MANDAR JOSHI IBM RESEARCH, BANGALORE mandarjoshi.90@gmail.com 2 +91 9665055898 & www.cse.iitb.ac.in/alumni/~mandar12

Research Interests

Machine learning, data mining, and natural language processing.

Education

July 2012 - June 2014	Master of Technology Computer Science & Engineering Indian Institute of Technology (IIT) Bombay CGPA 9.40
July 2008 - May 2012	Bachelor of Technology Computer Science & Engineering National Institute of Technology (NIT) Nagpur CGPA (major) 8.03 CGPA (overall) 7.73 Class rank 27 / 100

Publications

- Mandar Joshi, Uma Sawant, Soumen Chakrabarti. 2014. Knowledge Graph and Corpus Driven Segmentation and Answer Inference for Telegraphic Entity-seeking Queries. In Empirical Methods in Natural Language Processing (EMNLP) 2014, Qatar.
- Mandar Joshi, Rakesh Khobragade, Saurabh Sarda, Umesh Deshpande, Shiwali Mohan. 2012. *Object-Oriented Representation and Hierarchical Reinforcement Learning in Infinite Mario.* In Proceedings of the 24th IEEE International Conference on Tools with Artificial Intelligence (ICTAI), Greece.
- Mandar Joshi, Rakesh Khobragade, Saurabh Sarda, Umesh Deshpande, Shiwali Mohan. 2012. *Hierarchical Action Selection for Reinforcement Learning in Infinite Mario*. In Proceedings of the 6th Starting Artificial Intelligence Research Symposium (co-located with ECAI), France.

Research Experience

• LingVist: Generating concise captions for natural images IBM Research, August 2014 - present

- Summarizing the key aspects of an image is a cognitive tasks that human excel at.
- Objective: To create an automated system that can generate human-like captions for natural images.

• Answering entity-seeking Web search queries

Master of Technology Project, Adviser: Prof. Soumen Chakrabarti, December 2013 - June 2014

- Interpreting keyword queries is difficult due the lack of structure and context.
- Designed a scalable system for joint interpretation and response ranking using Freebase and an entity-annotated Web corpus. The work was accepted as a full paper at EMNLP 2014.

• Index tuning for PostgreSQL

Course Project: Implementation Techniques in Databases, Adviser: Prof. S. Sudarshan, Autumn 2012

- Developed an index tuning tool for PostgreSQL to find the optimal set of indices for a given workload of SQL queries.
- Extended the PostgreSQL source code to support the creation of "fake" indices for calculating index cost.

• Hadoop based system recommendation

Course Project: CS 709 Convex Optimization, Adviser: Prof. Ganesh Ramakrishnan , Autumn 2013

- Developed a movie recommendation system for large data sets.
- Implemented distributed stochastic gradient descent for matrix factorization on Hadoop.

• Entity annotation for Web queries Adviser: Prof. Soumen Chakrabarti, Jan 2014 - March 2014

- $_{\circ}\,$ Developed an entity annotation system to annotate Web queries with Wikipedia entities.
- The system outperformed established frameworks like TagMe on a dataset of 700 queries.

• Web query interpretation and representation

Master of Technology Seminar, Adviser: Prof. Soumen Chakrabarti, Spring 2013

- A semester long literature survey on classification, segmentation, and named entity recognition in Web queries as methods to predict user intent.
- Was awarded highest grade (AA), and the seminar report was acknowledged in Prof. Chakrabarti's tutorial on query interpretation at the 2013 World Wide Web conference, Brazil.

• Reinforcement Learning in Infinite Mario

Bachelor of Technology thesis, Adviser: Prof. Umesh Deshpande, Aug 2011 - May 2012

- Designed an agent that can learn to play Mario, a side-scrolling computer game, using reinforcement learning.
- Used hierarchical reinforcement learning and action selection methods to handle the key challenge of a large state space.

Employment

IBM Research Software Engineer, August 2014 - present LingVist. A machine learning based system that produces concise captions for natural images.

• OfficeBox Software

Intern, May 2011 - July 2011 Implemented a tool for comparing different schema versions and reporting differences.

• IIT Bombay

Graduate Teaching Assistant, July 2012 - June 2014 Graduate teaching assistant for Computer Programming (Autumn 2012, Spring 2013), Data Structures and Algorithms (Spring 2014), and Web Search and Mining (Autumn 2013).

Achievements

- All India Rank 2 in GATE (Graduate Aptitude Test in Engineering) 2012 amongst over 150,000 applicants.
- 21st in ACM ICPC (Inter Collegiate Programming Contest) Asia Regional Round held at Amritapuri in India in December 2010.

Skills

Programming Languages Databases Operating Systems Tools and Frameworks Java, C++, C, Python PostgreSQL, MySQL Linux, Windows Hadoop, scikit-learn, Latex

Selected Graduate Courses

Natural Language Processing, Foundations of Machine Learning, Implementation Techniques in Relational Databases, Advanced Machine Learning, Web Search and Mining, Organization of Web Information, Convex Optimization