting all the material related to the special issue out to IEEE.

BIR BHANU, *Guest Editor*
University of California at Riverside
Riverside, CA 92521 USA
bhanu@cris.ucr.edu

JING PENG, *Guest Editor*
Tulane University
New Orleans, LA 70118 USA
jp@eecs.tulane.edu

TOM HUANG, *Guest Editor*
University of Illinois
Urbana, IL 61801 USA
huang@ifp.uiuc.edu

BRUCE A. DRAPER, *Guest Editor*
Colorado State University
Fort Collins, CO 80521 USA
draper@cs.colostate.edu

REFERENCES

[1] R. S. Michalski, A. Rosenfeld, and Y. Aloimonos, "Machine Vision and Learning,", Harpers Ferry, WV, A Report on the NSF/ARPA Workshop on Machine Learning and Vision, Oct. 15–17, 1992.

[2] K. Boyer, L. Hall, P. Langley, B. Bhanu, and B. Draper, "Report of the AAAI fall symposium on machine learning and computer vision: What, why, and how?," in *Proc. ARPA Image Understanding Workshop*, Monterey, CA, Nov. 13–16, 1994, pp. 727–731.

[3] B. Bhanu and T Poggio, Eds., *IEEE Trans. Pattern Anal. Mach. Intell., Special Section on Learning in Computer Vision*, Aug. 1994, vol. 16, no. 9, pp. 865–919.

[4] S. K. Nayar and T. Poggio, Eds., *Early Visual Learning*: Oxford University Press, 1996.

[5] K. Ikeuchi and M. Veloso, Eds., *Symbolic Visual Learning*. Oxford, U.K.: Oxford Univ. Press, 1997.

[6] T. Caelli and W. F. Bischof, *Machine Learning and Image Interpretation*. New York: Plenum, 1997.

[7] B. D. Ripley, *Pattern Recognition and Neural Networks*. Cambridge, U.K.: Cambridge Univ. Press, 1996.

[8] C. M. Bishop, *Neural Networks for Pattern Recognition*. Oxford, U.K.: Clarendon, 1996.

[9] B. Bhanu and S. Lee, *Genetic Learning for Adaptive Computer Vision*. Norwell, MA: Kluwer, 1994.

[10] S. Haykin, *Neural Networks—A Comprhensive Foundation*. New York: IEEE Press, 1994.

[11] Y. Anzai, *Pattern Recognition and Machine Learning*. New York: Academic, 1992.
[12] V. N. Vapnik, *The Nature of Statistical Learning Theory*. New York: Springer, 2000.
[13] T. Hastie, R. Tibshirani, and J. H. Friedman, *The Elements of Statistical Learning*. New York: Springer, 2001.
[14] T. Mitchell, *Machine Learning*. New York: WCB/McGraw-Hill, 1997.
[15] P. Langley, *Elements of Machine Learning*. San Mateo, CA: Morgan Kaufmann, 1996.
[16] R. S. Sutton and A. G. Barto, *Reinforcement Learning—An Introduction*. Cambridge, MA: MIT Press, 1998.
[17] B. Bhanu, Y. Lin, and K. Krawiec, *Evolutionary Synthesis of Pattern Recognition Systems*: Springer, 2005.
[18] B. Bhanu and S. Das, *Computational Learning for Adaptive Computer Vision*. New York: Springer, 2005.
[19] *IEEE Workshops on Learning in Computer Vision and Pattern Recognition (2003–2005)*.
[20] *AAAI Workshop on Learning in Computer Vision*, July 2005.
[21] *Genetic and Evolutionary Computation Conference (GECCO)*, Washington, DC, June 25–29, 2005.
[22] *Evolutionary Computation in Image Analysis and Signal Processing*, Lausanne, Switzerland, Mar. 30–Apr. 1, 2005.
[23] *ECCV Workshop on Statistical Learning in Computer Vision*, Prague, Czech Republic, May 15, 2004.
[24] *Neural Information Processing Systems (NIPS) Conference*, Vancouver, BC, Canada, Dec. 6–8, 2005.
[25] *International Conference on Machine Learning (ICML)*, Bonn, Germany, Aug. 7–11, 2005.
[26] *IEEE Workshops on Machine Learning for Signal Processing*, Mystic, CT, Sept. 28–30, 2005.

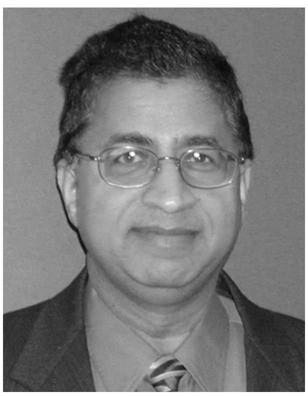

Bir Bhanu (S'72–M'82–SM'87–F'95) received the S.M. and E.E. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, the Ph.D. degree in electrical engineering from the Image Processing Institute, University of Southern California, Los Angeles, and the M.B.A. degree from the University of California, Irvine.

Dr. Bhanu was the founding Professor of electrical engineering and served its first Chair at the University of California, Riverside (UCR). He has been the Cooperative Professor of Computer Science and Engineering and Director of Visualization and Intelligent Systems Laboratory (VISLab) since 1991. He has also served as the founding Director of an interdisciplinary Center for Research in Intelligent Systems (CRIS) at UCR since 1998. Previously, he was a Senior Honeywell Fellow at Honeywell Inc., Minneapolis, MN. He was on the faculty of the Department of Computer Science, University of Utah, Salt Lake City, and has worked at Ford Aerospace and Communications Corporation, Newport Beach, CA; INRIA, Rocquencourt, France; and IBM San Jose Research Laboratory, San Jose, CA. He has been the principal investigator of various programs for the Defense Advanced Research Projects Agency, the National Aeronautics and Space Administration, the National Science Foundation, the Air Force Office of Scientific Research, the Army Research Office, and other agencies and industries in the areas of learning and vision, image understanding, pattern recognition, target recognition, biometrics, navigation, image databases, and machine vision applications. He is the co-author of the books *Evolutionary Synthesis of Pattern Recognition Systems* (New York: Springer, 2005), *Computational Algorithms for Fingerprint Recognition* (Boston, MA: Kluwer, 2004), *Genetic Learning for Adaptive Image Segmentation* (Kluwer, 1994), and *Qualitative Motion Understanding* (Kluwer, 1992), and the co-editor of the book *Computer Vision Beyond the Visible Spectrum* (New York: Springer, 2004).

Dr. Bhanu received two outstanding paper awards from the Pattern Recognition Society and has received industrial and university awards for research excellence, outstanding contributions, and team efforts. He has been on the editorial board of various journals and has edited special issues of several IEEE transactions and other journals. He holds 11 U.S. and international patents and over 230 reviewed technical publications in the areas of his interest. He was General Chair for IEEE Workshops on Applications of Computer Vision, Chair for the DARPA Image Understanding Workshop, General Chair for the IEEE Conference on Computer Vision and Pattern Recognition, and Program Chair for the IEEE Workshops on Computer Vision Beyond the Visible Spectrum. He is a Fellow of the American Association for the Advancement of Science, the International Association of Pattern Recognition, and the International Society for Optical Engineering.

Jing Peng (M'04) received the B.S. degree in computer science from the Beijing Institute of Aeronautics and Astronautics, Beijing, China, the M.A. degree in computer science from Brandeis University, Waltham, MA, and the Ph.D. degree in computer science from Northeastern University, Boston, MA in 1994.

Since 2001, he has been an Assistant Professor of electrical engineering and computer science at Tulane University, New Orleans, LA. Previously, he has been on the faculty of computer science at Oklahoma State University, Stillwater. He was a research scientist with the Center for Research in Intelligent Systems, University of California, Riverside. His research interests include machine learning, classification, image databases, and learning and vision applications.

Dr. Peng has served as the Program Co-Chair for the IEEE Workshop in Computer Vision and Pattern Recognition.

Thomas Huang (S'61–M'63–SM'76–F'79) received the Sc.D. degree from the Massachusetts Institute of Technology, Cambridge, in 1963.

He is William L. Everitt Distinguished Professor in the Department of Electrical and Computer Engineering and the Coordinated Science Lab (CSL), University of Illinois at Urbana Champaign (UIUC), and a full-time faculty member in the Beckman Institute Image Formation and Processing and Artificial Intelligence groups. His professional interests are computer vision, image compression and enhancement, pattern recognition, and multimodal signal processing. His research support includes the NSF, DOD, UIUC Research Board, and a number of industrial firms.

Dr. Huang honors include Member, National Academy of Engineering, Foreign Member, Chinese Academy of Engineering, Foreign Member, Chinese Academy of Sciences; IEEE Jack S. Kilby Signal Processing Medal (2000) (co-recipient with A. Netravali); International Association of Pattern Recognition, King-Sun Fu Prize (2002); Honda Lifetime Achievement Award (2000); Professor, Center for Advanced Study, UIUC; IEEE Third Millennium Medal (2000); IEEE Signal Processing Society Award (1991); UIUC University Scholar (1990); appointed by Swedish Science Board to evaluate research in image processing in Sweden (1987, 1990); *Pattern Recognition*, Best Paper Awards (1989, 1993, and 1997); Technical Achievement Award, IEEE ASSP Society (1987); listed as one of world's top 15 machine vision researchers in E. F. Hutton report (1986); Fellow, Japan Society for Promotion of Science (1986); A. V. Humbolt U.S. Senior Science Award (1976–1977); Guggenheim Fellow (1971–1972); Fellow, Optical Society of America (OSA); and International Association of Pattern Recognition (IAPR).

Bruce A. Draper (M'91) received the B.S. in computer science from Yale University, New Haven, CT, in 1984, and the M.S. and Ph.D. degrees in computer science from the University of Massachusetts (UMass), Amherst, in 1987 and 1993, respectively.

From 1993 to 1996, he served as a Senior Post-doctoral Scientist in the Computer Vision Laboratory at UMass. He is an currently an Associate Professor of computer science at Colorado State University, Fort Collins, with research interests in object recognition, including biomimetic vision, color vision, and control of computer vision.

Dr. Draper was general co-chair of the 1999 IEEE Conference on computer vision and pattern recognition, has served three times as an area chair for the International Conference on Machine learning (2000–2002), and has served on program committees for International Conference on Pattern Recognition (ICPR), the International Conference on Vision Systems (ICVS) and many other conferences. He is a reviewer for several professional journals.