Mayur Mudigonda

Evans #567 ,
UC Berkeley,
Berkeley, CA - 94704.

cell: +1 517 575 8900
e-mail: mudigonda@berkeley.edu
website: www.cse.msu.edu/~mudigon1

Research Interests

My research interests are in - vision, machine learning and computational neuroscience

Professional Positions

Redwood Center for Theoretical Neuroscience **Summer 2011 - Present**
Junior scientist/Visiting scholar[1]

Blindsight Corporation **Summer 2011- Fall 2012**
Research Scientist[2]

University of California, Berkeley **Fall 2010- Summer 2011**
Research Associate, Dept. of Psychology

Education

Michigan State University, **Summer 2010**
Masters of Science in Computer Science

Michigan State University, **Summer 2010**
Interdisciplinary Specialization in Cognitive Science

Anna University, Chennai, Tamil Nadu, India **Spring 2006**
Bachelor of Technology in Information Technology

Publications and Abstracts

- **Mayur Mudigonda, Thesis (Master of Science)**,"3D Face Parametrization", *Department of Computer Science, Michigan State University, Spring 2010*
- **Mayur Mudigonda**, Pavan Ramkumar, David Zhu, George Stockman,"Multi-Voxel Pattern Analysis identifies brain regions that discriminate between indoor and outdoor scenes", *Conference of Human Brain Mapping, Barcelona, June 2010*
- **Mayur Mudigonda**, George Stockman,"Superquadric Representation of Face: Towards Reduction of Time Complexity",*SPIE Defense, Sensing and Security,Biometrics Technology for Human Identification VII, April 4-5,2010*

[1] redwood.berkeley.edu
[2] www.blindsight.com

- **Mayur Mudigonda, Thesis (undergraduate)**, "Gesture Recognition for the SIMPUTER using 2D accelerometers", (advisor) Prof.V.Vinay (IISc), Prof Aravindan (SSNCE), 2006
- **Mayur Mudigonda**, Arjun V.S., Ashwin Kumar, Karthik Srinivasan,"VISION- Engineering Solutions For The Visually Challenged", *Computer Society International Design Contest, June 2005, Washington D.C.*,

In preparation/Under review

- **Mayur Mudigonda**, Ian Stevenson, Urs Koster, Chris Hillar and Bruno Olshausen, *Predicting V1 neural responses to natural movies using the shift-invariant bispectrum*, COSYNE 2013, Salt lake city, UT
- Jascha Sohl-Dickstein, **Mayur Mudigonda**, Mike Deweese, Surya Ganguli and Bruno Olshausen, *Reduced flipping in Hamiltonian Markov Chain Monte Carlo algorithms*, International Conference on Machine Learning (ICML) 2013, Atlanta,
- **Mayur Mudigonda**, Harold Williams, Risi Kondor, Chris Hillar and Bruno Olshausen, *Invariant object representation using bispectrums*, International Conference on Learning Representations (ICLR) 2013, Atlanta
- **Mayur Mudigonda**, Pavan Ramkumar, *Learning Machines Through Humans: An Inductive framework to learn new features*

Professional Experience

Teaching

- **GSI** VS 265: Neural Computation[3].The goal of this class is to familiarize students with the major theoretical frameworks and models used in neuroscience and psychology, and to provide hands-on experience in using these models. Topics include neural network models, supervised and unsupervised learning rules, associative memory models, probabilistic/graphical models, sensorimotor loops, and models of neural coding in the brain. *Fall 2012*
- **Head Instructor,** CSE 101 : Introduction to Computing and Concepts Department of Computer Science and Engineering, Michigan State University Summer 2009 - Summer 2010
- **Instructor,** CSE 101 : Introduction to Computing and Concepts Department of Computer Science and Engineering, Michigan State University Summer 2008 - Summer 2009
- **TA,** CSE 232 : Introduction to Programming II Department of Computer Science and Engineering, Michigan State University Spring 2008
- **TA,** CSE 420 : Computer Architecture Department of Computer Science and Engineering, Michigan State University Fall 2007

[3]graduate level course

Internships

- Research Intern: Picopeta Simputers (Affliated to The Indian Institute of Science), Spring 2006 - Summer 2006

Startup Experience

- Chairman **Vivartha** - A media portal that sought out non-mainstream content expressed through print and the web. Raised 5000$ and built a team of 12 people including writers, editors and salespersons. Spring 2006 - Spring 2007.

Research Writing

- Writing tutorials in Bayesian learning and image processing for undergraduate students in conjunction with Prof.Vinay[4]

Affiliations and Service

- Journal club organizer for Topics in Computational Neuroscience (TCN)[5]
- Institute of Electrical and Electronics Engineers (IEEE)
- Society for Photographic Instrumentation Engineers (SPIE)
- Reviewed papers (anonymously) for Systems, Man and Cybernetics
- Reviewer papers (anonymously) for IEEE WACV 2009
- Graduate representative for Distinguished Speaker Series in Cognitive Science
- Chairperson, Alumni Research Council

Programming

- C, C++, Matlab, Linux shell scripting, Python, LaTeX 2_ε, PHP, MYSQL
- Used AFNI for analyzing fMRI data (Analysis of Functional Neuro Images)

Honors and Awards

- Summer Research Fellowship 2008, *Deparment of Cognitive Science*
- Outstanding student achievement *Spring 2006*, Department of Information Technology, SSNCE
- Lead a team that was *Top 10 contestants* at the *IEEE-Computer Society International Design Contest 2005*

[4] www.limberlink.com
[5] redwood.berkeley.edu/wiki/TCN

References

Available upon request
Full list of projects in the next pages.

Projects

- **Reduced flipping in Hamiltonian Markov Chain Monte Carlo algorithms**

  Fall 2012 - Present

  Hamiltonian Monte Carlo with partial momentum refreshment, in the style of [Horowitz1991], explore the state space more slowly than they otherwise would due to the momentum reversals which occur on proposal rejection. These cause trajectories to double back on themselves, leading to random walk behavior on timescales longer than the typical rejection time, and leading to slower mixing. We explore a technique by which the number of momentum reversals can be reduced. This is accomplished by maintaining the net exchange of probability between states with opposite momenta, but reducing the rate of exchange in both directions such that it is 0 in one direction.

  This can further be extended by using a second pair of auxiliary momenta and expanding the list of operations that can be performed from each state to include Swaps, Flip-Swaps, etc. We expect that these extra momenta will better use the DOF (Degrees of Freedom) to explore the state space more efficiently.

- **Invariant object representations using bispectrums**

  Summer 2012 - Present

  The field of machine learning has benefited from algorithms of unsupervised statistical learning theory such as K-means, sparse dictionary learning, restricted Boltzmann machines, and deep belief networks. Notably, these tools have been applied successfully in computer vision object discrimination tasks, achieving small error rates on several benchmarking datasets. Specifically, a set of codes or features is learned in an unsupervised manner using image patches from the dataset; these features then form the input to efficient (supervised) hyperplane separating methods (e.g., support vector machines). One difficulty with such feature-extraction algorithms is that training is computationally expensive, making scaling to many object categories challenging. Here we explore an efficient one-time preprocessing step for use in these and other image algorithms that takes advantage of the natural invariances in object space. Specifically, we compute a set of rotationally and translationally invariant features from images using the mathematical machinery of the *bispectrum* in the representation theory of groups. Using these bispectrum features alone (i.e., without training over a dataset), we achieve low error rates with linear classifiers on datasets MNIST and NORB.

- **Predicting V1 neural responses to natural movies using shift-invariant bispectrums**

  Fall 2012 - Present

  Evidence from electrophysiology [Purpura1994, Felsen2005] suggests that the visual system is highly sensitive to higher-order stimulus statistics. However, most models for the stimulus response of V1 neurons are limited to first- and second-order statistics, i.e. features

defined on raw pixels or the power spectrum of the stimulus [David2005]. We explore the image bispectrum as a way to capture higher order features in the stimulus. We show that the performance of spiking response models can be improved by including these higher order features compared to only first and second order features. The bispectrum, which consists of the products of pairs of complex Fourier coefficients, has been used by researches in machine learning and image coding to produce invariant representations and characterize higher order image features such as curvature [Krieger1997]. The elements of the bispectrum are translation-invariant like the power spectrum, yet retain relative phase information. This allows the bispectrum to capture features such as sharp edges, corners and T-junctions, which may underlie response properties of V1 cells. We test this hypothesis by fitting models to 128 cells recorded from cat and primate primary visual cortex. Three different models were fit to each cell: 1) raw pixels, 2) power spectrum and 3) the bispectrum of the stimulus movies. For 27/128 cells, the bispectrum model outperforms the pixel model and the power spectrum model. Thus, while the majority of cells can be better described as either simple cells with the pixel model or complex cells with the power spectrum model, a significant fraction (21%) of cells have more complex receptive fields and can be better modeled in terms of bispectrum features.

- **Multi-view learning methods: How neuroscience and computer vision can benefit from each other**

Fall 2011 - Summer 2012

In the content based image retrieval (CBIR) literature, the semantic gap is defined as the elusive gap between low-level features and higher-order conceptual information. How do we bridge the semantic gap? To address this question, we study the feasibility of applying multi-view learning (MVL) methods to simultaneously classify natural stimuli from their computational (view 1) and neural (view 2) representations. Given two feature representations or views of a data set (such as computational and neural descriptors of a natural scene), MVL algorithms such as co-training (Blum & Mitchell,1998) attempt to train classifiers on each view by constraining them to agree. Studying how classifiers trained on both computational (image features) and neural representations (neuroimaging data) of the stimuli exchange information can potentially quantify the transformation that describes the semantic gap. Here, we focus on natural scene representation by analyzing magnetoencephalography (MEG) data acquired during a natural scene categorization task along with features extracted from the scene stimuli.

Classification on the combined feature space (OB+MEG/SpEn+MEG) resulted in accuracies of 80% or greater across all categories which was at least 15% better than a supervised classifier on any single feature space (MEG/SpEn/OB). This suggests that the two views offer complementary information about the scene. However, our experiments with Co-EM mostly performed poorer than supervised classification, except in a few cases. In future work, we intend to apply new MVL algorithms and define new measures that measure the nature of information exchange between classifiers.

- **Text/Object detection in the wild to aid the blind and aging**

Summer 2011 - Fall 2012

The task of text detection and recognition is a hard problem for computer vision. Challenges include extracting computationally tractable features that separates text from background in the wild. Lighting, contrast and skews further make this problem harder. The goal of the project was to create a mobile application that visually challenged people could use to interact with the world better. To this effect, we released TextDetective on the Apple store late in the summer of 2012. Lot of work involved making the algorithms faster by using more efficient models, algorithms and implementations. More details can be furnished upon request.

- **Correlating visual and fMRI features**

Sep 2009 –Jan 2010

Understanding how feature spaces (of stimuli) correlate with activation in different brain regions, could lead to insight about feature processing for the fields of neuroscience and computer vision. In this work, we wish to explore a novel line of analysis to compare feature spaces of the stimuli used in experiments, with their corresponding fMRI data. We propose a two stage procedure of decomposing inputs into their principal components(**PCA**) and performing a canonical correlational analysis (**CCA**). We applied this model to a visual processing study, and confirmed that colour histograms extracted from the stimulus images presented during a 25 s block, correlated ($\rho = 1$) strongly with fMRI data from regions of early visual processing such as the lingual gyrus and the cuneus.Currently, we are extending this idea as comparison framework for other visual features commonly used in Computer Vision.

- **Improving Semantic Classification: A Semi-Supervised Approach**

Jan 2009 –June 2009

To improve the performance of classification, we experimented with a multi-view semi-supervised approach. We employed two views, fMRI data and the stimuli used while collecting the neuro images for the indoor-outdoor cortical study by Henderson, et al. We extracted features from both these views and ran co-training algorithms. We observed an improvement in performance of both classifiers under some conditions but performance was not always repeatable.

- **Parametric representation of 3D faces**

Aug 2008 –Present

Commercial implementation of biometric systems are bounded by the time taken to match each individual's scan with every other individual in the system (in the case of recognition problems). Thus, any real-life system would be ineffective without some intelligent partitioning of search spaces. The goal of the project is to parametrize human faces using a mixture of geometric model and filters, so that faces may then be binned based on their parametric forms. A geometric model also provides for a more robust normalization of face axes. In this

work, we choose to test the descriptive power of superquadrics for geometric modeling. Employing superquadrics and eigen decompositions, we show that we can bin the data into 10-20 bins with an accuracy of about 80%. This goes to show that the larger geometric features are not sufficient to create large number of bins. We also experimented with semi-supervised clustering methods proposed by Basu et al and found that we can cluster data into 30-40 bins with an accuracy of 80%. Lastly, we experimented with various Gabor features in an Active Shape Modeling (ASM) framework to model smiles on 3D Faces.

- **Whole Brain Classification: Face vs Scene**

  Aug 2008 –Present

  Traditional analysis of fMRI data (General Linear Model -GLM) require a certain amount of prior knowledge, in terms of regions of interest and are primarily hypothesis driven. GLM approaches also take a more binary approach to voxel activation,i.e. GLM requires a voxel to be significantly active for the entire experimental condition. Machine learning approaches such problems with minimal prior knowledge and models are data-driven. The advantage of machine learning approaches is that it is possible to find discriminating patterns that might consist of combinations of partially or fully activated voxels to distinguish between two cognitive states. Employing Support Vector Machines(SVM), we show that for the face vs scene experiments, that our approach classifies blocks of fMR scans accurately at nearly 90%. We are currently developing new models of classification using Graph Laplacians, to enforce similarity in the weights ousingf the classifier, employing geometrical (Euclidean) similarity. That is to say, the nearest neighbours of a voxel in 3D have the same weights.

- **CIPHER - Psychophysical Inspired Retrieval**

  Jan 2007 –May 2007

  The aim of the project was to automate *relevance feedback* for image retrieval using biophysical data. Relevance feedback is the process of *quantifying* how successful a retrieval was. This was done through extracting feature vectors from ECG (Electro Kardio Gram) and GSR (Galvanic Skin Response) signals and constructing a classifier. The project was implemented in Matlab. In our experiments, we employed auto correlation coefficients of EKG, GSR as parameters for classification. We experimented with a range of classifiers, such as Support Vector Machines and decision tree classifiers(CART,etc). We found better performance with decision trees, above 85% accuracy. We hypothesize that due to the excessive number of features(raw) and a small data set the support vectors are unable to align themselves.

- **Privacy preserving Data Mining**

  Aug 2006 –Dec 2006

  The project aimed at investigating new approaches that can better the current state of the art for privacy preserving data mining through random data perturbations. A lot of simple experiments were conducted to find perturbations that can provide privacy but at the same point allow data mining algorithms to be run on them. We showed that the state of the art algorithms are inturn dependent on the security of the filtering methods.

- **ENTS Environment Nurturing and Tracking System**

  Dec 2005 –April 2006

  The project proposed a design that would enable easy, effective and intelligent monitoring of water bodies with the option of being scalable to other ecosystems by suitably changing the sensor sub system. The project was a collaborative inter-disciplinary effort involving experts from Mechanical, Electrical, Geology and Computer Sciences. The center of the project was an innovative module called CABA, a case based intelligent monitoring system.

- **Gesture Recognition For SIMPUTER Using 2-D Accelerometer**

  Dec 2005 –June 2006

  The aim of the project was to develop a user-interface for the low cost hand held computer - SIMPUTER [6]; such that it will enhance the human-computer interaction experience. The proposed system employs inexpensive 2-D accelerometers to recognize gestures. The gestures are dynamic and the recognition system is identification based. The system employs a 2 layer feed forward neural network to perform recognition. Towards the end of the project we experimented with Hidden Markov Models (for classification) for recognizing digits. The project was guided by Dr.V.Vinay of Indian Institute of Science, co-inventor of the Simputer.

- **Vision- Engineering Solutions For The Visually Challenged**

  Dec 2004 –June 2005

  Lead a team that designed a comprehensive system with interactive voice enabled environment control and learning system. The system comprised of novel ideas that alleviated the problems faced by the visually challenged. Visual Basic and C were employed for the project. The project was adjudged top 10 world wide and won critical acclaim at the world finals of CSIDC[7] at Washington D.C.

[6] www.simputer.org

[7] http://www.computer.org/portal/site/ieeecs/menuitem.c5efb9b8ade9096b8a9ca0108bcd45f3/index.jsp?&pName=ieeecs_level1&path=ieeecs/education/csidc/CSIDC2005&file=2005WInners.xml&xsl=generic.xsl