

REPORT DOCUMENTATION PAGE

Form Approved
OMB NO. 0704-0188

Public Reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comment regarding this burden estimates or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188,) Washington, DC 20503.

1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE 3 AUG 04	3. REPORT TYPE AND DATES COVERED FINAL PROGRESS 9/20/96 - 3/19/02
4. TITLE AND SUBTITLE VISION STRATEGIES AND ATR PERFORMANCE: A MATHEMATICAL/STATISTICAL FRAMEWORK AND CRITIQUE		5. FUNDING NUMBERS DAAH04-96-1-0445	
6. AUTHOR(S) DONALD E. MCCLURE		8. PERFORMING ORGANIZATION REPORT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) BROWN UNIVERSITY 164 ANGELL STREET, PROVIDENCE, RI 02912		10. SPONSORING / MONITORING AGENCY REPORT NUMBER 36259.1-PH-mur	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) U. S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211		11. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.	
12 a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution unlimited.		12 b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The broad goal of this project is to formulate mathematical models which are sufficiently general to support an analysis of particular algorithms for target detection and recognition from the perspective of classical statistics and information theory. When approached from this viewpoint, questions about the performance of an algorithm for detection, recognition or identification translate into familiar problems in estimation, complexity and hypothesis testing. Consequently, an arsenal of powerful results from statistics and information theory, for example results about optimal codes, most-powerful tests, inference and efficiency, and the complexity of testing highly composite hypotheses, can be exploited to achieve a deeper understanding of the ATR problem. The focus of the research is on performance metrics, various measures of an algorithms performance such as probability of detection, probability of "false alarm," bias/variance tradeoffs for algorithms that learn from training data, and computational complexity.			
14. SUBJECT TERMS		15. NUMBER OF PAGES 16	
		16. PRICE CODE	
17. SECURITY CLASSIFICATION OR REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION ON THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

Enclosure 1

FINAL PROGRESS REPORT

*Vision Strategies and ATR Performance:
A Mathematical/Statistical Framework and Critique*

Funding Number: DAAH04-96-1-0445

September 1996 through March 2002

Principal Investigator: Donald E. McClure, Brown University

Yali Amit, Subcontract PI, University of Chicago

Roger Brockett, Subcontract PI, Harvard University

Donald Geman, Subcontract PI, University of Massachusetts

Sanjoy K. Mitter, Subcontract PI, MIT

TABLE OF CONTENTS

PROBLEM STUDIED	4
SUMMARY OF KEY RESULTS	4
Performance Bounds and Theoretical Performance Analysis.....	4
Statistics of Natural Imagery.....	4
Clutter Modeling	5
Compositional Models.....	6
Learning and Recognition	7
Compression.....	7
Variational Approach to Bayesian Estimation	8
Temporal Information in Recognition.....	8
Data-Driven Performance Optimization	8
PUBLICATIONS	9
Books/Research Monographs	9
Journal Articles and Edited Volumes.....	9
Conference Proceedings.....	11
Manuscripts Submitted and Pending	12
Dissertations.....	12
Selected Invited Addresses.....	13
SCIENTIFIC PERSONNEL.....	13
Senior Faculty Investigators	13
Brown University.....	13
Harvard University.....	14
MIT	14
University of Chicago	14
University of Massachusetts at Amherst.....	14
Postdoctoral Research Associates	14
Brown University.....	14
Harvard University.....	14
MIT	14
University of Chicago	14

Graduate Research Assistants.....	15
Brown University.....	15
Harvard University.....	15
MIT.....	15
University of Chicago.....	15
University of Massachusetts at Amherst.....	15
REPORT OF INVENTIONS	15
BIBLIOGRAPHY	16

Problem Studied

The broad goal of this project is to formulate mathematical models which are sufficiently general to support an analysis of particular algorithms for target detection and recognition from the perspective of classical statistics and information theory. When approached from this viewpoint, questions about the performance of an algorithm for detection, recognition or identification translate into familiar problems in estimation, complexity and hypothesis testing. Consequently, an arsenal of powerful results from statistics and information theory, for example results about optimal codes, most-powerful tests, inference and efficiency, and the complexity of testing highly composite hypotheses, can be exploited to achieve a deeper understanding of the ATR problem.

The focus of the research is on performance metrics, various measures of an algorithms performance such as probability of detection, probability of "false alarm," bias/variance tradeoffs for algorithms that learn from training data, and computational complexity.

Summary of Key Results

Performance Bounds and Theoretical Performance Analysis

An explicit goal of the project is the determination of *lower bounds* for various measures of performance: *What is the best performance that one can possibly achieve within a generic class of algorithms?*

D. Geman and colleagues have proved that coarse-to-fine (CTF) testing is computationally optimal under a certain statistical model. More specifically, there is a certain classifier for a generic object class (e.g., face or truck) and this classifier may be evaluated (computed) in a great many ways. The classifier is based on a family of "tests." Each way of evaluating the tests sequentially and adaptively, stopping when the state of the classifier is determined, corresponds to a decision tree. Each tree (evaluation) has therefore exactly the same error rates, as it is merely a way of evaluating a fixed function. But the trees have widely varying computational costs. Here one assumes each test has a certain "algorithmic cost." We have proven that under a certain model for the joint distribution of the tests under the hypothesis of "no object," a "coarse-to-fine" evaluation is optimal in the sense of minimizing expected cost. The coarse-to-fine refers to the fact that tests have different levels of coarseness in terms of the number of poses of the object they cover. Another, related accomplishment, was an analysis of the false positive rate of the classifier, obtained by making a transformation to a non-homogeneous branching process and then using more or less classical results.

Statistics of Natural Imagery

One cannot meaningfully speak of the *probability of detection* of a target or of the *probability of a false alarm* without making assumptions about an underlying probability distribution on natural scenes. In the language of statistics, this means we need an explicit *null hypothesis* which models generic backgrounds or generic images. Formulating such models has been a major thrust of this contract. A question of fundamental importance is how to specify an appropriate mathematical model for scenes.

In summer 1999, controlled experiments were carried out to collect and categorize a database of laser range images in different settings: forest, residential cityscapes, indoor

scenes, and miscellaneous others. Following the data collection, systematic data analysis was performed to learn empirical distributions of a variety of image statistics. The particular analyses done were motivated by earlier striking findings of scale-invariance properties of reflectance images. Finally, various mathematical models, both *descriptive* ones and *generative* ones, were proposed and analyzed for their potential to adequately represent the empirical results for the range image data set. (Descriptive models describe analytically the essential properties of the low-dimensional marginal distributions of the image statistics, while generative models incorporate, in addition, the direct connections between those marginal distributions and the geometry and spatial arrangement of objects in the 3D world giving rise to the distributions.)

Other groups have started with the assumption that wavelet decompositions of images are the best way to describe their statistics. We do not believe this is true because we find complex dependencies between pixels and between wavelet coefficients which reflect the preferred geometries of the world. This can be understood by contrasting it with the (false) assumption that natural images are white noise: this would say that all pixels are independent and there is no geometry in scenes. The wavelet point of view says there is very simple geometry but it, too, can be eliminated by taking a suitable filter basis for images. Our approach is to look at large numbers of local 3 by 3 patches both in optical and range images, and describe directly their full joint statistics. This was begun in Huang's thesis and has been extended greatly by Lee and now Pedersen, a visiting graduate student from Denmark. The results are very striking: the empirical distribution of high contrast 3x3 patches has a singular surface along which it has infinite density and in whose neighborhood about half such patches are located. This surface is highly curved and represents images of ideal edges. Thus even a simple geometry in the image plane, a geometry of straight edges, produces a very complex cluster in the vector space of local images.

Clutter Modeling

A principal motivation for studying statistics of natural images is the group's objective of developing realistic and usable mathematical models for clutter in natural images; indeed, the generative models for natural image statistics mentioned above are clutter models. The objective of understanding and modeling clutter was stated in the original proposal and received steady attention throughout the project. For example, Grenander's work on modeling clutter in SAR imagery was reported in an early review meeting. In addition to the new work on generative models for image statistics, Grenander recently introduced the so-called *transported generator model*. The model is a simple one and has the potential to form the foundation of more ambitious clutter-modeling work.

A variety of approaches to expand and improve the models were explored, representing different individual views. Grenander outlined a plan for attacking the clutter-modeling challenge—a plan that his recent work instantiated. Essential elements of the plan include (1) experiments with real data, (2) design of models tuned to different categories of scenes (forest, desert, cityscape...), (3) linkage of the models to the sensor modality (IR, laser radar, SAR, reflectance...), and (4) using the models for algorithm design and performance analysis.

Compositional Models

Compositionality refers to the evident ability of humans to represent knowledge through a hierarchy of part-whole relationships. This is widely believed to be the basis for language representation, and many believe that it is in fact a fundamental organizing principle of cognition. One of the main thrusts of the MURI effort has been to explore the implications of a compositional formulation of the machine vision problem: What are the fundamental computational limitations? What are the implications for performance, in terms of “learning” (parameter estimation) and recognition accuracy?

The first step is to formulate the idea of compositionality as a precise mathematical theory of representation. This has been largely achieved, through a rigorous probabilistic formulation, and there is now the beginning of a theory of inference (learning). Shih-Hsiu Huang, a Ph.D. student, completed a software implementation of a “composition machine.” The key contribution of the thesis is a coarse-to-fine computational strategy that yields, *at any instant*, an approximate image interpretation.

A continuum formulation was worked out, and this yielded a surprising connection to what is known as “scaling” in natural images. The latter refers to the remarkable observation that the statistics of natural images are essentially invariant to scale: “blowing up” or “reducing” pictures of natural scenes preserves the statistical structure. Scaling requires a very specific distribution on the sizes of objects in the image plane. We discovered that compositional systems require the same distribution in order to show scaling.

In the language of formal grammars, our composition system amounts to a probabilistic treatment of context-sensitive grammars. This connection is the basis for an ongoing collaboration with computational linguists. The principles of compositionality make specific predictions about the nature of neuronal circuitry and the physiological solution to the so-called “binding problem.” This connection is the basis for an ongoing collaboration with neuroscientists.

Given a (probabilistic) compositional model, the problem of scene interpretation becomes one of assigning each element of the scene to a compositional structure in such a way that the resulting collection of such structures maximizes probability. This is formally equivalent to the well studied “covering problems,” which is known to be NP-Complete. This brings the computational aspect of computer vision into sharp focus: the basic problem of segmentation, which is a problem that must be solved, simultaneously, at each of many levels of abstraction (which edges go together, which contours are related, which textured patches are part of the same surface, what are the object delineations, which objects are related in a larger “context”...) amounts to the problem of choosing a best covering, and there is no known *general* polynomial-time solution.

Of course, compositional systems have a very special structure and, furthermore, it is clear enough that natural vision systems continuously find good (perhaps nearly optimal) interpretations of scenes. What sort of computational engine can attack the covering problem so effectively? Greedy algorithms can not work---the problem of “what goes with what” is locally ambiguous, which essentially rules out incremental optimization. Monte Carlo methods are universal, but much too slow for high-dimensional problems with complex structure. Dynamic programming, *per se*, doesn't apply because there is no Markov structure---the problem of scene interpretation is fundamentally global.

The scaling properties of natural imagery, and the closely related scaling properties of formal compositional systems, together with the apparent multi-resolution aspect of

feature detection in natural vision systems, strongly suggests a coarse-to-fine computation engine. How efficiently can coarse-to-fine processing solve the image interpretation problem?

This leads, more broadly, to the study of coarse-to-fine computation, and to an effort to make precise the achievable gains in computational speed. These issues were explored both for compositional systems and for closely related Markov systems with very large state spaces. The latter amounts to an analysis of exact coarse-to-fine dynamic programming for general graphical models.

Dynamic programming is the basic computational engine behind the use of hidden Markov models and their generalizations. Brian Lucena, a Ph.D. student, completed a mathematical analysis of a wonderful "coarse-to-fine" approach to dynamic programming, introduced recently by one of our former students, Chris Raphael. Brian is now working on a new class of computationally efficient error-correcting codes.

In ongoing work, an analysis of the fundamental computational limitations inherent in the vision problem will continue. Specifically, coarse-to-fine algorithms will be the focus of experiments and theoretical analysis. Preliminary experiments suggest that, as a rule, exponential speed-ups are available both in the compositional and the Markov (dynamic programming) settings.

Learning and Recognition

A comprehensive neural network model for learning, detecting and recognizing objects has been developed, and applied to the analysis of complex visual scenes. Amit and D. Geman base the network model on the sparse binary feature representations, which have been used in the detection and recognition algorithms developed. Learning is based on local Hebbian learning rules and is carried out in a central module. Robustness to variations in pose is obtained by using 'complex' units that perform an ORing operation over small neighborhoods of the input feature layers over a coarser resolution array. Translation invariant recognition and detection are obtained by hard wiring the appropriate shifting mechanisms. Every shift of the reference grid on the coarse resolution array is copied to allow processing in terms of interactions with the central module. The massive input from the lower layers into higher layers is dealt with through gating mechanisms. Either a specific set of feature/location pairs is gated to allow for detection, a particular shift is gated to allow for classification of the data at a particular location in the scene.

The ability to find stable features of varying degrees of complexity on objects and their sparsity in background allows us to choose from a family of algorithms according to various specifications of failed detection and false positive probabilities. The false positive analysis follows very cleanly from the Poisson statistics, and does not require massive testing on background images.

This approach has been applied to face detection in real scenes, detection of rigid 3d objects in real scenes, detection of symbols in highly cluttered artificially generated scenes. False positive / false negative curves can be predicted using the statistical properties in all these cases.

Compression

Associated with any connected planar region having a sufficiently nice boundary, there is a set of orthonormal functions that are adapted to the region, namely the eigenfunctions

of the Laplacian. These functions reflect the shape of the region and thus suggest themselves possible basis for representing the image. The merit of such an expansion can be evaluated in terms of the number of nonzero coefficients required to achieve a desired level of fidelity. For rectangular regions this specializes to the usual Fourier basis. Of course Fourier series is not particularly useful for non rectangular regions and this is the natural modification. Our experimental work supports the idea that these ideas can significantly improve the bits per pixel and it has been suggested that the coding scheme developed may be consistent with the most recent MPEG standards.

Variational Approach to Bayesian Estimation

Sanjoy Mitter in joint work with Nigel Newton of the University of Essex developed a Variational Theory for Bayesian Estimation which characterizes the conditional distribution as the solution of a variational problem of minimizing a certain Free Energy. This theory is very general and applies to Hidden Markov Models based on a Markov Random Field. This research makes non-trivial connections to recent work on Inference on Graphs, Coding theory and Non-Equilibrium Statistical Mechanics. It has enabled us to solve the long-standing open problem of giving a Variational View of Non-linear Filtering. Several papers describing this work are in preparation.

Temporal Information in Recognition

We developed a coherent statistical/Bayesian framework for tracking of moving objects in highly cluttered environments on the basis of video image sequences. The framework involves three basic components: (i) An object representation, i.e. a model that articulates the overall shape architecture of an object together with the objects random deformations (rigid and non-rigid); (ii) a dynamic model, i.e. a prior on the set of possible trajectories of a moving object; and (iii) an observation model that relates, at each video frame, the image gray-level data to the object and dynamic models, and articulates the random variability of the image data due to clutter, occlusion, and other image degradation. The combination of these three components leads to a nonlinear filtering problem which is equivalent to a Hidden Markov Model (HMM). We have explored two object representations (a deformable template model, and a hierarchical syntactic model), and two observation models-- a nonlinear one that explores the HMM representation of the filtering problem, and a linear one that employs the hierarchical/syntactic models. The nonlinear observation model is combined with a Monte Carlo based tracking algorithm and runs in real time. The linear observation model is combined with the Extended Kalman Filter (EKF), but at the present time it does not run in real time. Our experiments demonstrate that the Monte Carlo filter performs considerably better than the EKF in cluttered environments; the performance of the two filters is comparable in environments with limited image degradation.

Data-Driven Performance Optimization

The guiding principle of *minimum description length* has been espoused as means of inferring compositional structure in scenes by members of the group from all five participating universities. R. Brockett and colleagues at Harvard have developed ways of using *feedback* about the gradient of description length (or other measures of the complexity of the description of an image) during an iterative computation as a way of tuning or adapting a vision algorithm's parameters optimally for the particular scene being processed. The use of feedback in this setting is reminiscent of the use of feedback loops in control theory. The versatility of the approach to data-driven

performance optimization has been demonstrated on algorithms for (1) region-based image compression and (2) extraction of coherent structure from highly cluttered scenes.

Publications

Books/Research Monographs

Amit, Y. (2000), '2d Object Detection and Recognition: models, algorithms and networks', MIT Press.

Journal Articles and Edited Volumes

Amit Y. (2000) 'A neural network architecture for visual selection', *Neural Computation*, vol. 12, pp. 1141-1164.

Y. Amit and Geman D. (1998), 'Discussion of 'Arcing Classifiers' by Leo Breiman', *Annals of Statistics*, vol. 26 (3).

Amit Y. and D. Geman, 'A computational model for visual selection,' *Neural Computation*, 11, 1691-1715, 1999.

Y. Amit, D. Geman and B. Jedynek, "Efficient focusing and face detection," *Face Recognition: From Theory to Applications*, eds. H. Wechsler et al, NATO ASI Series F, Springer-Verlag, Berlin, 157-173, 1998.

Amit, Y. and Mascaro, M, (2001) Attractor networks for shape recognition, *Neural Computation*, vol. 13, pp. 1415-1442.

Amit Y. and Murua A. (2001), 'Speech recognition using randomized relational decision trees', *IEEE Transactions on Speech and Audio Processing*, vol. 9, pp. 333-342.

R. Brockett and D. Liberzon, "Nonlinear Feedback Systems Perturbed by Noise: Steady-State Probability Distributions and Optimal Control," *IEEE Transactions on Automatic Control*, vol. 45 n. 6, pp. 1116-1130(2000).

R. Brockett and D. Liberzon, "Quantized feedback stabilization of linear systems," *IEEE Transactions on Automatic Control*, vol. 45 n. 7 pp. 1279-1289 (2000).

R. Brockett and N. Ferrier, "Reconstructing the Shape of a Deformable Membrane from Image Data," *Robotics Research*, vol. 19 n. 9, pp. 795-816, (2000).

* S. Casadei and S.K. Mitter, "An Efficient and Provably Correct Algorithm for the Multiscale Estimation of Image contours by means on Polygonal Lines", *IEEE Trans. on Information Theory*, 45(3), April 1999.

* S. Casadei and S.K. Mitter, "Beyond the Uniqueness Assumption: Ambiguity Representation and Redundancy Elimination in the Computation of a Covering Sample of Salient Contour Cycles" Special issue on perceptual organization in computer vision, *Computer Vision and Image Understanding*, October 1999.

F. Fleuret and D. Geman, "Coarse-to-fine face detection," *Inter. J. Computer Vision*, Vol. 41, pp. 85-107, 2001.

D. Geman and B. Jedynek, "Model-based classification trees," *IEEE Trans. Info. Theory*, vol. 47, March 2001, pp 1075-1082.

- S. Geman, E. Bienenstock and D. Potter, 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, 1998.
- S. Geman and Z. Chi, Estimation of probabilistic context-free grammars, *Computational Linguistics*, 24, 1997, 299-305.
- S. Geman and Z. Chi, On the consistency of minimum complexity nonparametric estimation. *IEEE Trans. Inf. Theory*, 44, 1998, 1968-1973.
- S. Geman and K. Kochanek, Dynamic programming and the graphical representation of error-correcting codes. *IEEE Trans. Inf. Theory*, vol. 47, 2001, pp. 549-568.
- S. Geman, D. Potter and Z. Chi, Composition Systems, *Quart. Appl. Math.*, vol. LX, 2002, 707-736.
- B. Gidas and A. Murua, "Optimal Transformations for Prediction in Continuous -- Time Stochastic Processes", in *Stochastic Processes and Related Topics*, Birkhauser 1998, pp. 167-183, eds: I. Karatzas, B.S. Rajput, M.S. Taqqu.
- B. Gidas and D. Mumford, "Stochastic Models for Generic Images", *Quarterly of Applied Mathematics*, Vol. LIX, Number 1, 2001, pp. 85-111.
- U. Grenander, M.I. Miller and A. Srivastava, Hilbert-Schmidt Lower Bounds for Estimators on Matrix Lie Groups for ATR, *IEEE Trans. Pattern Analysis and Machine Intelligence*, 20 (8), 1998, pp. 790-802.
- U. Grenander and A. Srivastava: Probability Models for Clutter in Natural Images, *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 23 (4), April 2001, pp. 424-429..
- U. Grenander, A. Srivastava, M.I. Miller: Asymptotic Performance Analysis of Bayesian Target Recognition, *IEEE Trans. Information Theory*, 46 (4), July 2000, pp. 1658-1665.
- A. D. Lanterman, U. Grenander, M.I. Miller: Bayesian Segmentation via Asymptotic Partition Functions, *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 22 (4), April 2000, pp. 337-347.
- * Ann Lee, Kim Pedersen and David Mumford, The Complex Statistics of High-Contrast Patches in Natural Images, submitted to *Int. J. of Comp. Vision*,
- Ann Lee, Jिंगgang Huang and David Mumford, Occlusion models for Natural Images: A statistical study of a scale-invariant dead-leaves model, *Int. J. of Comp. Vision*, vol. 41, 2001, pp. 35-59.
- Conglin Lu, Yan Cao and David Mumford, Surface Evolution under curvature flow, accepted, *Image Proc. Comp. Vision and Comp. Graphics*, 2001.
- * S.K. Mitter, Control with Limited Information : The role of Systems Theory and Information Theory. In *Newsletter of Information Theory Society*, December 2000.
- David Mumford and Basilis Gidas, Stochastic Models for Generic Images, *Quarterly of Appl. Math.*, vol. 59, 2001, pp. 85-111.
- A. Srivastava, X. Liu and U. Grenander, Universal Analytic Forms for Modeling Image Probabilities, *IEEE Trans. Pattern Analysis and Machine Intelligence*, 24 (9), 2002, pp. 1200-1214.

Conference Proceedings

Y. Amit, (1998), 'Deformable templates for object detection', Tutorial Notes for the IEEE International Conference on Image Processing.

Brockett, R., 'Stationary Covariance Realization with a Specified Distribution of Amplitudes,' Proceedings of the 1998 Conference on Decision and Control, Tampa FL, Dec 1998, pp. 3742-3745.

Brockett, R., 'Stochastic Analysis for Fluid Queuing Systems,' Proc. of the 1999 CDC Conference, Phoenix AZ, 1999, pp. 3077-3082.

Brockett, R., 'Explicitly Solvable Control Problems with Nonholonomic Constraints,' Proc. of the 1999 CDC Conference, Phoenix AZ, 1999, pp. 13-16.

R. Brockett, 'SubRiemannian Gradient Flows,' Proc. of MTNS 2000 Perpignan, France, June 2000.

Brockett, R. and Khaneja, N., 'Dynamic Feedback Stabilization of Nonholonomic Systems,' Proc. of the 1999 CDC Conference, Phoenix AZ, 1999, pp 1640-1645.

F. Fleuret and D. Geman, 'Graded learning for object detection,' Proc. IEEE Workshop on Statistical and Computational Theories of Vision, Fort Collins, CO, June, 1999.

Artur Fridman, The Max-product algorithm in the Ising model, The Learning Workshop, Snowbird, UT, 2000

Artur Fridman, Mixed Markov models, Proc. Second International ICSC Symposium on Neural Computation, Berlin, Germany, 2000.

Artur Fridman and S. Sloman, Categorization versus similarity: The case of container names, in: Similarity and Categorization, M. Ramscar, U. Hahn, E. Cambouropoulos, H. Pain (Eds.), 2000

D. Geman and A. Koloydenko, 'Invariant statistics and coding of natural microimages,' Proc. IEEE Workshop on Statistical and Computational Theories of Vision, Fort Collins, CO, June, 1999.

S. Geman, 'Hierarchy in machine and natural vision,' Proceedings of the 11th Scandinavian Conference on Image Analysis, 1999.

S. Geman, S. Canon, Z. Chi, M. Johnson, and S. Riezler 'Estimators for stochastic "unification-based" grammars,' In: The Proceedings of the 37th Annual Conference of the Association for Computational Linguistics, San Francisco. Morgan Kaufmann. 1999.

S. Geman and C. Raphael, A grammatical approach to mine detection. In: Detection and Remediation Technologies for Mines and Minelike Targets II, SPIE Vol. 3079, Proceedings of SPIE. A.C. Dubey and R.L. Barnard, Editors, 1997, 316-332.

B. Gidas, "Tracking of Moving Objects in Cluttered Environments via Monte Carlo Filter" (with M. Almeida and C. Robertson), Proceedings, International Conference on Pattern Recognition (ICPR2000), Barcelona, Spain, 2000, pp. 1123-1128.

B. Gidas, "Model-Based Simultaneous Tracking and Recognition of Moving Objects" (with M. Almeida), Proceedings of International Conference on Artificial Intelligence/Computer Vision, Rio de Janeiro, Brazil, November 2000, pp. 892-897

J.Huang and D.Mumford, The Statistics of Natural Images and Models, CVPR 1999, pp.541-547, Comp Sci Press

Jinggang Huang and David Mumford, The Statistics of Range Images, CVPR (Comp. Vision and Pattern Recognition) conference, 2000,

C. Li and D. Geman, "Active testing at multiple resolutions," Proc. ASA Conference, Baltimore, 1999.

* S.K.Mitter and Nigel Newton: Duality in Estimation and Control : in Festschrift for Alain Bensoussan' 60th birthday in book entitled Partial Differential Equations and Stochastic Control, December 2000.

D. Mumford, Modeling Perception and Inference in Intelligent Systems, address to "Mathematical Challenges of the 21st Century" (being written up), AMS meeting, UCLA, Aug. 2000.

A. Srivastava and U. Grenander, Metrics for Target Recognition, SPIE Conference: Neural Networks, San Jose, January 1998, vol. 3307, pp. 29-37.

A. Srivastava, X. Liu and U. Grenander, Analytic Image Models and Their Applications, ECCV (1) 2002, pp. 37-51.

Manuscripts Submitted and Pending

* B. Gidas and A. Murua, "Construction of Optimal Transformations for Prediction in Continuous--Time Stochastic Process: Finite Past and Present ", submitted to Probability Theory and Related Fields 2000.

B. Gidas and A. Murua, "Estimation of Nonparametric Linear Functionals of Continuous Time Processes from a Finite Data Set: Optimal Transformations for Prediction", submitted to The Annals of Statistics.

Hiro, Y. and Amit Y. (1998), "Computer-aided diagnosis in chest radiographs: Reduction of false positives by removal of normal anatomical structures based on symmetry between left and right lung regions", Submitted to Medical Physics.

* S.K. Mitter, Paper submitted to IEEE Trans. Information theory : Sequential Rate Distortion Theory : An Optimization Theory Viewpoint.

Dissertations

Louay Bazzi, *Robust Algorithms for Model Based Object Recognition and Localization*, Master thesis, EECS, MIT, Massachusetts, 1999.

* Z. Chi, Probability models for complex systems, Ph.D. Thesis, Brown University, 1998.

Artur Fridman, "Mixed Markov Fields", Ph.D. dissertation, Department of Mathematics, Brown University, May 2000.

* Jinggang Huang, Statistics of Natural Images and Models, PhD thesis, May 2000.

* K. Kochanek, Grammatical representation of algebraic codes, Ph.D. Thesis, Brown University, 1998.

* C. Li, "Classification by active testing with applications to imaging and change detection," Ph.D. Thesis, Univ. of Mass., 1998.

* D. Potter, Compositional vision models, Ph.D. Thesis, Brown University, 1998.

C. Robertson, "Tracking Objects from Image Sequences Using Lagrangian Dynamics and Nonlinear Filtering" Ph.D. Thesis, Division of Applied Mathematics, Brown University, 1998.

S. Wang (1998) A Statistical Model for Computer Recognition of Sequences of Handwritten Digits, with Applications to Zipcodes, Ph.D. Thesis, University of Chicago.

Selected Invited Addresses

Y. Amit, 'How far can we get with counting', Presented at the 'Mathematics in Image Analysis', co-organizer of meeting in Ecole Polytechnique, Paris, September 2000.

Y. Amit, 'A neural network for learning detecting and classifying objects' 'High-Level Vision' IMA workshop, University of Minnesota, November 2000.

S. Geman, Learning mixtures and learning hierarchies. Brown University IGERT Retreat, Little Compton, RI, March, 2000.

S. Geman, Mathematical methods in speech analysis, introduction and overview. IMA Short Course on Speech and Image Analysis, Institute for Mathematics and its Applications, Minneapolis, Minnesota, September, 2000.

S. Geman, Mathematical methods in image analysis, introduction and overview. IMA Short Course on Speech and Image Analysis, Institute for Mathematics and its Applications, Minneapolis, Minnesota, September, 2000.

S. Geman, Variance and Invariance in Machine Vision. Sixth U.S. Army Conference on Applied Statistics, Rice University, October, 2000.

S. Geman, Coarse-to-fine parsing for context-free grammars. Mathematical Foundations of Natural Language Modeling, Institute for Mathematics and its Applications, Minneapolis, Minnesota, October, 2000.

B. Gidas, " Hierarchical/Syntactic Methods for Object Recognition, Speech Recognition, and Protein Folding", International Conference on Mathematical Analysis and Its Applications, Athens, Greece, August 2000.

Ulf Grenander, "Representing Clutter in Computer Vision", Invited lecture, Proc. Intern. Stat. Inst., Helsinki, 1999.

U. Grenander: Cluttered Images Acquired by Laser Radar, invited lecture at Conference on Computer Vision, Goteborg, Sweden, August 2000.

D. McClure, Restoration and reformatting of motion images, IMA Workshop on Image Processing and Low Level Vision, October 2000.

Scientific Personnel

Senior Faculty Investigators

Brown University

Stuart Geman

Basilis Gidas

Ulf Grenander

Donald E. McClure

David Mumford

Harvard University

Roger Brockett

MIT

Sanjoy Mitter

University of Chicago

Yali Amit

Zhiyi Chi

University of Massachusetts at Amherst

Donald Geman

Christopher Raphael

Postdoctoral Research Associates

Brown University

Artur Fridman

Yan Gousseau

Michael Isard

Kevin Kochanek

Lionel Moisan

Ilya Pollak

Daniel Potter

Zhiyong Yang

Harvard University

Stefano Soatto

MIT

Stefano Casadei

Nicola Elia

University of Chicago

Bruno Jedynak

Massimo Mascarò

Kenneth Wilder

Graduate Research Assistants

Brown University

Ting-Li Chen

Zhiyi Chi, Ph.D. 1998

Artur Fridman, Ph.D. 2000

Yun Gao

Jinggang Huang, Ph.D. 2000

Shih-Hsiu Huang, Ph.D. 2001

Gershon Kutliroff, Ph.D. 2002

Ann Lee, Ph.D. 2002

Conglin Lu, Ph.D. 2002

Brian Lucena, Ph.D. 2002

Christopher Robertson, Ph.D. 1998

Harvard University

Mark Hoefer

Manuela Sousa

Andrew Carlson (undergraduate student)

MIT

Louay Bazzi

Maurice Chu

University of Chicago

Gilles Blanchard

Zafiris Spyropolous

Steve Wang, Ph.D. 1998

University of Massachusetts at Amherst

Gilda Garibotti

Alexey Koloydenko

Chunming Li, Ph.D. 1998

Report of Inventions

None.

Bibliography

Please see the previous list of Publications supported by this grant.