Yash Shah

🌐 ys1998.github.io · ✉ yashshah2398@gmail.com · ⌨ ys1998

<u>Research Interests</u>: Natural Language Processing, Speech Recognition, Machine Learning

EDUCATION

Indian Institute of Technology, Bombay July 2016 - Present

B.Tech in Computer Science & Engineering with *Honors*
Minor in Applied Statistics and Informatics
Major GPA: 9.24/10
Thesis: Active Adversarial Accent Adaptation (supervised by Prof. Preethi Jyothi)

INTERNSHIPS

Verisk | AI May 2019 - July 2019
Research Intern under Dr. Maneesh Singh Hyderabad, India
Project: Accent Adaptation for End-To-End ASR

Offered a research internship position at Lyrebird AI (now part of Descript) for May - July 2019 in Montreal[1]

SINAPSE Lab, National University of Singapore May 2018 - July 2018
Research Intern under Prof. Alcimar Soares Singapore
Project: 3D Shape Recognition Using Tactile Sensing [code]

PAPERS & PREPRINTS

1. Stem-driven Language Models for Morphologically Rich Languages [paper]
   Yash Shah[†], Ishan Tarunesh[†], Harsh Deshpande, Preethi Jyothi
   (under review at ICASSP 2020)

2. SEG-LM: A Factored-Output Language Model for Morphologically Rich Languages [paper]
   Yash Shah[†], Ishan Tarunesh[†], Preethi Jyothi
   (preprint; was submitted to EMNLP 2019)

   [†] *equal contribution*

RESEARCH EXPERIENCE

Member of Prof. Preethi Jyothi's Computational Speech and Language Technologies (CSALT) Lab at IIT Bombay since January 2018.

Active Adversarial Accent Adaptation | Undergraduate Thesis May 2019 - Present
Advisor: Prof. Preethi Jyothi, IIT Bombay [code]

- Critically analyzed existing architectures for active learning and domain adaptation in end-to-end ASR
- Developed and open-sourced OpenASR-py (based on OpenNMT-py), a minimal toolkit for end-to-end ASR
- Building adversarially trained, accent invariant models for ASR in active and semi-supervised settings
- Currently exploring the use of local and global discriminators for aggregating and distilling accent information

Language Models for Morphologically Rich Languages July 2018 - October 2019
Advisor: Prof. Preethi Jyothi, IIT Bombay [paper-II] [paper-I] [code]

- Worked with morphologically complex languages with low/no labeled resources, such as word segmentations
- Devised a novel, frequency-based unsupervised algorithm to get canonical segmentations of words
- Proposed a factored-output model with jointly learned mixture weights that predicted the next word using word and morpheme-level probability distributions; submitted the work as a short paper to EMNLP 2019
- Concluded that prefix/suffix information was insignificant compared to that of the stem in this setting; proposed simple stem-based LMs borrowing ideas such as multi-task learning (MTL) and mixture of softmaxes (MoS)
- Observed drastic reductions in perplexities (upto 50%) across Hindi, Tamil, Kannada and Finnish as compared to competitive baselines; submitted new findings to the LM track at ICASSP 2020

---

[1]couldn't attend due to unavailability of a short-term work permit

Semi-Supervised Learning for Sequences
Advisors: Prof. Preethi Jyothi, IIT Bombay & Dr. Maneesh Singh, Verisk AI

May 2019 - Present

- Working on extending existing algorithms such as MixUp and MixMatch to seq2seq models
- Contributed to the formulation of an initial hypothesis that extended manifold mixup to seq2seq models via sequence embeddings

Self-Supervised Representation Learning for Raw Audio
Advisor: Prof. Preethi Jyothi, IIT Bombay (course project, CS753)

August 2019 - November 2019

[report] [code]

- Hypothesized that context and order information is sufficient for learning good problem-agnostic features
- Introduced sorting and masking as two self-supervised tasks for enforcing features to be order and context aware
- Achieved $2\times$ accuracy in speaker identification (VCTK) using our pretrained features as compared to random initialization; also analyzed the importance of fine-tuning our pretrained encoder
- Empirically demonstrated that our tasks could potentially replace upto 5 tasks in PASE (Santiago et al. (2019)) without much degradation in performance on speaker identification and phoneme classification (TIMIT)

Exploring Online Algorithms for Causal Bandits
Advisor: Prof. Shivaram K., IIT Bombay (course project, CS747)

August 2019 - November 2019

[report] [code]

- Worked with bandit instances for which structure of the causal graph is known but its distribution is not
- Proposed 2 extensions to the OC-TS algorithm (pruning the Dirichlet distribution and using an empirical estimate) by Sachidananda and Brunskill (2017) for this setting and showed that they performed better than the original
- Introduced Graph ϵ_t-greedy, a novel online-learning algorithm for causal bandits which combined concepts from Bayesian graph inference and ϵ_t-greedy algorithm for conventional multi-armed bandits
- Extensively compared the proposed algorithms with existing ones across different graph topologies (linear, disjoint, random) and demonstrated that Graph ϵ_t-greedy performed best amongst all
- One of the few teams that were immediately awarded **AA** grade in the course solely based on their project

VAEs with Jointly Optimized Latent Dependency Structure
Advisor: Prof. Sunita Sarawagi, IIT Bombay (course project, CS726)

January 2019 - April 2019

[report] [code]

- Provided the first open-source PyTorch implementation of this paper (ICLR 2019) to the best of my knowledge
- Proposed replacing the top-down inference module with a recurrent network to get slightly better empirical results at the expense of latent structure interpretability
- Extended the authors' approach to sequential data following Markov assumption

Modeling 3D Human Dynamics for Motion Forecasting
Advisor: Prof. Arjun Jain, IIT Bombay

January 2019 - April 2019

[report] [code]

- Formulated a latent generative model for motion synthesis which modeled motion using Neural Ordinary Differential Equations (NODE) by decomposing it into a series of bands
- Performed experiments to judge NODE's applicability to the task, and also theoretically extended the notion to other variational recurrent neural network models

Alternate Loss Functions for Neural Language Modeling
Advisor: Prof. Preethi Jyothi, IIT Bombay

January 2018 - April 2018

[report] [code]

- Explored and proposed novel loss functions for leveraging n-gram statistics in language modeling by (statistical to neural) transfer learning and minimizing divergence between distributions
- Suggested simple approaches for decreasing overall training time and storage space (by pruning and sparsifying n-gram distributions) within this setup with minimal effect on performance

RESPONSIBILITIES

- **Teaching Assistant**, Artificial Intelligence and Machine Learning (CS 337) + Lab (CS 335)
  Duration: August 2019 - November 2019, *Instructor*: Prof. Ganesh Ramakrishnan
- **Teaching Assistant**, Abstractions and Paradigms in Programming (CS 152) + Lab (CS 154)
  Duration: January 2018 - April 2018, *Instructor*: Prof. Amitabha Sanyal
- **Department Web Secretary**, Computer Science and Engineering Association
  Duration: April 2017 - April 2018

ACHIEVEMENTS

- Campus Winner at Microsoft's code.fun.do Hackathon '16; finalist at code.fun.do Showcase, held at Microsoft India Development Center (MSIDC), Hyderabad in May 2017
- Received the Institute Academic Prize for academic excellence during the year 2016-17
- Awarded AP grades[2] in Linear Algebra, Quantum Physics & its Applications and Physical Chemistry
- All India Rank 47 in JEE Advanced '16 (out of 150,000 candidates)
  All India Rank 66 in JEE Mains '16 (out of 1.4 million candidates)
  Top 0.1% in Mathematics, AISSCE 2016 (received Certificate of Merit from the Hon'ble HRD Minister of India)
- Qualified for the Indian National Mathematics Olympiad (INMO) 2014, conducted by HBCSE
- Regional Winner at the Shell Junior National Science Awards in 2013
- Recipient of the National Talent Search Scholarship by Govt. of India since 2012

TECHNICAL SKILLS

Python (PyTorch, TensorFlow), C/C++, Bash, MATLAB/Octave, Java, Git, LaTeX

OTHER SELECTED PROJECTS

SPSIM - Superscalar Processor Simulator　　　　　　　　　　July 2018 - November 2018
Advisor: Prof. Bernard Menezes (course project, CS305-CS341)　　　　　　　[report] [code]

- Designed a simulator in C++ for a 7-stage pipelined superscalar processor with dynamic branch prediction and speculative execution, and data, control and branch hazards appropriately dealt with
- Implemented algorithms for register renaming and remapping, out-of-order instruction issuance and in-order instruction graduation

Efficient Digital Grading Environment　　　　　　　　　　July 2018 - November 2018
Advisor: Prof. S. Sudarshan (course project, CS387)　　　　　　　　　　[report] [code]

- Designed a Flutter-based Android app for simplifying the grading and feedback process of examinations for instructors, students and TAs, using PostegreSQL and Java servlets as backend
- Key features included splitting of answersheets according to questions, distribution of splits to assigned TAs, and displaying aggregated grades and corrected answers back to students

3D Shape Recognition using Tactile Sensing　　　　　　　　　May 2018 - July 2018
Advisor: Prof. Alcimar Soares, SINAPSE Lab (NUS)　　　　　　　　　　　[code]

- Developed an algorithm to construct and normalize a point cloud using tactile feedback obtained by palpating an object with a sensor-mounted robotic hand
- Proposed and implemented a surface reconstruction and shape recognition pipeline using a multi-view CNN operating on perspectives of the input point cloud from the enclosing cube's face-centers

Wirespace　　　　　　　　　　　　　　　　　　　　　　July 2017 - November 2017
Advisor: Prof. Kavi Arya (course project, CS251)　　　　　　　　　　　　　[code]

- Developed a Django/Python based utility to facilitate disk-space sharing amongst users over the same network with features such as access/visibility levels, shareable links and time slots
- One of the four projects showcased online on Web and Coding Club, IIT Bombay's Hall of Fame

CheerBot　　　　　　　　　　　　　　　　　　　　　　　　　November 2016
Team CodeX, Microsoft's code.fun.do Hackathon　　　　　　　　　　　　　[code]

- Developed a chatbot to tackle anxiety and depression using Microsoft's Bot Framework that interacted with the user using predefined questions and analyzed sentiment based on the responses
- Implemented keyphrase detection on user responses to motivate them using relevant quotes fetched from the internet using Bing Search APIs
- Adjudged the winner out of 36 submissions made by teams from across the IIT Bombay campus

---

[2]awarded to about top 1% of the class

RELEVANT COURSES

- **Artificial Intelligence & Machine Learning** - Automatic Speech Recognition[3], Foundations of Intelligent and Learning Agents[2], Speech Processing[2,4], Artificial Intelligence and Machine Learning, Advanced Machine Learning, Introduction to Machine Learning[3], Data Analysis and Interpretation
- **Computer Science** - Compilers, Automata Theory, Digital Image Processing, Computer Graphics, Operating Systems, Database and Information Systems, Computer Architecture, Computer Networks, Logic for CS, Design and Analysis of Algorithms, Digital Logic Design, Data Structures and Algorithms, Discrete Structures, Abstractions and Paradigms for Programming, Computer Programming and Utilization
- **Mathematics & Statistics** - Applied Stochastic Processes, Statistical Inference, Introduction to Probability Theory, Introduction to Numerical Analysis, Linear Algebra, Differential Equations, Calculus
- **MOOC** - CNNs for Visual Recognition (CS231n), Deep Learning for NLP (CS224d)

EXTRA-CURRICULAR ACTIVITIES

- Secured 3^{rd} position overall in Web Development General Championship held at IIT Bombay   (2018)
- Selected for the in-semester preparatory basketball camp for 52^{nd} Inter-IIT Sports Meet   (2017)
- Completed a year long course offered by National Sports Organization in Basketball   (2016)

---

[3] courses taken in Fall 2019
[4] offered by Department of Electrical Engineering