Sagalpreet Singh

Pre-Doctoral Researcher, Google DeepMind

Sagalpreet Singh | sagalpreet | Webpage
sagalpreet60@gmail.com | +91-xxxxxx.329

Education

[E1] Indian Institute of Technology Ropar 2019-2023
B.Tech. in Computer Science and Engineering 9.31/10.0
Concentration in Artificial Intelligence 10.0/10.0

Experience

[X1] Google DeepMind Aug 2024 - Present
Pre-Doctoral Researcher - Agents

[X2] Oracle Jun 2023 - Aug 2024
Member of Technical Staff - AI Services (Speech)

[X3] Guvi Apr 2023 - Aug 2024
Teaching Consultant - AI/ML & Python

[X4] Learning Affect and Semantic Image Analysis Group Jan 2023 - Apr 2023
Undergraduate Researcher - IIT Ropar

[X5] Game Theory & ML Lab Mar 2022 - Oct 2022
Undergraduate Researcher - IIT Ropar

[X6] Oracle Jun 2022 - Jul 2022
Software Engineering Intern - Oracle Integrations Cloud

[X7] Image Processing, Security and Analytics Lab Jan 2022 - May 2022
Undergraduate Engineer - IIT Ropar

Publications

C=Conference, S=Submitted

[S1] Dense and Diverse Goal Coverage in Reinforcement Learning
Sagalpreet Singh, Rishi Saket, Aravindan Raghuveer
Google Deepmind Booth @ NeurIPS 2025
AISTATS 2026 (Under Review)

[C1] Learning from Label Proportions and Covariate-shifted Instances
Sagalpreet Singh, Navodita Sharma, Shreyas Havaldar, Rishi Saket, Aravindan Raghuveer UAI 2025

[C2] On Subset Selection of Multiple Humans to Improve Human-AI Team Accuracy
Sagalpreet Singh, Shweta Jain, Shashi Shekhar Jha AAMAS 2023 (Oral)

Awards

#=Declined-by-me

[A1] **Google Research** (USD 2k) and **Microsoft Research**# (INR 120k) travel grants, AAMAS 2023

[A2] Selected for **Amazon ML Summer School** and offered **Applied Scientist internship**#

[A5] **Institute Merit Scholarship** at IIT Ropar for exceptional academic performance

[A3] **Best Poster Presentation Award** at IIT Ropar on Research Scholar's Day 2023

[A4] **Best B.Tech. Project Award** at IIT Ropar on National Technology Day 2022

[A6] **NTSE Scholarship** by the Government of India, awarded to top 1000 students nationwide

Research & Industry Projects

[$]=I am not the project lead

Retrieval for Tool Aware Planning
Oct 2025 - Present

Advisors: Dr. Aravindan Raghuveer, Dr. Satinder Baveja [X1]

Here we ask a simple question - can an agent have an external memory that can be easily updated without having to update the model weights (or not drastically at least)? This line of work is heavily inspired from tool use in LLMs and Hierarchical Reinforcement Learning. Specifically, I am working on developing a tool retrieval agent which can efficiently fetch tools from a large tool library that can be used to plan and solve the task at hand.

Verifiable Problem Discovery[$]
Oct 2025 - Present

Advisors: Dr. Aravindan Raghuveer, Dr. Rahul Madhavan [X1]

This is an extension of Unsaturating Benchmarks where we are trying to autonomously generate verifiable problems in Mathematics and Computer Science domains by using rejection sampling rather than RL. The outcome of this project is data which improves the performance of Gemini models across the board on various benchmarks. This work is on its path to be incorporated into Gemini post-training.

Dense and Diverse Goal Coverage in RL
Jun 2025 - Oct 2025

Advisors: Dr. Rishi Saket, Dr. Aravindan Raghuveer [X1, S1]

To prevent mode collapse in reinforcement learning on multi-goal environments, we proposed a novel optimization objective, and an algorithm with convergence guarantees. We also show improvements in diversity of goal states reached by our algorithm in comparison to SAC and other baselines. The paper is under review and a provisional patent has been filed.

Unsaturating Benchmarks
Jan 2025 - May 2025

Advisors: Dr. Rishi Saket, Dr. Aravindan Raghuveer, Dr. Satinder Baveja [X1]

With continuous self-improvement in mind, we came up with a novel problem generation framework to autonomously generate verifiable problems for LLMs to solve, specifically for planning in PDDL. We observed 60%-points improvement in problem generation efficiency by guiding the problem generation using RL with verifiable rewards to mimic the distribution of problems that the domain specific agent cannot solve.

Hybrid LLP
Aug 2024 - Dec 2024

Advisors: Dr. Rishi Saket, Dr. Aravindan Raghuveer [X1, C1]

Inspired from ad-conversion modeling, the project focused on a scenario where the source domain provided fully supervised data, while the target domain offered weakly supervised data, characterized by an aggregate label for a collection of instances. With the idea of learning domain invariant representations, we proposed a novel loss function with theoretical guarantees that led to SOTA results on real-world datasets. The paper was accepted at UAI 2025 and the work is projected to have a revenue impact in the order of millions of dollars.

Natural Text-to-Speech
Feb 2024 - Aug 2024

Mentors: Mr. Ankit Tyagi, Mr. Phanindra Mankale [X2]

Built Oracle's flagship natural text-to-speech system from scratch, achieving naturalness across several speech datasets. The system is capable of generating high-quality speech in real-time on CPU, and supports voice cloning with just 5 seconds of reference audio. The system also supports SSML tags for fine-grained control over prosody, pitch and speed of generated speech.

Diffusion Models for Audios
Jan 2023 - Apr 2023

Advisors: Dr. Abhinav Dhall [E1, X4]

We proposed two inference-time procedures for domain adaptation in diffusion models - (i) cross-diffusion: for conditional generation of audio on target instrument given the source instrument audio, and (ii) double-diffusion: to produce an unconditional paired set of instrumental audios for chorus generation from models trained on datasets of individual audios.

Human-AI Team
Mar 2022 - Oct 2022

Advisors: Dr. Shweta Jain, Dr. Shashi Shekhar Jha [E1, X5, C2, A3]

Developed an algorithm to intelligently combine multiple human labels with logit outputs from AI models, significantly improving classification accuracy over naïve methods. We particularly look at classification setting in this work. Our algorithm (which approximates failure modes for each human and represents those via a confusion matrix) is a direct outcome of optimizing for the lower bound on accuracy of the combined model. The paper was accepted for oral presentation at AAMAS 2023.

Bring your own Email Server
Jun 2022 - Jul 2022

Mentor: Mr. Renukaradhya Dakshinamurthy [X6]

Designed & implemented circuit breaker for email notification service of Oracle Integrations Cloud to allow increasing the number of emails that customers can send in a rolling 24-hour window by using their own customer tenancy instead of the default infra-tenancy, thereby opening up a monetization opportunity.

SAMPAN Android App Jan 2022 - May 2022

Advisors: Dr. Puneet Goyal [E1, X7, A4]

Co-developed an android app that is used by Anganwadi workers, Child Development Project Officers and other field level implementers to record and analyse data related to children malnourishment. The app has over 1k Play Store downloads and our work was featured in The Tribune. This project was also recognized with the Best B.Tech. project award at IIT Ropar on National Technology Day 2023.

Competitions

India Top 20
Oppo Inspiration Cup

Rank 7
AMEXpert Smartathon

India Top 50
NK Securities ML Hackathon

Rank 84
Codegoda (Agoda)

Rank 198 (18th India)
Google Hash Code

Rank 238
Google Kickstart

Rank 84
Newton's Grand Contest

Rank 10
Inter-IIT CP (IIT-BBS)

Rank 1108 / ∼1M
JEE Advanced

Top 1% (National)
NSEP (IAPT)

Top 1% (State)
NSEJS (IAPT)

Top 0.1% (National)
Senior Secondary Physics Exam

Technical Skills

Language Proficiency	English (TOEFL - 113/120), Hindi (Native), Punjabi (Native)
Programming Languages	Python, C/C++, Bash, RISC-V Assembly, SQL, LaTeX
Development Frameworks	Jax, PyTorch, Tensorflow, FastAPI, JupyterLab, Triton
DevOps Tools	Git, Docker, Kubernetes

Volunteering & Positions of Responsibility

Reviewing
AISTATS 2026, ACL ARR - May 2025

Volunteer
AAMAS 2023

Academic Council Representative
IIT Ropar CSE 2019-23 Batch

Representative
IIT Ropar Coding Club

Problem Setter
Competitive Programming, TechFest, IIT Ropar

Event Organizer
Fury Road, Robotics Club, IIT Ropar

References

Dr. Rishi Saket, Staff Research Scientist, Google Deepmind

Dr. Aravindan Raghuveer, Principal Researcher, Google Deepmind

Dr. Shashi Shekhar Jha, Associate Professor, Indian Institute of Technology Ropar

Dr. Shweta Jain, Assistant Professor, Indian Institute of Technology Ropar

Dr. Abhinav Dhall, Associate Professor, Monash University

Mr. Phanindra Mankale, Director, Oracle

Mr. Ankit Tyagi, Staff Engineer, Oracle