JOYKIRAT SINGH
+91 97175 22342 | joykirat19166@iiitd.ac.in | Website | Linkedin | Github | Google Scholar

EDUCATION

Indraprastha Institute of Information Technology Delhi — **New Delhi, India**
B.Tech in Computer Science and Design (***Silver medalist***); ***GPA: 9.24/10.0*** — *2019 – 2023*

PUBLICATIONS

Frugal LMs Trained to Invoke Symbolic Solvers Achieve Parameter-Efficient Arithmetic Reasoning
Subhabrata Dutta, **Joykirat Singh***, Ishan Pandey*, S. Manchanda, S. Chakrabarti, T. Chakraborty* - **Accepted to AAAI 2024**

EROS: Entity-Driven Controlled Policy Document Summarization
Joykirat Singh*, *Sehban Fazili*, Rohan Jain, Md. Shad Akhtar* - **Accepted to LREC-COLING 2024**

How to think step-by-step: A mechanistic understanding of chain-of-thought reasoning
Subhabrata Dutta, **Joykirat Singh***, S. Chakrabarti, T. Chakraborty* - **Accepted to TMLR**

Mechanistic Behavior Editing of language models
Joykirat Singh*,*Subhabrata Dutta*, T. Chakraborty* - **UnderReview**

Exposing the Achilles' Heel: Evaluating LLMs Ability to Handle Mistakes in Mathematical Reasoning
Joykirat Singh*, Akshay Nambi, Vibhav Vineet* - **ACL Main**

PROMPTWIZARD*: Task-Aware Prompt optimization framework*
*Eshaan Agarwal, **Joykirat Singh**, Vivek Dani, Raghav Magazine, Tanuja Ganu, Akshay Nambi* - **ACL Finding**

Self-Evolved Preference Optimization for Enhancing Mathematical Reasoning in Small Language Models
Joykirat Singh*, Tanmoy Chakraborty, and Akshay Nambi* - **UnderReview**

Agentic Reasoning and Tool Integration for LLMs via Reinforcement Learning
Joykirat Singh*, Raghav Magazine, Yash Pandya, and Akshay Nambi* - **UnderReview**

RESEARCH EXPERIENCE

Microsoft Research — **Bengaluru, India**
Research Fellow working with Principal Researcher Akshay Nambi — *Jul 2024 – Present*

Exposing the Achilles' heel: Evaluating LLMs ability to handle mistake in Mathematical reasoning

- Evaluated several LLMs such as GPT-4o, GPT-4, GPT-3.5Turbo, Claude-Opus and other SLMs such as Llama, Phi, etc. on their ability to detect and correct reasoning mistakes
- Introduced a novel dataset MWP-Mistake, incorporating MWPs with both correct and incorrect reasoning steps generated through rule-based methods and SLMs to perform benchmarking
- Performed a rigorous evaluation of 13 SOTA models on 3 metrics - mistake identification, performance and rectification
- Identified GPT-4o's superior performance in mistake detection and rectification, and challenges faced by smaller models
- Analysed the performance of LMs to identify key issues such as data contamination and memorisation of datasets which impacts the reliability of LLMs in real world MWPs solving abilities

Improving SLMs Math reasoning capabilities through dense reward at each reasoning step

- Utilized MCTS (Monte Carlo Tree Search) at train time to generate dense reward for each reasoning step, eliminating the limitation of various datasets, i.e. lack of reward for each step
- Taught LM to pick the optimal path from set of n-branches at each step using DPO (Direct Policy Optimization) while incorporating external verifiers and using FL for effective pruning, better reward and faster training time
- Trained extra verifiers on dense rewards from MCT, which acts as proxy q-value model to estimate the expected reward of each step at inference time, enabling efficient utilization of inference time computation

PromptWizard*:* Task-Aware Prompt optimization framework

- Conceptualized PromptWizard, a novel, fully automated framework for discrete prompt optimization, a self-evolving and self-adapting mechanism to automate prompt engineering, which is a labor-intensive and domain-specific task
- Introduced a feedback-driven critique-and-synthesis mechanism, refining prompts based on performance insight
- Generated human-readable, task-specific prompts by effectively balancing exploration and exploitation and iteratively refining both prompt instructions and in-context examples; resulted in superior performance across 45 tasks
- Achieved a substantial reduction in number of API calls, token usage, and overall cost; demonstrating the novel solution's efficiency, scalability, and advantages over existing prompt optimization strategies

Laboratory for Computational Social Systems (LCS2), IIT Delhi — **New Delhi, India**
Graduate Research Assistant (part-time) under Prof Tanmoy Chakraborty and Prof Soumen Chakrabarti — *Jun 2023 – Jun 2024*

Frugal LMs trained to invoke symbolic solvers achieve parameter-efficient arithmetic reasoning

- Trained LLMs to function as proficient translators by converting natural language arithmetic queries into formal language (FL) descriptions; finetuned the LLMs using low-rank adapters translate word problems into FL proficiently
- Employed Proximal Policy Optimization (PPO) reinforcement learning to further enhance the adapted LLM, combined with a non-differentiable symbolic solver; achieved +30.65% in accuracy (SVAMP dataset) using GPT-J 6B model over base LMs
- Achieved comparable performance with Vicuna13B (SVAMP and GSM8K dataset) versus GPT3.5 (8-shot prompting)

How to think step-by-step: Reverse engineering Chain-of-Thought (CoT) prompting in Large Language Models (LLMs)

- Studied the internal mechanism of a model that facilitates CoT generation; investigated the neural sub-structures within LLMs that manifest CoT reasoning from a mechanistic point of view
- Analysed Llama-2-7B on multistep reasoning over fictional ontology to demonstrate that LLMs deploy multiple parallel pathways of answer generation for step-by-step reasoning
- Observed a functional rift in the middle layers of the LLM, where token representations in the initial half remain strongly biased towards the pretraining prior, with the in-context prior taking over in the later half

Mechanistic Behavior Editing of language models

- Built a novel task adaptation method, TaRot (Task-aware Rotation) of token association, to resolve data inefficiencies in existing fine-tuning techniques
- Optimised rotation metrics using Baysian optimisation which intervenes in neural circuitries to align the model's head output for task adaptation, being as data-efficient as in-context learning
- Achieved extreme data- and compute-efficiency, using 6-20 supervised examples for each task, improving upon both zero- as well as few-shot performance with average improvements (across models and tasks) of 23.81% and 11.15%, respectively

INDUSTRY WORK EXPERIENCE

Expedia Group **New Delhi, India**

Software Development Engineer, Payment's Module *Aug 2023 – Jun 2024*

- Led the development and implementation of the complete pipeline for integration and smoke tests for all the payment modules using Karate Framework to increase the service's robustness
- Developed and released various internal releases using Kotlin and Java based services, collaborated with other stake holders for seamless integration and provided on-call support during releases
- Implemented Datadog custom tags in payments module using the custom metric API in Kotlin and Java based applications

Software Development Intern, Conversation Module *May – Jul 2022*

- Built a robust Java and Spring-based conversation service application using LanguageIO (Translation API) to facilitate the translation of messages, ensuring production readiness; impacted the processing of 42,000 chats per quarter
- Integrated message translation capabilities within a Kafka service, enabling seamless communication between agents and customers, especially during non-office hours

Google Summer of Code **Remote**

Intern, MusicBlock *May - Aug 2021*

- Created the palette section and the artboard functionality using MVVM architecture for the new application MusicBlock 2.0
- Employed React+Typescript and P5 to make the new canvas, and enhance several UX issues within MusicBlock 1.0

ACADEMIC PROJECTS

EROS: Entity-Driven Controlled Policy Document Summarization *(under Prof. Md Shad Akhtar, IIIT Delhi)* May 2023

- Proposed a novel BART-based entity-driven controlled policy document summarization model (EROS) to comprehend lengthy and complex privacy policy documents which include privacy related entities (e.g. data and medium)
- Developed the model to identify critical entities through a span-based entity extraction model and employ them to control the information content of the summaries using proximal policy optimization (PPO)
- Enhanced the interpretability and readability of policy documents by using controlled abstractive summarization – enforced the generated summaries to include critical privacy-related entities and organization's rationale in collecting those entities
- Curated a PD-Sum, a policy-document summarization dataset with marked privacy-related entity labels to perform a qualitative and human evaluation and establish the efficacy of EROS over other baseline models

Ohilo – a motion based game on Unity *(under Prof Aman Parnami, IIIT Delhi)* Sep 2022

- Led the ideation and implementation of the backend infrastructure for the first motion based Unity game for children – Ohilo, which provided a full-body playful experience
- Utilized a robust tech stack which included Unity engine for game development and AWS services such as EC2, Dynamo DB, SNS, Kinesis, and GameLift to create a scalable and efficient backend infrastructure
- Implemented IOS version of the game, integrated the backend and released Ohilo for beta testing on the app stores

QA generation from health insurance documents for Chistats Labs *(under Prof. Md Shad Akhtar,IIIT Delhi)* Feb 2022

- Developed an end-to-end pipeline for Chistats Labs to extract questions from health insurance documents and generate MCQs for users to gauge their understanding of the insurance document; automated the task to reduce hours of human annotation
- Pre-processed the data and Fine-tuned multiple DL models such as T5, PGN and BERT to extract questions from the insurance documents; leveraged SVM to categorise paragraphs in the documents to build question answer pairs
- Tech stack: Python, PyTorch, Sklearn

ACHIEVEMENTS

- Awarded silver medal for academic excellence at IIITD (2023) for achieving highest CGPA in computer science branch
- Awarded the Chanakya Fellowship by iHub Anubhuti IIITD (2022) to work on an independent research problem
- Presented the paper (top ~2%), Frugal LMs Trained to Invoke Symbolic Solvers Achieve Parameter-Efficient Arithmetic Reasoning at AAAI'24 Conference with an audience of approximately 300 people
- Volunteer, AAAI'24 Conference - received a travel grant to attend the conference
- Appointed Reviewer for ICLR'24 Conference

SKILLS

Programming Languages: Python, Java, C++, C, JavaScript, Swift, Kotlin

Frameworks/Databases: Pytorch, Keras, Tensorflow, Scikit-learn, Pandas, React, Node.js, Unity, Adobe Photoshop, Adobe Illustrator, IOS development, Android Studio, Unreal Engine, Kafka, Spring, HuggingFace, AWS, GRPC, GraphQL

MISCELLANEOUS

- **Head, Open source team, Google Developer Student Club (2021)** - Organized events to promote open source contributions
- **Volunteer, The Citizen Care Foundation (Aug 2023 – Present)** - Led ten drives distributing essentials to the underprivileged