

Ajay Joshi

<http://www.cs.umn.edu/~ajay>

520 2nd Ave SE Apt 205
Minneapolis MN 55414

E-mail: ajay@cs.umn.edu
Phone: (612) 817-1869

SUMMARY

I am interested in Machine Learning, Computer Vision, Pattern Recognition, and Data Mining. My experience includes classification algorithms, active and semi-supervised learning, clustering, probabilistic models, object recognition in images, image matching, segmentation, and video analysis. Along with research, I am heavily interested in writing code for building real systems.

EDUCATION

University of Minnesota, Twin Cities

Ph.D., Computer Science & Engineering Spring 2011 (expected)
Dissertation: Image Classification with Minimal Supervision

M.S., Electrical & Computer Engineering Aug 2005 – May 2008

Govt. College of Engineering Pune, India

B.E., Instrumentation and Control Engineering Jul 2001 – Jun 2005

EXPERIENCE

Fuji-Xerox Palo Alto Labs, CA Jun 2010 – Aug 2010

Intern

Developed an adaptive machine learning algorithm for accurately finding humans in images.

Mitsubishi Electric Research Labs, Cambridge, MA Jan 2010 – May 2010

Intern

Developed dictionary learning methods for efficient compression of images that have block structure.

Mitsubishi Electric Research Labs, Cambridge, MA May 2008 – Aug 2008

Intern

Developed scalable active and semi-supervised learning methods for multi-class classification with little training data.

University of Minnesota, Twin Cities Jan 2006 – present

Research assistant

Thesis research in Computer Vision and Machine Learning.

SKILLS

Programming: C++, Python, MATLAB, Objective-C, OpenCV, iPhone SDK

Platforms: Windows, Unix/Linux, Mac OS X

HONORS

- Doctoral Dissertation Fellowship, University of Minnesota, 2010-2011.
- Finalist, Microsoft Research Ph.D. Fellowship, 2009.
- NSF travel award for Doctoral Spotlight presentation at CVPR 2009, Miami.
- ITS Graduate Student Award in Computer Vision, 2008 (2 awards across 4 states).
- 1st rank (of about 600) for 2 years, University of Pune, India.

PROJECTS

- **SteadyShotCamera iPhone App (free):**

SteadyShot allows you to take blur-free pictures, especially in low illumination conditions. It uses the built-in accelerometer on the iPhone, and takes a picture only when the phone is steady enough so as to avoid blur due to hand shake. It is an ongoing hobby project, and has 3500+ downloads so far.

- **Yahoo! Learning to Rank Challenge:**

The task was to rank webpages for queries (training data: 473000 urls, 19900 queries, 700 dimensional). I used SVM ensembles (100 ranking SVMs) trained over subsets of training data, obtained via clustering. Obtained ERR of 0.4424 and NDCG of 0.7611 (0.4611 and 0.7995 respectively, for the winning team) <http://learningtorankchallenge.yahoo.com/leaderboard2.php?track=1&n=100>. Placed at rank 100 (“Xcoder”) with thousands of competitors worldwide. Hobby project – all code was written in Python.

- **Active learning for image classification** (Part of Ph.D. thesis research)

Proposed novel techniques for large multi-class classification tasks such as object recognition in images. Our methods produce high quality classifiers often achieving 4 to 5 times reduction in the amount of training required over conventional machine learning methods.

- **Visual descriptor search and matching** (Part of a surveillance system)

Developed and deployed a visual search system to identify images of people and objects in large databases using snapshots. The system is useful for applications such as searching for similar images on the web and camera-based surveillance in public places.

- **Scene-adaptive human detection with incremental learning** (Part of Ph.D. thesis research)

Developed a machine learning system to incrementally train classifiers for a detection task (such as person detection in images), in order to allow for continuous adaptation in changing scene conditions (e.g., changing illumination), and deployment in previously unseen scenarios.

- **Pedestrian counting in video**

Collaboratively developed and deployed a system to count the number of pedestrians in a potentially crowded scene. The method gives reliable people counts and real-time performance for video. Particular applications of interest are surveillance and crowd behavior modeling.

- **Graphical text and image modeling**

Developed a system for clustering images in an unsupervised fashion so as to allow clusters of semantically related content.

- **Detecting unusual crowd activity in video** (Part of a surveillance system)

Successfully developed a system that can identify unusual crowd activity in video, based on previously seen patterns of usual activity. Particularly useful for continuous monitoring of sensitive sites.

SELECTED PUBLICATIONS

- Ajay J. Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos, “Scalable active learning for multi-class image classification.” Under review at *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*.
- Ajay J. Joshi and Nikolaos Papanikolopoulos, “Learning to detect moving shadows in dynamic environments.” In *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 30(11): 2055–2063, 2008.
- Prahlad Kilambi, Evan Ribnick, Ajay J. Joshi, Osama Masoud, and Nikolaos Papanikolopoulos, “Estimating pedestrian counts in groups.” In *Computer Vision and Image Understanding (CVIU)*, 110: 43–59, 2008.
- Ajay J. Joshi and Fatih Porikli, “Scene-adaptive human detection with incremental active learning.” In *IEEE International Conference on Pattern Recognition (ICPR)*, 2010.
- Ajay J. Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos, “Breaking the interactive bottleneck in multi-class classification with active selection and binary feedback.” In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010.
- Ajay J. Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos, “Multi-class batch-mode active learning.” In *IEEE International Conference on Robotics and Automation (ICRA)*, 2010.
- Ajay J. Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos, “Multi-class active learning for image classification.” In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. **Doctoral Spotlight presentation.**
- Ajay J. Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos, “Multi-class active learning with binary user feedback.” In *NIPS Workshop on Analysis and Design of Algorithms for Interactive Machine Learning (ADAILM)*, 2009.

PATENTS

- *Active Learning Method for Multi-Class Classifiers*, Fatih Porikli and Ajay J. Joshi. U.S. Patent filed, 2009.
- *Multi-Class Active Learning with Binary User Feedback*, Fatih Porikli and Ajay J. Joshi. U.S. Patent filed, 2010.

COURSEWORK

Computer Vision
Advanced Topics in Graphical Models
Optimization Theory
Artificial Intelligence
Computational Vision and Robotics

Machine Learning
Nonlinear Optimization
Computational Aspects of Matrix Theory
Probability and Stochastic Processes
Digital Signal Processing

VISA STATUS

F1 (Student) Visa.

REFERENCES

Available upon request.