

# Ganesh Ramakrishnan, PhD

*Institute Chair Professor*  
Department of Computer  
Science & Engineering,  
IIT Bombay,  
Mumbai-400076, India  
Phone: (022) 25767728  
Url: <http://www.cse.iitb.ac.in/~ganesh>

*Professor-in-Charge*  
Koita Centre for Digital  
Health,  
IIT Bombay,  
Mumbai-400076, India  
Phone: (022) 25720430  
Url: <https://www.kcdh.iitb.ac.in/>

## *Education*

### **B.Tech, Computer Science and Engineering**

1996 – 2000

Indian Institute of Technology Bombay,  
Mumbai, India.

*Enrolled with an All India Rank of 186. Enrolled in the Department of Electrical Engineering and managed a branch change to CSE after the first year, with an institute rank of 3.*

### **PhD, Computer Science and Engineering**

2000 – 2005

Indian Institute Of Technology Bombay,  
Mumbai, India.

*Advisors: Prof. Pushpak Bhattacharyya and Prof. Soumen Chakrabarti*

*Title:* Bridging Chasms in Text Mining Using Word and Entity Associations

*Description:* The thesis poses the problem of underlying meaning extraction from text documents, coupled with world knowledge, as a problem of bridging the chasms by exploiting associations between entities. We utilize two types of entity associations, *viz.* paradigmatic (PA) and syntagmatic (SA). We present first-tier algorithms that use these two word associations in bridging the semantic and lexical chasms. We also propose second-tier algorithms for question answering, text classification, text summarization and word sense disambiguation which use the first-tier algorithms.

## *Areas of interest*

(i) Robust labelled data generation, (ii) data subset selection and summarization, (iii) effective incorporation of domain knowledge for realizing AIs use in resource-constrained operational settings. Ganesh has focused on adaptation of these contributions to real-world end applications such as optical character recognition and its post-editing, sequence-to-sequence tasks such as automated question generation and machine translation, human activity sensing in audio-visual data and video summarization, etc. These adaptations have resulted in technology transfers, products and startups.

I am currently serving as a Professor and prior to that, served as an Associate Professor (2015 - Nov 2018), Assistant Professor (2009 - 2014) and Adjunct Professor (2007 - 2009) at the Department of Computer Science and Engineering, IIT Bombay. I proposed two new courses (*Statistical Relational Learning* and *Convex Optimization*) in the CSE department and have taught 9 courses so far, *viz.*, Foundations of Machine Learning (CS725/CS419), Convex Optimization (CS709), Optimization in Machine Learning (CS769), Statistical Relational Learning (CS717), Data Interpretation and Analysis (CS215), Data Structures and Algorithms (CS213), Theory of Computation (CS208/SI501) and Introduction to Public Health Informatics (DH302). I have also delivered several guest lectures in courses by Prof. Pushpak Bhattacharyya (Natural Language Processing and AI) and Prof. Soumen Chakrabarti (Web Mining). I am a recipient of the Institute Chair Award (2020-2022), Dr. P.K. Patwardhan Technology Development Award 2020, IITB impactful research award 2017 was a recipient of the J.R. Isaac Chair award (2014-2016), an award granted at IIT Bombay to recognize a faculty member for his achievements at a young age. In 2011, I was one of the recipients of the IBM Faculty Award<sup>a</sup>. I have contributed to over 25 projects and advised/co-advised 9 PhDs who have graduated and am currently advising 5 PhD students.

•**Students graduated as main advisor:** (i) Dr. Ajay Nagesh: Post-Doctoral Research Associate at the Computational Language Understanding Lab at the University of Arizona, <https://sites.google.com/site/ajaynagesh/home>, (ii) Dr. Naveen Nair: Senior Scientist, Machine Learning, Amazon, Seattle <https://www.linkedin.com/in/naveen-nair-3340a445>, (iii) Dr. Ramakrishna Bairi: Microsoft India Research Labs, <https://www.microsoft.com/en-us/research/people/rbairi/> - received Excellence in Ph.D. Research Award 2018 (iv) Dr. Ashish Kulkarni: Machine Learning Scientist at Amazon, Bengaluru <https://in.linkedin.com/in/ashishakulkarni> (v) Dr. Rohit Saluja (Scene Text & OCR in Indian contexts) - (co-advised by Parag): Currently pursuing Post Doc with Prof. C.V. Jawahar at IIIT Hyderabad. (vi) Dr. Vishwajeet Kumar (Improving Sequence to sequence models in deep learning): Currently a Research Staff Member at IBM India Research Lab.

•**Current students as main advisor:** (i) Vishal Kaushal (Summarization problems in Machine Learning) (ii) Rishabh Dabral (Machine Learning models for improved home pose-estimation) (iii) Durga Sivasubramanian (Learning with Less Data) (iv) Ayush Maheshwari (Reducing Data labeling efforts using data programming)

I have advised more than 40 MTech projects and 18 BTech projects so far. I am currently advising 6 MTech projects and 4 BTech projects. I have also been jointly-advised a student from CTARA<sup>b</sup> on his PhD and in the past, have advised several Masters projects at CTARA. I have over 70 publications in referred conferences and journals with 90% featuring in A\* or A rated venues. The IRCC booklet titled ‘Glimpses of Research’, published<sup>c</sup> in March 2018, features several recent projects that I am leading. I worked extensively in the area of human interaction in machine learning, feature induction and relational learning in machine learning, including algorithms and data structures for scaling them up. I lead a project on Programmable Machine Translation<sup>d</sup>, ICT for rural areas, etc. Apart from these, I made significant contributions to Sandhan, an Indian Language search engine; Search over entities and relationships<sup>e</sup>; BET, a tool for Inductive Logic Programming that integrates several existing algorithms and induction frameworks (BET stands for Background + Example = Theory).

**IBM India Research Labs,  
Delhi, India**

December 2004 – March 2009

I worked as a Research Staff Member in the Unstructured Information Management (UIM) group at the IBM India Research (IRL) Labs until 2009 March. I have worked on the following projects during the five years of my affiliation with IRL. (i) eDiscovery: Risk management and Compliance Software, (ii) SystemText for Information Extraction (iii) IBM OmniFind Personal E-mail Search (IOPES) (iii) System Text for Information Extraction (iv) Mining Conversational Patterns (v) Study of ILP Procedures for Compact Feature Construction (vi) Scalable Systems for Information Extraction (vii) Manthun: Churning out Information

*Courses  
Taught*

**CS337+CS335: Artificial Intelligence & Machine Learning**

*I taught this course in autumn 2019 and autumn 2020. With the help of CDEEP, I have also made all the video recordings of my course available online on youtube <http://bit.ly/cs337-2019>*

**CS725: Foundations of Machine Learning**

*I taught this course in autumn 2010 and autumn 2011, Autumn 2015, Spring and Autumn 2016 as well as 2017. More details can be found at <http://www.cse.iitb.ac.in/~CS725>. With the help of Prof. D.B. Phatak and his team, I have also made all the videos, slides, tutorials etc of my course available online on youtube <http://bit.ly/cs725-2016>.*

**CS709: Convex Optimization & CS769 Optimization for Machine Learning**

*I introduced this course in our department in autumn 2008 and am teaching it almost once every year since. The course had accompanying course notes, homework, programming tutorials and assignments. More details can be found at <http://www.cse.iitb.ac.in/~CS709>. The accompanying video lectures for the 2018 offering can be found at <http://bit.ly/cs709-2018> and those for the 2017 offering at <http://bit.ly/cs709-2017>. Starting 2021, I will be offering a new version of this course as ‘Optimization for Machine Learning’*

**CS101: Computer Programming and Utilization**

*I taught this course in Spring 2019. Video recordings of the course can be found at [http://www.cdeep.iitb.ac.in/vod/vodCloud/course\\_intra.php?ccode=290](http://www.cdeep.iitb.ac.in/vod/vodCloud/course_intra.php?ccode=290). I have also made all the slides, tutorials etc available on <http://bodhitree2019.cse.iitb.ac.in/courseware/course/5/content/>*

**CS213m: Data Structures and Algorithms (Minor)**

*I taught this course in Spring 2017. More details can be found at <http://www.cse.iitb.ac.in/~CS213m>. We have also made all the videos, slides, tutorials etc of the course available through an online course on Edx: <https://www.edx.org/course/algorithms-iitbombayx-cs213-3x>*

**CS215: Data Analysis and Interpretation**

*More details about the course can be found at <http://www.cse.iitb.ac.in/~cs215/index.html>. I have also made all the videos, slides, quizzes etc of my course available through an online course on Microsoft mix.*

**CS717: Statistical Relational Learning**

*I introduced this course in our department in spring 2008 and taught it multiple times later. The course had accompanying course notes, homework, programming tutorials and assignments. More details can be found at <http://www.cse.iitb.ac.in/~CS717>*

**CS419: Machine Learning**

*I taught this minors course in autumn 2012.*

**CS208: Theory of Computation**

*I taught this course in Spring 2010 and also for the Math department in 2007.*

**Guest Lectures in CS635 and CS705 by Prof. Soumen Chakrabarti and CS334 and CS621 Courses by Prof. Pushpak Bhattacharyya**

**TD695: Appropriate Technology**

*I have co-taught this CTARA course thrice with Prof. A.W. Date. I take care of motivating multi-criteria decision making and explaining the details of the Analytic Hierarchical Process<sup>a</sup>. I substantiate this with several case studies and provide a rigorous assignment for the same.*

---

<sup>a</sup>You can access the online service created by me from scratch at <http://10.129.141.100:8080/AHP/login>.

**Training Data Subset Selection for Regression With Controlled Generalization Error**

*Durga Sivasubramanian, Rishabh Iyer, Ganesh Ramakrishnan, Abir De*  
Proceedings of The 38th International Conference on Machine Learning (ICML 2021).

**GRAD-MATCH: Gradient Matching based Data Subset Selection for Efficient-Deep Model Training**

*Krishnateja Killamsetty, Durga Sivasubramanian, Ganesh Ramakrishnan, Abir De, Rishabh Iyer*  
Proceedings of The 38th International Conference on Machine Learning (ICML 2021).

**GLISTER: Generalization based Data Subset Selection for Efficient and Robust Learning**

*Krishnateja Killamsetty, Durga Sivasubramanian, Ganesh Ramakrishnan and Rishabh Iyer*  
Proceedings of The 35th AAAI Conference on Artificial Intelligence, AAAI 2021

**Semi-Supervised Data Programming with Subset Selection**

*Ayush Maheshwari, Oishik Chatterjee, Krishnateja Killamsetty, Ganesh Ramakrishnan, Rishabh Iyer*  
Proceedings of The 59th Annual Meeting of the Association for Computational Linguistics (ACL 2021 Findings)

**Automatic Speech Recognition in Sanskrit: A New Speech Corpus and Modelling Insights**

*Devaraja Adiga, Rishabh Kumar, Amrith Krishna, Preethi Jyothi, Ganesh Ramakrishnan, Pawan Goyal*  
Proceedings of The 59th Annual Meeting of the Association for Computational Linguistics (ACL 2021 Findings)

**Rule Augmented Unsupervised Constituency Parsing**

*Atul Sahay, Anshul Nasery, Ayush Maheshwari, Ganesh Ramakrishnan, Rishabh Iyer*  
Proceedings of The 59th Annual Meeting of the Association for Computational Linguistics (ACL 2021 Findings).

**Select, Substitute, Search: A New Benchmark for Knowledge-Augmented Visual Question Answering**

*Aman Jain, Mayank Kothiyari, Vishwajeet Kumar, Preethi Jyothi, Ganesh Ramakrishnan, Soumen Chakrabarti*  
Proceedings of The 44th International ACM Conference on Research and Development in Information Retrieval (SIGIR), Resource Track, 2021.

**Joint Learning of Hyperbolic Label Embeddings for Hierarchical Multi-label Classification**

*Soumya Chatterjee, Ayush Maheshwari, Ganesh Ramakrishnan and Saketha Nath Jagaralpu*  
Proceedings of The 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2021).

**Meta-Learning for Effective Multi-task and Multilingual Modelling**

*Ishan Tarunesh, Sushil Khyalia, Vishwajeet kumar, Ganesh Ramakrishnan and Preethi Jyothi*  
Proceedings of The 16th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2021).

**Wisdom of (Binned) Crowds: A Bayesian Stratification Paradigm for Crowd Counting**

*Sravya Shivapuja, Mansi Khamkar, Divij Bajaj, Ganesh Ramakrishnan, Ravi Kiran Sarvadev-abhatla*

In Proceedings of The 29th ACM International Conference on Multimedia (ACMM 2021).

**Cross-Modal learning for Audio-Visual Video Parsing**

*Jatin Lamba, Jayaprakash Akula, Abhishek ., Rishabh Dabral, Ganesh Ramakrishnan and Preethi Jyothi*

Proceedings of The 22nd INTERSPEECH Conference (Interspeech 2021)

**Exploration of Spatial and Temporal Modeling Alternatives for HOI**

*Rishabh Dabral, Srijon Sarkar, Sai Praneeth Reddy, Ganesh Ramakrishnan*

In Proceedings of The 9th IEEE Winter Conference on Applications of Computer Vision, WACV 2021

**LIGHTEN: Learning Interactions with Graph and Heirarchical Temporal Networks for HOI in videos**

*Sai Praneeth Sunkesula, Rishabh Dabral, Ganesh Ramakrishnan*

In Proceedings of The 28th ACM International Conference on Multimedia (ACMM 2020), Seattle, USA.

**Data Programming using Continuous and Quality-Guided Labeling Function**

*Oishik Chatterjee, Ganesh Ramakrishnan, Sunita Sarawagi*

In Proceedings of The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI 2020), New York, USA.

**Caption Alignment for Low Resource Audio-Visual Data**

*Vighnesh Reddy Konda, Mayur Warialani, Rakesh Prasanth Achari, Varad Bhatnagar, Japarakash Akula, Preethi Jyothi, Gnesh Ramakrishnan, Gholamreza Haffari and Pankaj Singh*

In Proceedings of The 21st INTERSPEECH Conference (Interspeech 2020), Shanghai, China.

**Vocabulary Matters: A Simple yet Effective Approach to Paragraph-level Question Generation**

*Vishwajeet Kumar, Manish Joshi, Ganesh Ramakrishnan, Yuan-Fang Li*

In Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing, ACL/IJCNLP 2020.

**Watch Hours in Minutes: Summarizing Video with User Intent**

*Saiteja Nalla, Mohit Agrawal, Vishal Kaushal, Rishabh Iyer, Ganesh Ramakrishnan*

In Proceedings of the The 2nd Workshop on Video Turing Test: Toward Human-Level Video Story Understanding, ECCV 2020

**Realistic Video Summarization through VISIOCITY: A New Benchmark and Evaluation Framework**

*Vishal Kaushal, Suraj Kothawade, Rishabh Iyer, Ganesh Ramakrishnan*

ACMM Workshops 2020

**Cross-Lingual Training for Automatic Question Generation**

*Vishwajeet Kumar, Nitish Joshi, Arijit Mukherjee, Ganesh Ramakrishnan and Preethi Jyothi*

In Proceedings of The 57th Annual Meeting of the Association for Computational Linguistics (ACL), 2019, Florence, Italy.

**Sub-word Embeddings for OCR Corrections in highly Fusional Indic Languages**

*Rohit Saluja, Mayur Punjabi, Mark Carman, Ganesh Ramakrishnan and Parag Chaudhuri*

In Proceedings of The 15th International Conference on Document Analysis and Recognition (ICDAR 2019), Sydney, Australia

**OCR On-the-Go: Robust End-to-end Systems for Reading License Plates and Street Signs**

*Rohit Saluja, Ayush Maheshwari, Ganesh Ramakrishnan, Parag Chaudhuri and Mark Carman*

In Proceedings of The 15th International Conference on Document Analysis and Recognition (ICDAR 2019), Sydney, Australia

**A Framework towards Domain Specific Video Summarization**

*Vishal Kaushal, Sandeep Subramanian, Rishabh Iyer, Suraj Kothawade, Ganesh Ramakrishnan*

In Proceedings of The 7th IEEE Winter Conference on Applications of Computer Vision (WACV), 2019, Hawaii, USA.

**Learning From Less Data: Diversified Subset Selection and Active Learning in Image Classification Tasks**

*Vishal Kaushal, Rishabh Iyer, Anurag Sahoo, Khoshrav Doctor, Ganesh Ramakrishnan*

In Proceedings of The 7th IEEE Winter Conference on Applications of Computer Vision (WACV), 2019, Hawaii, USA.

**Demystifying Multi-Faceted Video Summarization: Tradeoff Between Diversity, Representation, Coverage and Importance**

*Vishal Kaushal, Rishabh Iyer, Anurag Sahoo, Pratik Dubal, Suraj Kothawade, Rohan Mahadev, Kunal Dargan, Ganesh Ramakrishna*

In Proceedings of The 7th IEEE Winter Conference on Applications of Computer Vision (WACV), 2019, Hawaii, USA.

**Putting the Horse Before the Cart: A Generator-Evaluator Framework for Question Generation from Text**

*Vishwajeet kumar, Ganesh Ramakrishnan and Yuan-Fang Li*

In Proceedings of The SIGNLL Conference on Computational Natural Language Learning, CoNLL 2019, Hong Kong.

**ParaQG: A System for Generating Questions and Answers from Paragraphs**

*Vishwajeet kumar, Sivaanandh Muneeswaran, Ganesh Ramakrishnan and Yuan-Fang Li*

In Proceedings of The 2019 Conference on Empirical Methods in Natural Language Processing, EMNLP 2019, Hong Kong (Demo paper).

**An Interactive Multi-Label Consensus Labeling Model for Multiple Labeler Judgments**

*Ashish Kulkarni, Narasimha Raju Uppalapati, Pankaj Singh, Ganesh Ramakrishnan*

In Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI), 2018, New Orleans, Louisiana, USA.

**Synthesis of Programs from Multimodal Datasets**

*Shantanu Thakoor, Simoni Shah, Ganesh Ramakrishnan, Amitabha Sanyal*

In Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI), 2018, New Orleans, Louisiana, USA.

**Time Aggregation Operators for Multi-label Audio Event Detection**

*Pankaj Joshi, Digvijay Gautam, Ganesh Ramakrishnan, Preethi Jyothi*

In Proceedings of Interspeech 2018, Hyderabad, India.

**Automating reading comprehension by generating question and answer pairs**

*Vishwajeet Kumar, Kreeti Boorla, Yogesh Meena, Ganesh Ramakrishnan, Yuan Fang Li*

In Proceedings of the 22nd Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD) 2018, Melbourne, Australia.

**Entity Resolution and Location Disambiguation in Ancient Hindu Temples Domain Using Web Data**

*Ayush Maheshwari, Vishwajeet Kumar, Ganesh Ramakrishnan and J. Saketha Nath*

In Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT Demo Track), 2018, New Orleans, Louisiana, USA.

**Open-domain question answering using a knowledge graph and Web corpus**

*Uma Sawant, Soumen Chakrabarti and Ganesh Ramakrishnan*

Information Retrieval Journal, Presented at ECIR 2020. Short version published in ACM SIGWEB Newsletter (invited). 2018.

**Improving the learnability of classifiers for Sanskrit OCR corrections**

*Devaraja Adiga, Rohit Saluja, Vaibhav Agrawal, Ganesh Ramakrishnan, Parag Chaudhuri, K. Ramasubramanian and Malhar Kulkarni*

In Proceedings of the 17th World Sanskrit Conference, Vancouver (WSC), 2018.

**Scalable Optimization of Multivariate Performance Measures in Multi-instance Multi-label Learning**

*Apoorv Aggarwal, Sandip Ghoshal, Ankith M S, Suhit Sinha, Ganesh Ramakrishnan, Purushottam Kar and Prateek Jain*

In Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI), 2017, San Francisco, USA.

**Error Detection and Corrections in Indic OCR using LSTMs**

*Rohit Saluja, Devaraj Adiga, Parag Chaudhuri, Ganesh Ramakrishnan and Mark Carman*

In Proceedings of the International Conference on Document Analysis and Recognition (ICDAR) 2017, Kyoto, Japan.

**A Framework for Document Specific Error Detection and Corrections in Indic OCR**

*Rohit Saluja, Devaraj Adiga, Ganesh Ramakrishnan, Parag Chaudhuri and Mark Carman*

In Proceedings of the 1st International Workshop on Open Services and Tools for Document Analysis (ICDAR-OST) 2017, Kyoto, Japan.

**Beyond clustering: Sub-DAG Discovery for Categorising Documents**

*Ramakrishna Bairi, Mark Carman and Ganesh Ramakrishnan*

In Proceedings of the 25th International Conference on Information and Knowledge Management (CIKM), 2016, Indianapolis, USA.

**A Framework for Task-specific Short Document Expansion**

*Ramakrishna Bairi, Raghavendra Udupa and Ganesh Ramakrishnan*

In Proceedings of the 25th International Conference on Information and Knowledge Management (CIKM), 2016, Indianapolis, USA.

**Query Expansion in Resource Scarce Languages: A Multilingual Framework Utilizing Document Structure**

*Arjun Atreya, Ashish Kankaria, Pushpak Bhattacharyya and Ganesh Ramakrishnan*

In Proceedings of the ACM Transactions on Asian and Low-Resource Language Information Processing (TALLIP), 2016.

**Explicit Query Interpretation and Diversification for Context-driven Concept Search across Ontologies**

*Chetana Gavankar, Yuan-Fang Li, Ganesh Ramakrishnan*

In Proceedings of the 15th International Semantic Web Conference (ISWC), 2016, Kobe, Japan.

**Interactive Martingale Boosting**

*Ashish Kulkarni, Pushpak Burange, Ganesh Ramakrishnan*

In Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI) 2016, New York, USA.

**Learning to Collectively Link Entities**

*Ashish Kulkarni, Kanika Agarwal, Pararth Shah, Sunny Raj Rathod, Ganesh Ramakrishnan*

In Proceedings of the third IKDD Conference on Data Science (CoDS), 2016, Pune, India.

**Building Compact Lexicons for Cross-Domain SMT by mining near-optimal Pattern Sets**

*Pankaj Singh, Ashish Kulkarni, Himanshu Ojha, Vishwajeet Kumar, Ganesh Ramakrishnan*

In Proceedings of the 20th Pacific Asia Conference on Knowledge Discovery and Data Mining (PAKDD) 2016, Auckland, New Zealand

**Numerical Relation Extraction with Minimal Supervision**

*Aman Madaan, Ashish Mittal, Mausam, Ganesh Ramakrishnan, Sunita Sarawagi*  
In Proceedings of the Thirtieth Conference on Artificial Intelligence (AAAI) 2016, Phoenix, Arizona USA.

**Summarizing Multi-Document Topic Hierarchies using Submodular Mixtures**

*Ramakrishna Bairi, Rishabh Iyer, Ganesh Ramakrishnan and Jeff Bilmes*  
In Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL), 2015, Beijing, China.

**Optimizing Multivariate Performance Measures for Learning Relation Extraction Models**

*Gholamreza Haffari, Ajay Nagesh and Ganesh Ramakrishnan*  
In Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics - Human Language Technologies (NAACL-HLT), 2015, Denver, Colorado, USA.

**Generalized Hierarchical Kernel Learning.**

*Pratik Jawapuria, J. Saketha Nath and Ganesh Ramakrishnan*  
Journal of Machine Learning Research, 16(Mar):617652, 2015.

**Context-driven Concept Search across Web Ontologies using Keyword Queries**

*Chetana Gavankar, Yuan-Fang Li, Ganesh Ramakrishnan*  
In Proceedings of the Eighth International Conference on Knowledge Capture (Short paper), 2015.

**A Machine Assisted Human Translation System for Technical Documents**

*Vishwajeet Kumar, Ashish Kulkarni, Pankaj Singh, Ganesh Ramakrishnan, Ganesh Arnaal*  
In Proceedings of the Eighth International Conference on Knowledge Capture (Short paper), 2015.

**Thinking, Pairing, and Sharing to Improve Learning and Engagement in a Data Structures and Algorithms (DSA) Class.**

*Patil Deepti Reddy, Shitanshu Mishra, Ganesh Ramakrishnan, Sahana Murthy*  
In Proceedings of the 2015 International Conference on Learning and Teaching in Computing and Engineering (LaTiCE), 2015, Taipei, Taiwan

**Efficient Reuse of Structured and Unstructured Resources for Ontology Population**

*Chetana Gavankar, Ashish Kulkarni and Ganesh Ramakrishnan*  
In Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014, Reykjavik, Iceland.

**Personalized classifiers: evolving a classifier from a large reference knowledge graph**

*Ramakrishna B. Bairi, Ganesh Ramakrishnan and Vikas Sindhwani*  
In Proceedings of 18th International Database Engineering & Applications Symposium, IDEAS 2014, Porto, Portugal.

**Noisy Or-based model for Relation Extraction using Distant Supervision**

*Ajay Nagesh, Gholamreza Haffari and Ganesh Ramakrishnan*  
In Proceedings of 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP) 2014, Doha, Qatar.

**Enriching Concept Search across Semantic Web Ontologies**

*Chetana Gavankar, Vishwajeet Kumar, Yuan-Fang Li and Ganesh Ramakrishnan*  
In Proceedings of the 12th International Semantic Web Conference, 2013 (Poster as well as a Demo), Sydney, Australia

**Semi-automatic Dictionary Curation for Domain-specific Ontologies**

*Ashish Kulkarni, Chetana Gavankar, Ganesh Ramakrishnan and Sriram Raghavan*  
In Proceedings of the IEEE International Conference on Tools with Artificial Intelligence, ICTAI 2013, Washington DC, USA



**Comparison between Explicit Learning and Implicit Modeling of Relational Features in Structured Output Spaces**

*Ajay Nagesh, Naveen Nair and Ganesh Ramakrishnan*

In Proceedings of the 23rd International Conference on Inductive Logic Programming (ILP), 2013, Rio De Janerio, Brazil.

**Learning to Generate Diversified Query Interpretations using Biconvex Optimization**

*Ramakrishna Bairi, Ambha, Ganesh Ramakrishnan*

In Proceedings of the 6th International Joint Conference on Natural Language Processing (IJCNLP), 2013, Nagoya, Japan.

**Structure Cognizant Pseudo Relevance Feedback**

*Arjun Atreya, Pushpak Bhattacharyya, Ganesh Ramakrishnan*

In Proceedings of the 6th International Joint Conference on Natural Language Processing (IJCNLP), 2013, Nagoya, Japan.

**SATTY : Word Sense Induction Application in Web Search Clustering**

*Satyabrata Behera, Upasana Gaikwad, Ramakrishna Bairi and Ganesh Ramakrishnan*

In Proceedings of the Seventh International Workshop on Semantic Evaluation (SemEval), 2013, Atlanta, Georgia

**Data-based research at IIT Bombay**

*Soumen Chakrabarti, Ganesh Ramakrishnan, Krithi Ramamritham, Sunita Sarawagi, S. Sudarshan*

SIGMOD Record 42(1), 2013

**Towards Efficient Named-Entity Rule Induction for Customizability**

*Ajay Nagesh, Ganesh Ramakrishnan, Laura Chiticariu, Rajasekar Krishnamurthy, Ankush Dharkar, Pushpak Bhattacharyya*

Proceedings of the 2012 Conference on Empirical Methods in Natural Language Processing (EMNLP), 2012, Jeju, Korea

**Compressed Data Structures for Annotated Web Search**

*Soumen Chakrabarti, Sasidhar Kasturi, Bharath Balakrishnan, Ganesh Ramakrishnan, and Rohit Saraf*

Proceedings of the 21st World Wide Web Conference (WWW), 2012, Lyon, France

**Efficient Rule Ensemble Learning in Structured Output Spaces**

*Naveen Nair, Amrita Saha, Ganesh Ramakrishnan, Shonali Krishnaswamy*

Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (AAAI), 2012, Toronto, Canada

**What Kinds of Relational Features are Useful for Statistical Learning?**

*Amrita Saha, Ashwin Srinivasan, Ganesh Ramakrishnan*

Proceedings of the 22nd International Conference on Inductive Logic Programming (ILP), 2012, Dubrovnik

**Challenges in Learning Optimum Models for Complex First Order Activity Recognition Settings**

*Naveen Nair, Ganesh Ramakrishnan, Shonali Krishnaswamy*

Proceedings of AAAI-12 Workshop on Activity Context Representation: Techniques and Languages

**Using Sequential Unconstrained Minimization Techniques to Simplify SVM Solvers**

*Sachindra Joshi, Jayadeva, Ganesh Ramakrishnan, Suresh Chandra*

Neurocomputing 77(1): 253-260 (2012)

**Discovering Customer Intent in Real-time for Streamlining Service Desk Conversations**

*Ullas Nambiar, Tanveer Faruque, Venkata Subramaniam, Sumit Negi, Ganesh Ramakrishnan*

Proceedings of the 20th ACM Conference on Information and Knowledge Management, CIKM 2011, Glasgow, United Kingdom

**Efficient Rule Ensemble Learning using Hierarchical Kernels**

*Pratik Jawanpuria, Jagarlapudi Saketha Nath, Ganesh Ramakrishnan*

Proceedings of the 28th International Conference on Machine Learning (ICML) 2011, Bellevue, Washington, USA

**Parameter Screening and Optimisation for ILP using Designed Experiments**

*Ashwin Srinivasan, Ganesh Ramakrishnan*

Journal of Machine Learning Research 12: 627-662 (2011)

**Web-scale entity-relation search architecture**

*Soumen Chakrabarti, Devshree Sane, Ganesh Ramakrishnan*

Poster Paper in Proceedings of the 20th International Conference on World Wide Web, WWW 2011, Hyderabad

**Enhancing Activity Recognition in Smart Homes Using Feature Induction**

*Naveen Nair, Ganesh Ramakrishnan, Shonali Krishnaswamy*

Data Warehousing and Knowledge Discovery - 13th International Conference, DaWaK 2011, Toulouse, France

**Pruning Search Space for Weighted First Order Horn Clause Satisfiability**

*Naveen Nair, Chander Jayaraman, Kiran TVS and Ganesh Ramakrishnan*

In Proceedings of ILP 2010, Florence, Italy

**BET: An Inductive Logic Programming Workbench**

*Srihari Kalgi, Chirag Gosar, Prasad Gawde, Ganesh Ramakrishnan, Chander Iyer, Kiran T V S, Kekin Gada and Ashwin Srinivasan*

In Proceedings of ILP 2010, Florence, Italy

**An Investigation into Feature Construction to Assist Word Sense Disambiguation**

*Lucia Specia<sup>1</sup>, Ashwin Srinivasan, Ganesh Ramakrishnan, Sachindra Joshi and Maria das Gracas Volpe Nunes*

Machine Learning Journal, 2009

**Using Entity Annotations to Improve Quantity Concensus Queries**

*Amit Singh, Sayali Kulkarni, Ganesh Ramakrishnan and Soumen Chakrabarti*

15<sup>th</sup> ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Demo Track, SIGKDD 2009, Paris, France

**Tunable Feature Weights for Flexible Text Retrieval**

*Natwar Modani, Ganesh Ramakrishnan and Shantanu Godbole*

SNA-KDD 2009, Paris, France

**Relational Learning Assisted Construction of Rule Base for Indian Language NER**

*Anup Patel, Pushpak Bhattacharyya, Ganesh Ramakrishnan*

ICON, Hyderabad, India, 2009

**Learning to rank for quantity consensus queries**

*Somnath Banerjee, Soumen Chakrabarti and Ganesh Ramakrishnan*

32<sup>nd</sup> Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2009, Boston, Massachusetts, USA

**Parameter Screening and Optimisation for ILP using Designed Experiments**

*Ashwin Srinivasan and Ganesh Ramakrishnan*

ILP, Leuven, Belgium, 2009

**Application of Theory of Optimal Search to ILP**

*Srihari Kalgi, Chirag Gosar, Ganesh Ramakrishnan and Ashwin Srinivasan*

ILP, Leuven, Belgium, 2009

**Incorporating Linguistic Expertise Using ILP for Named Entity Recognition in Data Hungry Indian Languages**

*Anup Patel, Ganesh Ramakrishnan and Pushpak Bhattacharyya*

ILP, Leuven, Belgium, 2009

**Collective annotation of Wikipedia entities in Web text**

*Sayali Kulkarni, Amit Singh, Ganesh Ramakrishnan and Soumen Chakrabarti*  
15<sup>th</sup> ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, SIGKDD 2009, Paris, France

**Identification of Class Specific Discourse Patterns**

*Anup Kumar Chalamalla, Sumit Negi, L. Venkata Subramaniam, Ganesh Ramakrishnan*  
ACM 17th Conference on Information and Knowledge Management (CIKM), 2008, Napa Valley, California.

**Feature Construction using Theory-Guided Sampling and Randomized Search**

*Sachindra Joshi, Ganesh Ramakrishnan, and Ashwin Srinivasan*  
18th International Conference on Inductive Logic Programming (ILP), 2008, Prague, Czech Republic.

**Optimization Issues in Inverted Index-based Entity Annotation**

*Ganesh Ramakrishnan, Sachindra Joshi, Sanjeet Khaitan, Sreeram Balakrishnan*  
The Third International ICST Conference on Scalable Information Systems (Infoscale), 2008, Napoli, Italy.

**RAD: A Scalable Framework for Annotator Development**

*Sanjeet Khaitan, Ganesh Ramakrishnan, Sachindra Joshi, Anup Chalamalla*  
The 24<sup>th</sup> International Conference on Data Engineering (ICDE), 2008, Cancun, Mexico.

**Learning Decision Lists with Known Rules for Text Mining**

*Venkatesan Chakravarthy, Sachindra Joshi, Ganesh Ramakrishnan, Shantanu Godbole and Sreeram Balakrishnan*  
The Third International Joint Conference on Natural Language Processing (IJCNLP), 2008, Hyderabad, India.

**Book Chapter: Question Answering Using Word Associations**

*Ganesh Ramakrishnan and Pushpak Bhattacharyya*  
Handbook of Research on Text and Web Mining Technologies, Edited by Min Song and Yi-Fang Wu, Published by Idea Group Inc., USA

**Towards Interactive Learning by Concept Ordering**

*Shantanu Godbole, Sachindra Joshi, Sameep Mehta, Ganesh Ramakrishnan*  
The Eighteenth ACM Conference on Hypertext and Hypermedia (HT), 2007, Manchester, UK.

**Using ILP to Construct Features for Information Extraction from Semi-Structured Text**

*Ganesh Ramakrishnan, Sachindra Joshi, Sreeram Balakrishnan, Ashwin Srinivasan*  
The 17<sup>th</sup> International Conference on Inductive Logic Programming (ILP), 2007, Oregon State University - Corvallis, OR - USA.

**USP-IBM-1 and USP-IBM-2: The ILP-based Systems for Lexical Sample WSD in SemEval-2007**

*Lucia Specia, Ashwin Srinivasan, Ganesh Ramakrishnan and Maria das Gracas Volpe Nunes*  
SemEval-2007 - 4<sup>th</sup> International Workshop on Semantic Evaluations

**Word Sense Disambiguation using Inductive Logic Programming**

*Lucia Specia, Ashwin Srinivasan, Ganesh Ramakrishnan, Maria das Gracas Volpe Nunes*  
The 16<sup>th</sup> International Conference on Inductive Logic Programming (ILP), 2006, Santiago, Spain.

**Information Extraction using Non-consecutive Word Sequences**

*Sachindra Joshi, Ganesh Ramakrishnan, Sreeram Balakrishnan, Ashwin Srinivasan*  
IJCAI Workshop on Text-Mining and Link-Analysis, TextLink 2007, Hyderabad, India.

**Entity Annotation based on Inverse Index Operations**

*Ganesh Ramakrishnan, Sreeram Balakrishnan, Sachindra Joshi*

Conference on Empirical Methods in Natural Language Processing (EMNLP), 2006, Sydney, Australia.

**Automatic Sales Lead Generation from Web Data**

*Ganesh Ramakrishnan, Sachindra Joshi, Sumit Negi, Raghu Krishnapuram, Sreeram Balakrishnan*

The 22nd International Conference on Data Engineering (ICDE), 2006, Atlanta, GA, U.S.A

**Text Classification with Evolving Label-sets**

*Shantanu Godbole, Ganesh Ramakrishnan, Sunita Sarawagi*

The Fifth IEEE International Conference on Data Mining (ICDM), 2005, New Orleans, Louisiana, U.S.A.

**A Model for Handling Approximate, Noisy or Incomplete Labeling in Text Classification**

*Ganesh Ramakrishnan, Krishna Prasad Chitrapura, Raghu Krishnapuram, Pushpak Bhattacharyya*

The 13<sup>th</sup> International Conference on Machine Learning (ICML), 2005, Bonn, Germany.

**A Structure-sensitive framework for Text Categorization**

*Ganesh Ramakrishnan, Deepa Paranjpe, Byron Dom*

Conference on Information and Knowledge Management (CIKM), 2005, Bremen, Germany.

**VisualRDR: A general framework for creating, maintaining and learning of ripple down rules for Information Extraction**

*Delip Rao, Sachindra Joshi, Ganesh Ramakrishnan, Avishkar Misra, Sreeram Balakrishnan, Ashwin Srinivasan*

12<sup>th</sup> International Conference on Management of Data, COMAD 2005b, IIIT, Hyderabad.

**Is Question Answering an acquired skill ?**

*Ganesh Ramakrishnan, Soumen Chakrabarti, Deepa Paranjpe, Pushpak Bhattacharya*

The Word Wide Web Conference (WWW), 2004, New York, U.S.A.

**A Gloss Centered Algorithm for Word Sense Disambiguation**

*Ganesh Ramakrishnan, Pushpak Bhattacharya, Prithviraj*

Proceedings of ACL Senseval (Senseval), 2004, Barcelona, Spain.

**Generic Text Summarization Using WordNet**

*Ganesh Ramakrishnan, Kedar Bellare, Navneet Loiwal, Vaibhav Mehta, Atish Das Sarma, Anish Das Sarma, Pushpak Bhattacharyya*

Language Resource Evaluation Conference (LREC), 2004, Lisbon, Portugal.

**Passage Scoring for Question Answering via Bayesian Inference on Lexical Relations**

*Deepa Paranjpe, Ganesh Ramakrishnan, Sumana Srinivasan*

The Twelfth Text REtrieval Conference (TREC 2003), National Institute of Standards and Technology (NIST), Gaithersburg, Maryland, 2003.

**Generic Text Summarization Using Wordnet for Novelty and Hard**

*Ganesh Ramakrishnan, Kedar Bellare, Chirag Shah, Deepa Paranjpe*

The Twelfth Text REtrieval Conference (TREC 2003), National Institute of Standards and Technology (NIST), Gaithersburg, Maryland, 2003.

**Soft Word Sense Disambiguation**

*Ganesh Ramakrishnan, Pushpak Bhattacharya, Prithviraj, Deepa Paranjpe, Soumen Chakrabarti*

Global WordNet Conference (GWC), 2003, Czech Republic.

**Text Representation with WordNet synsets: A soft sense disambiguation approach**

*Ganesh Ramakrishnan and Pushpak Bhattacharyya*

Proceedings of 8<sup>th</sup> International Conference on Applications of Natural Language to Information Systems (NLDB 2003), Burg, Germany. Extended version published in ISI-NIS Journal, Special Issue on Natural Language Interface to Information Systems, 2003.

*Publications  
continued*

**Question Answering using Bayesian Inferencing on Lexical Relations**

*Ganesh Ramakrishnan, Apurva Jadhav, Ashutosh Joshi, Soumen Chakrabarti, Pushpak Bhattacharyya*

Proceedings of the ACL Workshop on *Role of Machine Learning in Question Answering and Summarization*, 2003, Sapporo, Japan.

**Using WordNet Based Semantic Sets for Word Sense Disambiguation**

*Ganesh Ramakrishnan and Pushpak Bhattacharyya*

Workshop on Application of Semantics in Information Retrieval and Filtering (LREC), 2002, Canary Islands, Spain.

**Using WordNet Based Semantic Sets for Word Sense Disambiguation and Keyword Extraction**

*Ganesh Ramakrishnan and Pushpak Bhattacharyya*

Proceedings of International Conference on Knowledge Based Computer Systems (KBCS), 2002, Mumbai, India.

*Granted  
Patents*

**US8447766B2: Method and system for searching unstructured textual data for quantitative answers to queries**

*Somnath Banerjee, Soumen Chakrabarti, Ganesh Ramakrishnan*

**US7904399B2: Method and apparatus for determining decision points for streaming conversational data**

*L V Subramaniam, Ganesh Ramakrishnan and Tanveer A Faruque*

**US20080072134A1: Annotating token sequences within documents**

*Sreeram Balakrishnan, Ganesh Ramakrishnan and Sachindra Joshi*

**US8706730B2: System and method for extraction of factoids from textual repositories**

*Sachindra Joshi, Raghuram Krishnapuram, Nimit Kumar, Kiran Mehta, Sumit Negi, Ganesh Ramakrishnan, Scott R Holmes*

*Draft Book*

**Handbook for Statistical Relational Learning**

*Ganesh Ramakrishnan and Ashwin Srinivasan*

[https://www.cse.iitb.ac.in/~ganesh/papers/HandbookForSRL\\_upcoming.pdf](https://www.cse.iitb.ac.in/~ganesh/papers/HandbookForSRL_upcoming.pdf)

*Tutorials*

**Combinatorial Approaches for Data, Feature and Topic Selection and Summarization**

*Rishabh Iyer and Ganesh Ramakrishnan*

Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020 January, 2021, <https://sites.google.com/view/ijcaitutorial2020summarization/home>

**A Submodular Optimization Framework for Data, Feature and Topic Summarization**

*Rishabh Iyer and Ganesh Ramakrishnan*

24th European Conference on Artificial Intelligence, ECAI 2020, September, 2020, <https://sites.google.com/view/ecaitutorial2020summ/home>

**Human Assisted Machine Learning: Consensus, Domain Knowledge and Performance Measures**

*3<sup>rd</sup> Summer School On Machine Learning: Advances In Modern AI, IIIT Hyderabad,*

July 11, 2018, <http://cvit.iiit.ac.in/mlsummerschool2018/>

**Tutorial on Optimal Subset Selection over DAGs with Applications in Machine Learning**

*Ganesh Ramakrishnan*

2<sup>nd</sup> Indian Workshop on Machine Learning, 2016: <http://www2.cse.iitk.ac.in/~iwml/2016>

**Tutorial on Distant Supervision for Information Extraction: Modeling and Learning Challenges**

*Ganesh Ramakrishnan*

Xerox Research Innovation Challenge and Winter School on Machine Learning: <http://xrci.xerox.com/xerox-research-innovation-challenge>

**Tutorial on Graphical Models for Learning in Natural Language Processing**

*Pushpak Bhattacharyya and Ganesh Ramakrishnan*

International Joint Conference on Artificial Intelligence, January 2007, IJCAI '07

**Tutorial on Graphical Models for Learning in Natural Language Processing**

*Pushpak Bhattacharyya and Ganesh Ramakrishnan*

International Conference on Natural Language Processing, December 2005, ICON '05

*Invited Talks  
(Partial list)*

**Knowledge Understanding and Representation for Automatic Question Generation**

*Invited talk at the MIT AI week in the 'Knowledge Representation Reasoning Meets Machine Learning' Workshop (<https://kr2ml.github.io/ibm-2019/schedule/>). See video at <https://ibm.ent.box.com/v/kr2ml-ibm-19-videos/file/543970929408>*

19 September, 2019

**AI Mentoring Circles**

*Chaired a Panel at an AI mentoring circle (<https://www.research.ibm.com/artificial-intelligence/ai-research-week/schedule/>)*

16 September, 2019

*Invited Talks*  
(*Partial list*)

**Machine Learning for Analyzing Video Content for Internal Security**  
*Invited talk for Training IPS officers at the National Police Academy, Hyderabad*  
10 December, 2019

**Machine Learning for Analyzing Video Content for Internal Security**  
*Invited talk as part of Panel Discussion on ‘Next Generation / Futuristic Smart Policing using IOT (Internet of things), AI (Artificial Intelligence) and related Cyber Security’ at the All India Heads of Police Communication Conference, Vigyan Bhawan*  
19-20 November 2018

**Human Assisted Machine Learning: Consensus Driven Data Curation, Domain Knowledge and Performance Measures**  
*Talk at IIT Bombay Faculty Alumni Network Meeting, Stanford, USA*  
October 13, 2018

**Human Assisted Machine Learning: Consensus, Domain Knowledge and Performance Measures**  
*Keynote talk at SYNAPSE, Microsoft-India wide AI and ML meet, Hyderabad*  
July 5th-6th 2018

**Human Assisted Machine Learning: Consensus, Domain Knowledge and Performance Measures**  
*Talk at Google AI/ML Workshop, Bengaluru*  
March 16, 2018

**AI Solutions for Smart Cities**  
*The Fourth Indian FAN Symposium on Smart and Sustainable Cities, Faculty Alumni Network (FAN) Meet, Taj Exotica, Goa*  
20th January 2018

**Optimizing Performance Measures that Matter in ML: Some Challenges and Successes**  
*FUSS Talk Series at Dept of CSE, IIT Bombay*  
March 8, 2017

**Optimization of real-world performance measures for real-world ML problems**  
*Large Scale Computing and its Applications, Faculty Alumni Network (FAN) Meet, Taj Exotica, Goa*  
20th and 21st January 2017

**Optimal Subset Selection over DAGs with Applications in Machine Learning**  
*2nd Indian Workshop on Machine Learning, IIT Kanpur*  
July 2016

**Optimizing Multivariate Performance Measures for Learning Relation Extraction Models**  
*Microsoft Research India Labs*  
June 2016

**Optimizing Multivariate Performance Measures for Learning Relation Extraction Models**  
*General Electrical (GE) Research*  
June 2016

**Distant Supervision for Information Extraction: Modeling and Learning Challenges**  
*Microsoft Research India (MSRI)*  
December 2016

**Scaling Up Information Extraction and Disambiguation**  
*Talk at Google Research, Mountainview, USA and Yahoo Labs, Sunnyvale*  
July 2013

*Invited Talks*  
(*Partial list*  
- *contd.*)

**Efficient and Optimal Feature Induction: Opportunities and Challenges**  
*Talks at IBM Almaden Research Center, San Jose, USA and Yahoo Labs, Sunnyvale*  
July 2012

**Rule Ensemble Learning Using Hierarchical Kernels**  
*Yahoo! Labs, Bangalore, India*  
December 2011

**Rule Ensemble Learning Using Hierarchical Kernels**  
*IBM India Research Labs, Bangalore, India*  
December 2011

**Efforts on Information Extraction at IBM**  
*ICWIS09 - International Conference on Web Intelligent Systems*  
January 2009

**Scalable techniques for IE**  
*Summer Workshop on Ontology, NLP, Personalization and IE/IR at IITB, sponsored by HP*  
*Labs, Bangalore*  
August 2008

**Scalable techniques for Information Extraction**  
*NLP Winter School, organized at IIIT Hyderabad*  
January 2008

**Efficient Information Extraction using Inverse Index Operations**  
*Ganesh Ramakrishnan*  
*IRL IIT Bombay Joint Workshop on Information Integration, September 2006, IIT Bombay,*  
*Mumbai, India.*

**Language Models for Text**  
*Ganesh Ramakrishnan*  
*The First National Symposium on Modeling and Shallow Parsing of Indian Languages, April*  
*2006, IIT Bombay, Mumbai, India*



*Other  
Professional  
Activities*

- Organization committee member of SubsetML-21 workshop at ICML 2021 <https://sites.google.com/view/icml-2021-subsetml/home>
- Workshop Co-chair for the 25th Pacific Asia Conference on Knowledge Discovery and Data Mining (PAKDD 2021), <https://www.pakdd2021.org/Call/workshop>
- Organization Committee Member for ICML 2021 workshop on ‘SubSetML: Subset Selection in Machine Learning: From Theory to Practice’ (<https://sites.google.com/view/icml-2021-subsetml/home>)
- Senior Program Committee (SPC) Member for the AAAI Conference on Artificial Intelligence, 2021, 2020, 2019, 2018.
- Senior Program Committee (SPC) Member for the International Joint Conference on Artificial Intelligence (IJCAI), 2021, 2020, 2019, 2018.
- Program Committee (PC) Member for the AAAI Conference on Artificial Intelligence, 2017, 2016, 2015.
- Program Committee (PC) Member for the International Joint Conference on Artificial Intelligence (IJCAI), 2016, 2015.
- Program Committee (PC) Member for the International Conference on Knowledge Discovery and Data mining (KDD) 2013, 2014, 2015, 2016, 2017, 2018.
- Program Committee (PC) Member for the International Conference on Computational Linguistics (COLING) 2014, 2016.
- Program Committee (PC) Member for Annual Meeting of the Association for Computational Linguistics (ACL) 2013, 2014, 2015, 2016, 2017, 2018, 2019, 2020, 2021.
- Program Committee (PC) Member for the 25<sup>th</sup> International Conference on Machine Learning (ICML) 2008, 2014, 2015, 2019, 2020, 2021
- Program Committee (PC) Member for International Semantic Web Conference (ISWC) 2014, 2015, 2016.
- Program Committee (PC) Member for SDM 2018 and for the 42<sup>nd</sup> International Conference on Very Large Databases (VLDB), 2016
- Program Committee (PC) Member for the Conference on Inductive Logic Programming (ILP) 2012, 2103, 2014.
- Program Committee (PC) Member for the Pacific-Asia Conference on Knowledge Discovery and Data Mining (PAKDD) 2010, 2011, 2012, 2015.
- Organization Committee Member for COLING 2012 workshops on Information Extraction & Entity Analytics on Social Media Data (<https://sites.google.com/site/coling12iesocialmedia/home>) and Question Answering for Complex Domains (<https://sites.google.com/site/qacd2012/home>)
- Tutorial Co-chair for the Pattern Recognition and Machine Intelligence Conference, 2009.
- Program Committee (PC) Member for the International Joint Conference on Natural Language Processing (IJCNLP) 2009, 2011, 2013.
- Program Committee (PC) Member for ICWIS09 - International Conference on Web Intelligent Systems

**DECILE: Data efficient Machine Learning** <http://www.cse.iitb.ac.in/~decile:2020> – Present

This is a massive open source software effort toward principled human-machine interaction for machine learning. A short presentation of how Decile has been influenced and in-turn also has influenced uses of AI in the field/operational settings or products/startups I have been closely associated with can be found here: <http://bit.ly/decile-deck>. DECILE has 4 important components listed below:

1. Submodlib (in C++ with Python wrappers) <https://github.com/vishkaush/submodlib/>: Submodlib is an efficient and scalable library for submodular optimization which finds its application in summarization, data subset selection, hyper parameter tuning etc.
2. DISTIL (<https://github.com/decile-team/distil>): This is a library in python for Deep dIversified inTeraCTive Learning
3. CORDS (<https://github.com/decile-team/cords>): This is a library in python for COResets and Data Subset selection
4. SPEAR (<https://github.com/decile-team/spear>): This is a library in python for Semi-suPervisEd dAta pROgramming which will also eventually include rule induction through human Interaction.

**Udaan: A Indian Language End-to-End Translation Ecosystem for Breaking the Language Barrier in Education:** <https://www.udaanproject.org/> 2017-present

This project began with a vision to build an End-to-End ecosystem to translate from English to Hindi and all Indian Languages, textbooks and learning materials in Engineering and all main streams of Higher learning. Our approach has been that it will be aided by human effort. We started building lexicons of various technical domains. In parallel to our efforts on researching and developing machine learning, we also set about developing robust bilingual OCR technology and several post-editing tools by which we now have access to digital bilingual dictionaries in the original format. We are therefore able to use the appropriate scientific and technical terms available in Hindi instead of transliterating the English terms. Additionally, by employing our AI-based post-editing workbench, we are now able to translate a technical book in one-sixth the time it would take for a team consisting of domain and linguistic experts working manually. In due course, as our AI and ML engine learns with every page and every book being edited in each domain, we expect to achieve a much shorter turnaround time.

**Learning Shared Representations for Audio and Visual Data: Part of IBM AI Horizons Project:** <https://www.cse.iitb.ac.in/~malta/> 2017-present

Video event analysis aims to identify the events depicted in a video clip. This is a very important task with many applications in higher-level tasks such as video captioning, summarization and classification. Most research in this field has largely focused on the use of visual features from videos for the analysis. However, videos also have audio streams which carry a great amount of information about the video events. While the sound tracks from videos have been the focus of *audio event detection*, jointly using them with the video stream is a relatively unexplored area. We have developed new multi-modal analysis techniques, combining visual and audio features for video event analysis, and further incorporate them into tasks like video summarization.

In our INTERSPEECH 2018 work, we investigated and developed a suite of operators to aggregate evidence of audio-only events over the time domain, yielding better *embeddings* of the audio features. Subsequently, at ACMM 2020 and WACV 2021, we have also published models for efficient detection of human-object interaction by learning Interactions with graph and hierarchical temporal networks using visual data alone. In our interspeech 2020 paper (and a subsequent paper under review), we have, showed the effectiveness of jointly learning embeddings for audio and video features on tasks such as caption alignment, text-to-video search and video-to-text search. See <https://www.cse.iitb.ac.in/~malta/> and <https://www.cse.iitb.ac.in/~ganesh/videosurveillanceanalytics/>

This is an ongoing project, initiated by me in 2016, which got incorporated in 2017 as part of National Centre for Excellence in Technology for Internal Security: [ncetis.iitb.ac.in](https://www.ncetis.iitb.ac.in). As part of this, we have developed several software tools for Video Analytics <https://www.cse.iitb.ac.in/~ganesh/videosurveillanceanalytics/> and a large part of the technology has been officially Licensed to SrivisifAI Technologies Pvt Ltd, a startup based out of Pune. Cameras have emerged as a very effective and important aspect of security and monitoring. In this talk, we present software solutions for analysing surveillance videos, providing summarization and helping raise real time actionable alerts. The solutions also facilitate smart video indexing and search. The solution turns a camera into a tool allowing prevention and detection of events in a proactive manner. This AI-based indigenous solution is backed by extensive research on algorithms for domain specific efficient summarization, query driven summarization, algorithms for generalization based data subset selection for efficient and robust learning, efficient detection of human-object interaction by learning Interactions with graph and hierarchical temporal networks, real-time anomaly detection, crowd counting, etc. The tools have been deployed at several locations including Naval Dockyard Vishakhapatnam (NDV), UP Police, CISF, SPG at PMO, IB (MHA) and our own IIT Bombay campus. A detailed slide deck on all the professional engagements is here: <http://bit.ly/vidan-slides>. For more details on the products please visit <https://www.cse.iitb.ac.in/~ganesh/videosurveillanceanalytics/> and visit <https://www.cse.iitb.ac.in/~ganesh/Publications.html> for the associated publications. Some datasets that have resulted from this work include VISIOCITY (<http://visiocity.github.io/>) - a new benchmark dataset for video summarization. Even during the lockdown we traveled and interacted extensively with ND(V). The software was also used for contactless surveillance on the IITB campus: see <https://www.insightiitb.org/contactless-surveillance/>. Some engagements with state ATS (anti-terrorist squads also involve development of social media analytics tools. These have been highly effective deployments and have earned us a lot of appreciation. This work has now been licensed to *Simulate learning solutions Pvt. Ltd*. Please see the last page for certificate of appreciation from ATS Mumbai for this work.

### **Video Analytics for Monitoring and Performance Evaluation of Skill Development Centers** 2016 – Present

Deen Dayal Upadhyaya Grameen Kaushalya Yojana (DDU-GKY), a scheme by the Ministry of Rural Development (MoRD), has set up a Compliance and Quality Monitoring System (CQMS) for monitoring and performance evaluation of various skill development centers. The objective of this project is to help CQMS by automated or semi-automated analysis of videos from the surveillance cameras installed at these skill development centers by leveraging state-of-the-art machine learning and computer vision techniques for video analytics. The proposed software solution attempts to automate analysis of videos to produce the following statistics about them:

- 1.For classroom / domain lab / IT lab videos:** Name of trainer - assumes availability of trainer faces database, Percentage of time class was conducted, How late a class started?, How early a class was left?, Is it video of a legitimate class?, Number of people wearing boys uniform, Number of people wearing girls uniform, Number of people seen in the video, Heatmap/flowmap of motion in the frame, Presence/absence of DDU-GKY signage
- 2.Only for classroom videos:** Count number of tables/chairs in a classroom
- 3.Only for IT lab videos:** Number of computers in the lab
- 4.Only for domain labs:** Domain specific equipments for other domain labs
- 5.For biometric punching videos:** (people punching and looking at camera one by one), Number of faces detected over the time period, Verify duplicate faces,
- 6.Summary videos**

More about this can be read at [https://www.ircc.iitb.ac.in/IRCC-Webpage/rnd/PDF/GlimpseIITBResearch/Nov2017/N\\_334.pdf](https://www.ircc.iitb.ac.in/IRCC-Webpage/rnd/PDF/GlimpseIITBResearch/Nov2017/N_334.pdf).

**Automating Reading Comprehension by Generating Question and Answer Pairs**  
2017 – Present

Given a piece of text as a sequence of words, the project focused on generating syntactically correct, meaningful and natural questions along with answers to those questions. Such a system has many applications in a myriad of areas such as FAQ generation, intelligent tutoring systems, and virtual assistants. We have used this question system for generating questions for improved reading comprehension as well as self-assessment by the user for several tutorials as in this EMNP 2019 demo paper: <https://www.aclweb.org/anthology/D19-3030/>. We have also built the first question generation data<sup>a</sup> and system<sup>b</sup> in Hindi by leveraging cross-lingual data. More information about our project can be found at [https://www.ircc.iitb.ac.in/IRCC-Webpage/rnd/PDF/GlimpseIITBResearch/Nov2017/N\\_331.pdf](https://www.ircc.iitb.ac.in/IRCC-Webpage/rnd/PDF/GlimpseIITBResearch/Nov2017/N_331.pdf) as well as in our research papers.

**Venter: Intelligent Complaint Resolution System** 2017 – Present

This project was partially supported by Microsoft India Research Labs with Prof. Sunita Sarawagi as a co-PI. The goal of this project has been to create a community platform for analyzing complaints of varied types (broken taps, cutting trees, noise, etc) and various levels (a workplace building, a university campus, or a city). As part of this project, we have deployed the code ([https://github.com/VenterProject/Venter\\_CMS](https://github.com/VenterProject/Venter_CMS)) on a portal where our clients that include NGOs such as <https://www.ichangemycity.com>, <https://speak-up.in> and civic authorities such as MCGM (<http://dm.mcgm.gov.in/central-complaint-registration-system>) can use our services to train their own complaint classification and analysis systems.

**Lokacart: Technology Licensed to Strategic ERP:** <https://lokacart.com> 2017 – Present

Lokacart is an e-commerce platform developed for farmers, and MSMEs in India (with applications in android and iOS as well as web service applications). The Lokacart application is available in three flavours on android - Lokacart for buyers<sup>c</sup>, Lokacart Admin for sellers<sup>d</sup> and Lokacart plus for bulk buyers<sup>e</sup> and in one flavour in iOS - Lokacart. Currently with 194 vendors onboard, the platform automates the process of receiving orders, bill generation and delivery processing through mobile. The buyer can select the products and place an order with the store registered with her/him. The seller keeps an account of orders and the consumer keeps track of orders and billing. Some details of the story behind the Lokacart application are available at the insight-IITB's article <https://www.insightiitb.org/lokacart-app-institutes-innovations/> containing testimonials citing the success stories for the Lokacart App. You can also read the following: short slide deck<sup>f</sup>, example media coverage<sup>g</sup> as One of the projects initiated by IIT Bombay for COVID-19 mitigation (see pages 17 and 65-69)<sup>h</sup>.

**Lokavidya: Technology Licensed to Lokavidya Technologies Pvt. Ltd.:** <https://lokavidya.com/> 2017 – Present

Lokavidya is an open educational ICT architecture that helps to capture, complement, supplement, and disseminate knowledge of existing integral practices. In the present circumstances, the efficient and reliable techniques are needed for collecting, preserving, organizing, and disseminating knowledge. The app has been adopted at a large scale by Ekal Vidyalaya operating in 55000 Indian villages. You can also read more in this Insight Writeup: <https://www.insightiitb.org/lokavidya-institutes-innovations/> as One of the projects initiated by IIT Bombay for COVID-19 mitigation (see pages 18 and 73-77)

<sup>a</sup><https://www.cse.iitb.ac.in/~ganesh/HiQuAD/clqg/> described in our ACL 2019 paper

<sup>b</sup><https://github.com/vishwajeet93/clqg>

<sup>c</sup><https://play.google.com/store/apps/details?id=com.mobile.ict.cart2>

<sup>d</sup><https://play.google.com/store/apps/details?id=admin.lokacart.ict.mobile.com.adminapp3>

<sup>e</sup><https://play.google.com/store/apps/details?id=com.mobile.ict.lokacartplus2>

<sup>f</sup><http://bit.ly/lokacart-deck>

<sup>g</sup><https://www.moneycontrol.com/news/business/startup/a-look-at-how-iit-bombay-developed-app-lokacart-is-looking-to-solve-inventory-woes-of-small-businesses-and-farmers-5759851.html>

<sup>h</sup><https://www.ircc.iitb.ac.in/IRCC-Webpage/rnd/ProjectsInitiatedByIITBombayForCOVID-19Mitigation.pdf>

Optical Character Recognition (OCR) is the process of converting the document images into an editable electronic format. This has many advantages like data compression, enabling search or edit options in the images/text, and creating the database for other applications like Machine Translation, Speech Recognition, and enhancing dictionaries and language models. OCR in Indian Languages is quite challenging due to richness in inflections.

Using Open Source and Commercial OCR systems, we have observed the Word Error Rates (WER) of around 20-50% on typewriter printed documents according to our experiments. Also, developing a highly accurate OCR system with an accuracy as high as 90% is not useful unless aided by the mechanism to identify errors. For Error Detection and Corrections in Indic-OCR, we have outperformed state-of-the-art for languages with varied inflections and have solved the Out of Vocabulary problem for Error Correction in Indic-OCR. Please find link to code, demo videos, etc at <https://www.cse.iitb.ac.in/~ocr/>. The link to download our framework is also available therein, including other details such as associated publications.

**Distant Supervision and Multi Instance Multi Label Learning** 2013 – Present

The problem of multi-instance multi-label learning (MIML), an extremely frequent problem in machine learning requires a bag of instances to be assigned a set of labels most relevant to the bag as a whole. The MIML problem finds numerous other applications in machine learning, computer vision, and natural language processing settings where only partial or distant supervision is available. As a specific case, the label/class (eg: "sports") that is assigned to an object such as document/image/video is triggered by some specific segments of that object (eg: performance of India at the olympics). And in general, there are multiple labels associated with a single object. Further, the set of labels could be structurally correlated (such as the "open directory" or Wikipedia's hierarchical categories). We have looked into frameworks for interactively learning models for document classification with topic hierarchies and under MIML settings. Further, we have also looked at summarizing document collections through topic hierarchies, with the additional requirement that summarization has additional requirements such as diversity and coverage. An example application is the automatic generation of Wikipedia disambiguation pages (see our ACL '15, CIKM '16 papers). A recently popular instance of the MIML problem is that of relation extraction using distant supervision. We have looked at various novel models for relation extraction under distant supervision, including inducing explainable rules, incorporating world knowledge into distant supervision based inference and most recently, optimizing the F1 multivariate performance measure, that is of actual interested in real world settings (see our EMLNP '14 and NAACL '15 papers). Specifically, we have developed novel methods for optimizing multivariate performance measures in the MIML setting that use novel plug-in techniques and offer seamless ways to optimize a vast variety of performance measures such as macro and micro-F measure, average precision, etc which are performance measures of choice in multi-label learning domains (see our AAAI '17 paper for more details). Across a diverse range of benchmark tasks, ranging from relation extraction to text categorization and scene classification, it offers superior performance as compared to state of the art methods designed specifically for these tasks. Secondly, it operates with significantly reduced running times as compared to other methods, often by an order of magnitude.

Existing Statistical Machine Translation (SMT) systems have high coverage but very less accuracy. Also, SMT system does not have expressiveness and explainability. As publishable quality machine translation becomes more and more important for generating high-quality data for resource scarce languages, machine translation rules have become increasingly important. Developing high quality translation rules is primarily a manual process, requires lots of manual effort and time, and is prone to errors. Each step in machine translation has different set of rules. At each step erroneous set of rules have to be examined and corrected. Provenance of a sentence explains how a sentence is transformed into an intermediate representation and which rules have been applied to it. As such provenance can help in understanding, debugging and addition of new translation rules. This project aims to build a highly interactive user-friendly rule based machine translation system which will act as a toolkit for human translators. This system will enable domain experts to modify/add new translation rules in order to improve translation quality. This system will also provide an intuitive provenance of a particular translation output generated by the system.

Translation systems are known to benefit from the availability of a bilingual lexicon for a domain of interest. A system, aiming to build such a lexicon from source language corpus, often requires human assistance and is confronted by conflicting requirements of minimizing human translation effort while improving the translation quality. We have developed discrete optimization methods that exploit redundancy in the source corpus and extract recurring patterns which are: frequent, syntactically well-formed, and provide maximum corpus coverage. The patterns generalize over phrases and word types. Our interactive framework leverages these patterns in translation and post-editing, thus enabling machine assisted human translation (see <https://www.cse.iitb.ac.in/~pmt/usage.html> for snapshots).

#### **Typed Query Searches over Enterprise Data (Funded by IRCC) July 2011 – Present**

The goal of this project is to enable queries on enterprise data and facilitate retrieval of precise information. There are several challenges in high precision intranet search, in contrast to the standard internet search. Classical Information Retrieval has involved 40+ years of work in designing better models such as Vector space models, Binary independence models, Network models, Logistic regression models, Bayesian inference models, and Hyperlink retrieval models. The analysis (parsing, tokenization, apply synonyms) on queries and documents have remained fairly simple while complex and intricate ranking algorithms have come to the rescue to do all the heavy lifting. They have been largely characterized by a single formula that uses document and term statistics to capture the numerous factors that determine relevance in terms of **search intent** (real meaning of the search query), **match location** (which part of a document are we matching (title, body, footer, heading, bold text, *etc.*)), **match precision** (nature of the match between query terms and document text (identical, partial match, approximate match, *etc.*)) and **overall importance of document/page** (is the document the landing page of a site?, is it heavily linked to from other important pages?, when was the page last changed? *etc.*). What was the price paid for all this in terms of say, adapting them to searching intranets such as IITBs? These systems are complex and monolithic, making it difficult for non-experts too understand and maintain them. Current versions of off-the-shelf search systems are opaque – they have no mechanisms to customize by providing domain knowledge. And most importantly, they can be unstable: answers to the same query change dramatically as the underlying collection changes

Achieving high search result quality in an intranet is not a one-shot install configure run task. It will requires continuous monitoring and customization to respond to new data, new users, and new queries. The challenge is: How does one design a search system that is amenable to such Search Quality Management? The key principles behind the new search system will be: (a) Deeper analysis of queries and documents (b) Deeper analysis of documents to judge quality and extract index terms (c) Deeper analysis of queries to ascertain user intent (d) Transparent Rule-driven Relevance Computation (e) Many rules built into the system, some exposed for customization. And the advantages it should offer are: (i) Explainability (ii) Know precisely why every result item is being brought back (iii) Dependability/Integrity (iv) As underlying data increases, top-quality results for existing queries continue to show up (not skewed by changes in underlying statistics) (v) Maintainability/Debuggability (vi) Search logic is guided by explicit rules as opposed to weighting and scoring functions.

**Real time contact tracing through Corontine application : Geofenced Tracking, Tracing and Alerting (<https://corontine.in/> )** *2020 – Present*

This app was used successfully in states such as Meghalaya<sup>a</sup> and part of the solution was adapted in Orissa as well as Mumbai. In fact, in Mumbai, we set up a large volunteer team from the campus and outside (of the order of 30-40) to help MCGM track patients and their status. Also see this Insight writeup<sup>b</sup> on the same, as well as articles on several other news sites<sup>c</sup> in addition to a detailed wiki<sup>d</sup>.

Further, on request from the Niti Ayog, AI/ML based algorithms were developed for detection of (anomalous) red zone violation at Telecom Service Providers (TSPs) based on their internal VLR/CDR logs<sup>e</sup>. The anomaly detection model/code<sup>f</sup> was shared with individual TSPs. Further, corintine has been integrated with IoT<sup>g</sup>.

**A Platform for Cross Lingual and Multilingual Event Monitoring in Indian Languages** *2017 – 2020*

This is a project funded by MHRD's IMPRINT initiative and I was a co-PI with the main PI being Prof. Sudeshna Sarkar at IIT Kharagpur. The project aimed to build a cross-lingual, multi-lingual, multi-source and multi-domain platform for information extraction (IE) in Indian languages (ILs) to extract and classify events of interest from news media and to demonstrate the application in a few domains.

**CLIA: Cross Lingual Information Access (Funded by DIT)** *May 2011 – 2016*

The Cross Lingual Information Access (CLIA) project<sup>h</sup>, funded by the Ministry of IT, aims to develop a cross lingual information retrieval system (initially focused on the tourism domain) for 6 languages- Bangla, Telugu, Tamil, Marathi, Punjabi, Hindi and English. This project is run by a consortium of 12 institutions with IITB in the lead. I contributed to the advancement of algorithms, system reorganization and scalability issues for CLIA since May 2011.

**BET: A Workbench for Inductive Logic Programming (Funded by Seed Grant)**

*June 2009 – Present*

With contributions from Ashwin Srinivasan<sup>i</sup>, the author of Aleph<sup>j</sup> and two MTech students at IIT Bombay, we have developed an open source workbench for Inductive Logic Programming, called BET<sup>k</sup>. BET has been developed along the lines of the Weka<sup>l</sup> machine learning tool box. Along with two of my students, the workbench includes implementations of two papers recently co-authored by me on optimally and efficiently learning feature conjunctions for classification models (ICML 2011 and AAAI 2012). Some of the packages also work as plugins with Weka.

**Curating and Searching the Annotated Web: CSAW (Funded by HP Labs)** *April 2008 – Present*

This work was done in collaboration with Prof. Soumen Chakrabarti and several Masters students, results in publications in KDD 2009, SIGIR 2009, WWW 2011 and WWW 2012. This project investigates the use of new models and algorithms for: (i) annotating unstructured Web text with links to uniquely identified entities in an internet encyclopedia such as the Wikipedia and (ii) efficiently indexing and accurately searching annotated document collections. This project was funded by HP Labs.

**UIMA Annotators for Indian Languages (Funded by IBM Research)** *August 2007 – 2009*

This work was done in collaboration with Prof. Pushpak Bhattacharyya. The work was funded by the IBM UIMA grant for developing UIMA annotators for Indian Languages. We supervised an MTP student who built a set of UIMA annotators for Indian languages.

**eDiscovery: Risk management and Compliance Software (at IBM Research)** *Jan 2008 – Dec 2008*

Financial institutions face increased scrutiny and potentially devastating legal and financial liability in connection with their management, use and disclosure of electronic data. IBM solutions for Enterprise Content Management (ECM) address this issue by providing a single platform that delivers integrated policy, process and evidence collection, and retention capabilities. This functionality can be woven into a current application fabric. A single investment can be spread across many regulations, reducing costs and maintenance while providing an integrated information repository and reporting framework.

<sup>a</sup> <https://nenow.in/north-east-news/meghalaya/no-fresh-covid19-case-meghalaya-uses-app-to-monitor-persons-in-home-quarantine.html>

<sup>b</sup> <https://www.insightiitb.org/the-corontine-platform-institutes-innovations/>

*Projects  
Undertaken  
Continued*

**SystemText for Information Extraction (at IBM Research)** Jan 2007 – November 2008

This system enables text-centered enterprise applications by extracting structured information from unstructured text. Unlike previous systems for information extraction, System Text incorporates AQL, a declarative rule language that makes it easy to express precise specifications for complex patterns in text. Thanks to System Text's sophisticated, cost-based optimizer, these complex rules can run on enterprise-scale workloads with minimal hardware. System Text technology provides state-of-the-art information extraction for Lotus Notes, Live Text, IBM OmniFind, Personal Email Search (also available at alphaWorks), and several forthcoming IBM products. This release of the System Text for Information Extraction includes the Development Environment component, which provides support for building and testing extraction rules in AQL. Also included are example rules in AQL, as well as documentation for the rule language and the Development Environment. More details can be found at <http://www.alphaworks.ibm.com/tech/systemt>. I have been collaborating with the IBM Team till date in developing open source rule learners in the SystemT framework (EMNLP 2012).

*Companies  
Co-founded:*

<https://icall.co.in>: **Product from Ethical Soft chain Pvt Ltd, incubated at SINE, IIT Bombay** 2020

Through icall, we are trying to address the need for a video conferencing solution, that is highly customizable, provisioning for segment-specific needs and accessibility even in remote areas while being completely device friendly. Over and above all of these, the solution is completely made in India.

**INNOVATION-HUB: Company being incubated at SINE, IIT Bombay** 2020

Innovation Hub aims to have a large scale impact on on the educational industrial eco-system, in India. It focuses on Connecting the Innovation Ecosystem of India comprising Students, Experts (Faculty & Retired Professional), Institutes Industry by way of a social network and apps that enable a lifelong access to each others works, needs and capabilities.



*Awards and  
Honors*

## **Institute Chair Professorship**

*18 March 2021-present*

<https://drive.google.com/file/d/1gTg94cycFW0mPVhVnpWZJh0cShpjXPk-/view?usp=sharing>

## **Dr. P.K. Patwardhan Technology Development Award 2020 for Video Analytics for Security and Compliance Applications**

Acceptance talk: <https://youtu.be/Lx1GM7mH418?t=3068>, Director's FB Post:

<https://www.facebook.com/directoriitbombay/posts/1317747405353883>

## **Certificate of Appreciation from IIT Bombay for Exceptional Contribution at Significant Personal Risk in the Fight Against Covid-19**

*26th January 2021*

[https://drive.google.com/file/d/1ulXCU5GRZqxxj1\\_JfARD-LPx\\_xiKJLY/view?usp=sharing](https://drive.google.com/file/d/1ulXCU5GRZqxxj1_JfARD-LPx_xiKJLY/view?usp=sharing)

## **IIT Bombay impactful research award 2017 for 'Development of an adaptive framework for end-to-end corrections in Indic OCR' (October 2018)**

<https://rnd.iitb.ac.in/awards/prof-ganesh-ramakrishnan-2017>

## **Certificate of Appreciation from ATS Mumbai for deploying tools that use AI for Data Analytics (October 2017)**

<https://drive.google.com/file/d/1qg7vhY9bp4x1PDaD1kCt8jw4Ua8RjHun/view?usp=sharing>

## **Phd student Dr. Ramakrishna Bairi received Excellence in Ph.D. Research Award 2018**

<https://www.iitbmonash.org/news/excellence-in-phd-research/>

## **J.R. Isaac Chair**

*2014-2016*

[https://www.alumni.iitb.ac.in/sites/www.alumni.iitb.ac.in/files/article/files/Fund\%20Utilization\%20Report\%202014-2015\\_0.pdf](https://www.alumni.iitb.ac.in/sites/www.alumni.iitb.ac.in/files/article/files/Fund\%20Utilization\%20Report\%202014-2015_0.pdf)

## **IBM Faculty Award (2011)**

[http://download.boulder.ibm.com/ibmdl/pub/software/dw/university/facultyawards/2011\\_faculty\\_recipients.pdf](http://download.boulder.ibm.com/ibmdl/pub/software/dw/university/facultyawards/2011_faculty_recipients.pdf)

## **Awards from Qualcomm, Microsoft, Adobe, etc.**

[https://drive.google.com/drive/folders/12Yzh5JBI8F3Zh\\_xCKQSh-YouLcrBFwRS?usp=sharing](https://drive.google.com/drive/folders/12Yzh5JBI8F3Zh_xCKQSh-YouLcrBFwRS?usp=sharing)

## **Computerworld Horizon Awards 2006 Honoree**

<https://www.computerworld.com/article/2546738/computerworld-horizon-awards-2006-honorees.html?page=3>

## **Best Application Paper Award**

*Word Sense Disambiguation using Inductive Logic Programming<sup>a</sup>*

Best Application Paper Award, for Word Sense Disambiguation using Inductive Logic Programming, at The 16th International Conference on Inductive Logic Programming, ILP 2006, Santiago, Spain, subsequently invited to and published at the Machine Learning Journal, Volume 76(1): 109-136 (2009)

## **Best Demo Paper Award**

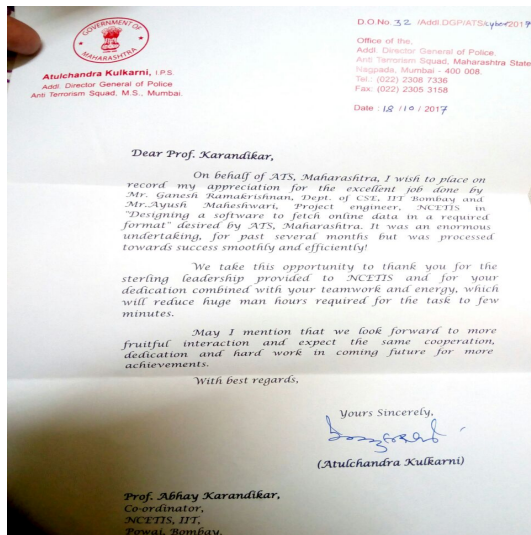
*Anomaly Detection in Surveillance Videos<sup>b</sup>*

Sukalyan Bhakat, Ganesh Ramakrishnan, In Proceedings of the ACM India Joint International Conference on Data Science and Management of Data, CoDS-COMAD 2019, Kolkata, India.

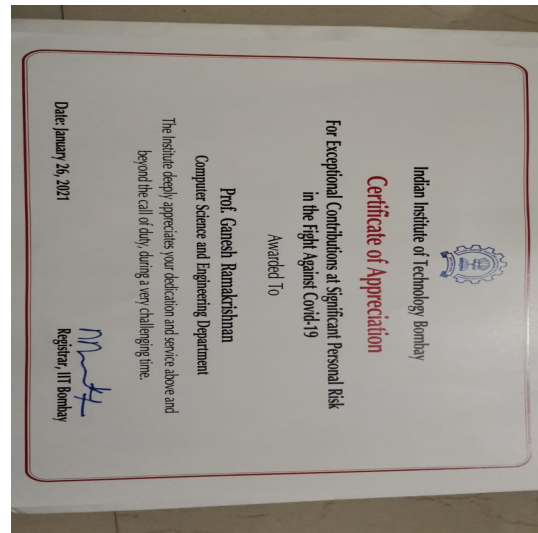
**Identified as an IBM Top Talent (within the entire organization, research inclusive)**

## **BM Bravo! awards**

*I was awarded 4 IBM Bravo awards.*



(a) Certificate of Appreciation from ATS Mumbai



(b) Certificate of Appreciation from for Covid-19 Contributions from IIT Bombay

Figure 1: Certificates of Appreciation

Software &  
Data Con-  
tributions

**Data Efficient Machine Learning**

<https://decile.org/>

**Video Surveillance Analytics for National Centre for Excellence in Technology for Internal Security (NCETIS: <http://www.ncetis.iitb.ac.in>)**

<https://www.cse.iitb.ac.in/~ganesh/videosurveillanceanalytics/>

**Open Source Software for OCR, Post-editing Output, Versioning, etc**

<https://www.cse.iitb.ac.in/~ocr/>

**Multi-Modal and Multilingual Data Analytics**

<https://www.cse.iitb.ac.in/~malta/> & <https://rudder-2021.github.io/>

**Knowledge-Augmented Visual Question Answering**

<https://s3vqa.github.io/>

**Udaan: Breaking the Language Barrier in Education**

<https://www.udaanproject.org/>

**Analytic Hierarchical Process (AHP) for CTARA**

<http://10.129.141.100:8080/AHP/login.html>

**System Text for Information Extraction**

<http://www.alphaworks.ibm.com/tech/systemt/requirements>

**IBM OmniFind Personal E-mail Search**

<http://www.alphaworks.ibm.com/tech/emailsearch>

## Design, Development and Testing of a Novel Digital Learning Platform for Skilling Rural India

January 2015 – Present

This project is focused on designing an online digital platform which can assist in creating an ecosystem for local livelihood generation. We aim to create ecosystems that would help create a pull mechanism (by mapping of local demand) to promote online training. A long term vision is also to help sustain and preserve the traditional knowledge and skill sets on this platform. Empowering the people and inspiring them to create their own skill development module will be the penultimate success of the platform. A summary presentation can be found at <https://docs.google.com/presentation/d/10G-1am5VtjC8cC9U6zAaf4tS8X50SysLLxgrxA701U8/edit>. For both rural development initiatives such as at the large number of NGOs in our country and for educational programs such as at CTARA *etc.*, we have developed an ICT tool for **extremely low footprint** audio-visual documentation of best practices. The goal is to evolve it into an equivalent of Wikihowfor the development sector of India.

This tool provides complete interoperability between a smart phone (which allows for quick prototyping) and desktop (wrapper built on open office) with the facility of modular dubbing (from one language to another), localization etc. A sample repository of documentations created using this tool has been created<sup>a</sup> along with *App on the play store*<sup>b</sup>, *Presentation on user stories*<sup>c</sup>, *Presentation on the Desktop Application*<sup>d</sup>, *Further links to the Desktop app*<sup>e</sup>. Starting November 2017, Microsoft Research India has supported this project and are helping align this work with their Sangam training platform.

## ECommerce for Cottage Industries: Lokacart

January 2015 – Present

This is a an ecommerce portal (online market) for cottage industries, especially farmers clubs, farmer cooperatives and farmer producer companies. We have both an admin (the cottage industry) side app<sup>f</sup> as well as the client or customer-side app<sup>g</sup>. NABARD acknowledged this contribution from us and funded us for a year to investigate ways of scaling up this effort. In June 2017, we organized a workshop for more than 100 Farmer Producer Organizations (FPO) to help reach out the application. Following this workshop, we have had two master's students study the strengths, weaknesses and dynamics of FPOs. See [https://www.ircc.iitb.ac.in/IRCC-Webpage/rnd/PDF/GlimpseIITBResearch/Nov2017/N\\_335.pdf](https://www.ircc.iitb.ac.in/IRCC-Webpage/rnd/PDF/GlimpseIITBResearch/Nov2017/N_335.pdf) for a summary. A summary presentation is also available<sup>h</sup>

## Skill development in Biogas plant construction through Audio-visual tutorials and assessments

January 2015 – Present

Beginning early 2015, we started creating extremely detailed audio-visual tutorials on how to procure, construct, maintain and troubleshoot (2 cubic metre) ferrocement biogas plants<sup>i</sup>. Mid of 2016, we piloted a first biogas plant in Wada Taluka, Palghar district, constructed only by watching the tutorials. We have now initiated a continuous training<sup>j</sup> and assessment program.

## SchemePravaha for Workflow management of Schemes at Tribal project offices

January 2015 – Present

Developed and deployed a system (both computer based and App based) for Workflow management of Schemes at Tribal project offices<sup>k</sup>. The goal here is to significantly reduce the paperwork involved in (a) applying for, (b) approving, (c) sanctioning and (d) following up on schemes by people in tribal areas. We overcame several database (merging of several identification numbers including Adhar id etc) and system design issues (including some facilities on normal phones, most on smart phones as well as on computer).

The APP is available on Playstore<sup>l</sup> but will need approval of the cell phone number into the workflow system for a particular location.

<sup>a</sup><https://drive.google.com/folderview?id=0B4JG1o0uAKt0VTE5SGhUMH1YUU0> and [http://spoken-tutorial.org/tutorial-search/?search\\_foss=Biogas+Plant&search\\_language=](http://spoken-tutorial.org/tutorial-search/?search_foss=Biogas+Plant&search_language=) and the OLI (online learning initiative) of NSS IIT Bombay, <https://www.youtube.com/channel/UCiRgHJGPSMsMI3hcQKrm4yA>

<sup>b</sup><https://play.google.com/store/apps/details?id=com.iitb.mobileict.lokavidya>

<sup>c</sup>[https://docs.google.com/presentation/d/1pAh1D8AaA-nzf499kCpHCBP0s7GRJBiksXTughc3JhaM/edit#slide=id.g11e3a68c94\\_0\\_33](https://docs.google.com/presentation/d/1pAh1D8AaA-nzf499kCpHCBP0s7GRJBiksXTughc3JhaM/edit#slide=id.g11e3a68c94_0_33)

<sup>d</sup>[https://docs.google.com/presentation/d/1TYNDe8I5kQcT1MgZK1Skv2aM2k1wgjV6kkzLuPIUhwk/edit#slide=id.g141df2d195\\_2\\_59](https://docs.google.com/presentation/d/1TYNDe8I5kQcT1MgZK1Skv2aM2k1wgjV6kkzLuPIUhwk/edit#slide=id.g141df2d195_2_59)

**IOPEs: IBM OmniFind Personal E-mail Search (at IBM Research)** *Jan 2006 – Dec 2007*

Simple keyword or text search is not always effective for quickly finding what you need. IBM has gone beyond keywords by inventing a fast and accurate semantic search system for personal e-mail. IBM OmniFind Personal E-mail Search enables semantic searching by extracting and organizing concepts and relationships from personal e-mail. Any business e-mail user who must search in order to accomplish a business purpose will find this tool invaluable. Customization of semantic concepts and the ability to share these concepts with colleagues make this tool especially useful for large enterprise customers. More details can be found at <http://www.alphaworks.ibm.com/tech/emailsearch>.

**Mining Conversational Patterns (at IBM Research)** *April 2008 – Dec 2008*

In this work, we address the problem of extracting important (and unimportant) discourse patterns from call center conversations. Call centers provide dialog based calling-in support for customers to address their queries, requests and complaints. A Call center is the direct interface between an organization and its customers and it is important to capture the voice-of-customer by gathering insights into the customer experience. We have observed that the calls received at a call center contain segments within them that follow specific patterns that are typical of the issue being addressed in the call. We present methods to extract such patterns from the calls. We show that by aggregating over a few hundred calls, specific discourse patterns begin to emerge for each class of calls. Further, we show that such discourse patterns are useful for classifying calls and for identifying parts of the calls that provide insights into customer behavior.

**Study of ILP Procedures for Compact Feature Construction** *Jan 2007 – Present*

We conceptually view the task as searching through subsets of all possible features that can be constructed by the ILP system. Clearly an exhaustive search of such a space will usually be intractable. We resort instead to a randomised local search which repeatedly constructs randomly (but non-uniformly) a subset of features and then performs a greedy local search starting from this subset. The number of possible features usually prohibits an enumeration of all local moves. Consequently, the next move in the search-space is guided by the errors made by the model constructed using the current set of features. This can be seen as sampling non-uniformly from the set of all possible local moves, with a view of selecting only those capable of improving performance. The result is a procedure in which a feature subset is initially generated in the pre-processing style, but further alterations are guided actively by actual model predictions. We test this procedure on language processing task of word-sense disambiguation. Good models have previously been obtained for this task using an SVM in conjunction with ILP features constructed in the pre-processing style. Our results show an improvement on these previous results: predictive accuracies are usually higher, and substantially fewer features are needed.

**COBRA: Corporate Brand and Reputation Analysis** *August 2007 – August 2008*

At IRL, we owned and developed the sentiment analysis capabilities within the solution called COBRA (Corporate Brand and Reputation Analysis). More information can be obtained at <http://domino.research.ibm.com/odis/odis.nsf/pages/solution.18.html>.

### Scalable Systems for Information Extraction

Jan 2006 – Dec 2008

We proposed a new paradigm for rule-based entity annotation, that operates on the inverted index of a document collection and achieves an order of magnitude speed-up over the document-based counterpart. In addition the index based approach permits collection level optimization of the order of index operations required for the annotation process. We develop a polynomial time algorithm that, based on estimated cost, can optimally select between different logically equivalent evaluation plans for a given rule. Additionally, we prove that this problem becomes NP-hard when the optimization has to be performed over multiple rules and provide effective heuristics for handling this case. Our empirical evaluations show a speed-up factor upto 2 over the baseline system without optimizations. We developed RAD: a tool for Rapid Annotator Development on a document collection. RAD translates entity annotation rules into equivalent operations on the inverted index of the collection, to directly generate an annotation index (which can be used in search applications). To make the framework scalable, we use an industrial strength indexer, Lucene and introduce some modifications to its API. The index also serves as a suitable representation for making quick comparisons with an indexed ground truth of annotations on the same collection to evaluate precision and recall of the annotations. RAD achieves at least an order of magnitude speedup over the standard approach of annotating a document-at-a-time as adopted by GATE. The speedup factor increases with increase in the size of the collection, making RAD scalable. We cache intermediate results from the index operations, enabling quick update of the annotation index as well as speedy evaluation when rules are modified. This makes RAD suitable for rapid and interactive development of annotators.

### Manthun: Churning out Information from Text

Jan 2005 – Dec 2006

The project Manthun aimed at analyzing natural language text for information extraction (IE). The task of information extraction involves identification of entities (such as organizations, places, and people) and relationships among entities (such as ). In this project we focused on problems that are central to a large-scale adoption of information extraction in practice, such as (1) Automatic discovery and engineering of effective features using disparate knowledge sources for relationship extraction, (2) Development and organization of rules for the named entity and relationship extraction task, and (3) A scaleable and efficient framework for the IE task. The techniques employed included Inductive Logic Programming, use of Inverted Indices for matching regular expressions, Ripple Down Rules and techniques from data mining.

### Electronic Trigger Alert Program (eTAP)

June 2004 – Dec 2004

This project involved automated generation of sales leads for proactive marketing. Example categories for which sales leads generation was targeted were *leadership change*, *company mergers*, *revenue assets*, *change in location*, *product launches*, etc. A prototype system for eTAP was developed.

### Question Answering: Searching on entities and relations

June 2002 – Dec 2004

In this project, we posed the task of Question Answering as a search over entities and relations. Techniques that were employed included Bayesian inferencing over the lexical relations over entities in the WordNet, learning scoring models that score based on entity and relation affinities, identifying answer types within questions as well as mapping question prefix trees to answer types, etc. This project involved participation in the QA Track at TREC 2003 and resulted in four conference papers and a book chapter.

### Design of slides for Prof. Soumen Chakrabarti's book - 'Mining the Web' Jan - May, 2003

Refer site : <http://www.cse.iitb.ac.in/~soumen/mining-the-web/>

*Challenges  
Hosted and  
Competition  
Participa-  
tion*

**S3VQA: Select, Substitute, Search - A New Benchmark for Knowledge-Augmented Visual Question Answering, 2021**

*Aman Jain, Mayank Kothyari, Vishwajeet Kumar, Preethi Jyothi, Ganesh Ramakrishnan, Soumen Chakrabarti*

<https://s3vqa.github.io/> and <https://arxiv.org/abs/2103.05568>

**RUDDER: A New Benchmark for cRoss lingUal viDeo anD tExt Retrieval, 2021**

*Jayaprakash A, Abhishek, Rishabh Dabral, Ganesh Ramakrishnan, Preethi Jyothi*

<https://rudder-2021.github.io/> and <https://arxiv.org/abs/2103.05457>

**VISIOCITY: A New Benchmark and Evaluation Framework for Summarization, 2020**

*Vishal Kaushal, Suraj Kothawade, Anshul Tomar, Rishabh Iyer, Ganesh Ramakrishnan*

<https://visiocity.github.io/>

**Multilingual PostOCR Competetion at ICDAR, 2019**

*Satyabrata Behera, Upasana Gaikwad, Ramakrishna Bairi, Ganesh Ramakrishnan*

Our team 'CLAM' secured 2nd position. Our model achieved highest corrections of 44% in Finnish, which is significantly higher than overall topper (8% in Finnish). See <https://drive.google.com/file/d/15mxN0-M9PiXBnffi7M0a8wUw33nj1xBp/view>

**NAACL Semeval Word Sense Induction Competition, 2013**

*Satyabrata Behera, Upasana Gaikwad, Ramakrishna Bairi, Ganesh Ramakrishnan*

Ranked 2<sup>nd</sup>. See <https://www.aclweb.org/anthology/S13-2037/>

**Opportunity Activity Recognition Challenge, 2011**

*Naveen Nair, Ganesh Ramakrishnan, Shonali Krishnaswamy*

Ranked 1<sup>st</sup> in Task C

**ACL Semeval Competition, 2007**

*Lucia Specia, Ashwin Srinivasan, Ganesh Ramakrishnan and Maria das Gracas Volpe Nunes*

Ranked 4<sup>th</sup>

**Question Answering: Text Retrieval Conference, 2003 (The QA Track)**

*Ganesh Ramakrishnan, Deepa Paranjpe, Soumen Chakrabarti*

Ranked 4<sup>th</sup>

**ACL Senseval Competition, 2004**

*Ganesh Ramakrishnan, B Prithviraj and Pushpak Bhattacharyya*