Rishubh Singh

Lausanne, Switzerland
rishubhsingh135@gmail.com
rishubhsingh.github.io
+41 77 281 20 62

Education

Ph.D. in Computer Science, EPFL — Lausanne, 2023 – present
Laboratory: VILAB, Supervisor: Amir Zamir

M.Tech. in Computer Science, Indian Institute of Technology Delhi — Delhi, 2014 – 2019
Dual degree (B.Tech. + M.Tech.). Thesis supervisor: Sumeet Agarwal.

B.Tech. in Computer Science, Indian Institute of Technology Delhi — Delhi, 2014 – 2019
Exchange semester (5th) at KTH Royal Institute of Technology, Stockholm, Sweden.

Work Experience

Doctoral Assistant & Ph.D. Candidate, EPFL — Lausanne, 2023 – present
My research is broadly focused on multimodal learning, large-scale modeling and training, low-resolution perception and scientific discovery. Currently, I'm working on learning multimodal representations that enable scalable training and reasoning and closed-loop optimization for scientific discovery. I'm also interested in low-resolution perception for robotic control and bringing multimodal learning to edge setups. Findings of my work have been published at NeurIPS'26, ICLR'26, ECCV'24, CVPR'23, and CVPR'22.

Pre-Doctoral Researcher, Google Research India — India, 2020 – 2023
Cognitive Modeling and Machine Learning Group. Worked on robust deep vision models through shape sensitization (CVPR'23) and human-inspired deep learning models for multi-object multi-part scene parsing (CVPR'22). Advised a student from IIIT Hyderabad, leading to a publication at ECCV'24.

Software Engineer – Quantitative, Graviton Research Capital LLP — Gurgaon, 2019 – 2020
Optimised C++ code at the compiler, OS and hardware levels for the fastest processing of market events. Wrote and maintained the backend of four markets traded by the company and single-handedly integrated two new markets into the trading setup.

Master's Thesis, Indian Institute of Technology Delhi — Delhi, 2018 – 2019
Studied structurally and biologically more plausible recurrent networks, proposed the Decay RNN, a biologically more plausible architecture while performing almost on par with LSTMs (ACL SRW 2020). Supervised by Sumeet Agarwal.

Software Engineering Intern, Google — Mountain View, 2018
Designed an API that generates weather customizable images and videos from Street View data given camera specifications and orientation, and parallelised its data generation with Flume for a 10x efficiency gain. The synthesized videos (100+ GB) are used to train and test internal algorithms and ML models at scale.

Software Engineering Intern, Google — Mountain View, 2017
Built a system to filter false on-call escalations and remove noise from the monitoring and feedback system, and improved the ad serving system to reduce lost ad impressions, with direct revenue impact.

Publications

* denotes equal contribution.

OmniDiMM: Bridging 1D, 2D, and 3D with Any-to-Any Multimodal Modeling — **NeurIPS 2026**
J. Toskov*, O. Barbany*, **R. Singh***, J. Sakurai, E. Tarhan, O. F. Kar, R. Bachmann, A. Zadeh, J. Allardice, C. Li, C. Torras, A. Dehghan, A. Zamir
[webpage] [paper] [code]

Multimodality as Supervision: Self-Supervised Specialization to the Test Environment via Multimodality — **ICLR 2026**
K. P. Singh*, A. Garjani*, **R. Singh**, M. U. Khattak, E. Tarhan, J. Toskov, A. Atanov, O. F. Kar, A. Zamir
[webpage] [arxiv] [code]

How Far Can a 1-Pixel Camera Go? Solving Vision Tasks using Photoreceptors and Computationally Designed Morphology — **ECCV 2024**
A. Atanov*, J. Fu*, **R. Singh***, I. Yu, A. Spielberg, A. Zamir
[webpage] [paper]

OLAF: A Plug-and-Play Framework for Enhanced Multi-object Multi-part Scene Parsing **ECCV 2024**

P. Gupta, **R. Singh**, P. Shenoy, R. K. Sarvadevabhatla

[webpage] [paper] [code]

Edges to Shapes to Concepts: Adversarial Augmentation for Robust Vision **CVPR 2023**

A. Tripathi, **R. Singh**, A. Chakraborty, P. Shenoy

FLOAT: Factorized Learning of Object Attributes for Improved Multi-object Multi-part Scene Parsing **CVPR 2022**

R. Singh, P. Gupta, P. Shenoy, R. K. Sarvadevabhatla

[webpage] [arxiv] [code]

How Much Complexity Does an RNN Architecture Need to Learn Syntax-Sensitive Dependencies? **ACL SRW 2020**

G. Bhatt*, H. Bansal*, **R. Singh***, S. Agarwal

Academic Activities and Services

Reviewer: NeurIPS'26, *ICLR'26* (recognized among top 200 reviewers), NeurIPS'25, CVPR'25, ECCV'24, CVPR'24, CVPR'23

Head Teaching Assistant: CS-503 Visual Intelligence, heading the teaching team, lecturing on multimodal modeling, and advising master's students' research projects (EPFL, 2024–present)

Teaching Assistant: CS-500 AI Product Management (EPFL, 2024–present); Artificial Intelligence, Quantum Computing, Parallel Computing, Computer Architecture (IIT Delhi, 2017–2019)

Awards

EPFL EDIC Fellowship for Ph.D. students, 2023
National Rank 36 and fellowship, Kishore Vaigyanik Protsahan Yojana (KVPY) Young Scientist Program, 2013
ACM-ICPC Amritapuri Regional, qualified and ranked 71st in India, 2015
National Top 25, Indian National Olympiad in Informatics (INOI); IOI Training Camp, 2013 and 2014
National Top 40, Indian National Astronomy Olympiad (INAO); Astronomy Orientation Camp, 2014
National Talent Search Examination (NTSE) scholarship, Government of India, 2010
JENESYS India–Japan cultural exchange program, selected among the top 10 students of CISCE schools