





Rakesh Bal

 rakeshbal99@ucla.edu  (424) 440-9976  rakeshbal99  Rakesh Bal  Rakesh Bal

Education

University of California Los Angeles (UCLA)

M.S. in Computer Science | GPA - 3.81/4.0

Indian Institute of Technology (IIT) Kharagpur

B.Tech. in Computer Science and Engineering | GPA - 9.04/10.0

Los Angeles, CA

Sep'21 - Jun'23

Kharagpur, India

Jul'16 - May'20

Work Experience and Internships

Stovell AI Systems

Machine Learning Operations Engineer

- Used Snowflake for data engineering, KubeFlow pipelines to train ML models, and TF/Keras to test various model architectures
- Obtained 5-10% performance lift in marketplace time series forecasting models in multiple datasets and deployed to production

Kanini Software Solutions

DevOps Engineer

- Migrating scalable infrastructure for a healthcare application using Terraform and AWS services EC2, S3, Lambda, & RDS

Scalable Analytics Institute (ScAI), UCLA

Research Assistant, Advisor: Prof. Wei Wang

- Augmented Protein & Molecule contact maps using Diffusion Docking and AlphaFold to DTI models and attained 10% boost
- Experimented with different novel cross-attention and contrastive loss architectures for modeling Drug-Protein Interaction

Amazon Web Services (AWS)

Applied Scientist Intern

- Analysed the performance of Protein LLMs (from Transformers - BERT, T5, GPT) on Drug Target Interaction (DTI) problem
- Obtained 20% & 12% improvement with ProtBERT over two datasets compared to the baseline with significant cost reductions
- Employed large AWS GPU Clusters, EC2, SageMaker for training pipelines; published the work in Amazon ML conference

Goldman Sachs

Software Engineer

- Integrated two internal bug & issue tracking softwares as a full-stack development project and worked on their cloud migration
- Implemented backend (with RESTful APIs) using Java/SpringBoot and frontend using TypeScript/Angular & Redux
- Coordinated with end-users for the entire SDLC; deployed project to production and handled adoption by over **10000** users

Accenture Technology Labs

Research Intern

- Designed stock price prediction models using news articles and knowledge graphs to incorporate real-world domain knowledge
- Applied GCNs with events for real-world stock scenarios leading to performance improvement of 5% over the baseline models

University of California Los Angeles

Graduate Teaching Associate

- Managed and led discussions/office hours for 500 undergrads in CS32 & Chem 20A, totaling over 500 hours in 5 quarters

Publications

Analysing the Extent of Misinformation in Cancer Related Tweets | Rakesh Bal *et al.*

- *14th International AAAI Conference on Web and Social Media (ICWSM 2020)* | **31 citations**

Two-Sided Fairness in Non-Personalised Recommendations | Rakesh Bal *et al.*

- *35th AAAI Conference on AI (AAAI 2021) Student Abstract and Poster Program* | **10 citations**

PGraphDTA: Improving DTI Prediction using Protein Language Models & Contact Maps | Rakesh Bal *et al.*

- *Accepted in 37th Advances in Neural Information Processing Systems (NeurIPS 2023) Workshop on AI for Science*

Relevant Projects

CLIP for Visual Question Answering (VQA)

- Harnessed OpenAI CLIP in VQA models like MCAN and Pythia in both zero-shot and finetune settings with 2% upgrade
- Added Language Driven Semantic Segmentation (LSeg) to pipeline for answering number-based questions in VQA2.0 dataset

Text Graph Convolutional Networks (GCNs)

- Investigated TextGCN by reproducing the model's results and adding new components on 5 different text classification datasets
- Designed new graph construction algorithms and improved the time cost of graph construction in TextGCN model by 5 times

Stereo Vision based NeRF

- Built NeRF framework for rectified stereo vision on **NeRF synthetic dataset**, with superior performance over monocular vision

Skills

Languages: Python, Java, C, C++, C#, JavaScript, TypeScript, R, Go, Matlab, Lisp

Frameworks: PyTorch, Tensorflow, Keras, Pandas, KubeFlow, Sklearn, Snowflake, Lightning, AWS, Azure, GCP, Angular, React, Docker, K8s, Node.js, SpringBoot, Kafka, GraphQL, MySQL, Flask, MongoDB, Git, Bash, Jenkins, Terraform