

CURRICULUM VITAE

Kanishk Paul, Kolkata.

RESEARCH INTERESTS

Rigorous evaluation of AI systems; world models and causal induction in interactive agents; multilingual and low-resource NLP; efficient inference under tight compute.

EDUCATION

B.Tech, Computer Science and Engineering (Internet of Things and Cyber Security including Block Chain Technology), Heritage Institute of Technology, Kolkata. Aug 2023 to Jun 2027

National Gems Higher Secondary School, Kolkata. 2008 to 2022

RESEARCH EXPERIENCE

Jadavpur University, Kolkata. Research collaborator, multilingual NLP. Co-authoring a manuscript (in preparation) with Prof. Sudip Naskar and a TU Munich collaborator. Aug 2026 onward

- ☞ Built a reproducible round-trip evaluation (14 Indic languages, two translation-model families, 56,000 translations) measuring whether honorific register survives a pivot through English, on 4-bit local models and free-tier GPUs.
- ☞ Identified a systematic deference floor: IndicTrans2 kept 99% of deferential source register (1,659/1,677) but 6% of less-deferential register (71/1,188) among classified outputs; native-speaker validation of cue tables for 12 of 14 languages is pending.
- ☞ Showed that dedicated MT's near-zero pivot penalty is a floor artefact, since its direct path already drops almost all less-deferential register, while LLM translators lose 20–22 points through their own English pivot.
- ☞ Wrote LLM-free, rule-based register extractors (one cue table per language) and a Unicode normalization policy that fixed four silent under-counting bugs; froze datasets behind hash checks and made every pipeline resumable and seeded.
- ☞ Preregistered predictions before each phase and reported two that failed: the original contextual-fraction predictor did not survive robustness checks, and the target-decoder-prior test failed with a flawed metric; a corrected confirmatory re-test is preregistered.

MANUSCRIPTS IN PREPARATION

Register obstruction, with S. Naskar and a TU Munich collaborator. In preparation.

Mechanic Composition Worlds. In preparation.

RESEARCH PROJECTS

- | | |
|--|------|
| World models for ARC-AGI-3: Virgil and its predecessors, with every score and its rules. | 2026 |
| ButterflyGate: an $O(n \log n)$ token mixer in a pretrained model's early layers. | 2026 |
| ResonatorLM audit: an independent replication. | 2026 |

Earlier: NeuroSymbolic Chain of Thought, Synthesis of Thought, Debate of Thought, a photonic Shor's algorithm design, an MCMC sampler, a fee-less AI-miner blockchain, federated learning with proof-of-work. See History.	2024 to 2025
--	--------------

INDUSTRY AND OTHER

Unify (YC W23). Design partner. Tested a pre-release agent platform and reported reproducible gaps in OAuth permission visibility, learned-memory review, action audit trails and per-action cost attribution.	Aug 2026
--	----------

Heritage Institute of Technology. Undergraduate researcher, humanoid robotics. Wrote and defended a proposal for a low-budget assistive humanoid robot through several faculty review rounds; informal funding commitment of ₹20,000 from the department.	Jul to Oct 2025
---	-----------------

buildspace, Season 5. Built an offline LLM for Android that acts (messages, calls, music) solo in about three weeks.	Jun to Jul 2024
--	-----------------

SELECTED PROJECTS

Blindfold. A browser agent that redacts pages on device and sends only stand-in values through one audited exit. Co-led a six-person SIH 2026 team.

maiku. One runtime for Python functions, local LLMs, PyTorch modules, MCP servers, A2A agents, CLIs and people, through one streaming event contract.

ChromeClaw. A browser agent with one schema-validated action per step, a policy layer and a JSONL trace of every step.

ReasonTrace. A visual debugger that turns agent logs into a seven-event graph.

Sefai. Offline GGUF inference in about 150 lines of Rust.

Also: Casca, a desktop agent runner. Stockfih, Stockfish 18 in the browser. ParaReport, Kolkata civic reports. CuteTuner, a local vocal tuning studio. lifesim, a Monte Carlo life simulator. Horizon, a local decision workstation.

Earlier builds, 2020 to 2025: an answer engine, English to SQL, no-code fine-tuning, an AI music critic, a map you talk to, a speaking coach, flood and drought mapping, a deep-fake detector, a university advisor, a disaster-response map, a lunar game, a multimodal assistant, trading systems. See Work.

VENTURES

OpenMind. CEO.	Dec 2023 to Feb 2024
----------------	----------------------

faculza.	2023
----------	------

NashBase, with a potential co-founder.	Jan to Apr 2026
--	-----------------

A research collective. Wound down.	2026
------------------------------------	------

A college AI research group.	Feb 2024
------------------------------	----------

HACKATHONS

8 hackathon wins.

YEAR	EVENT	FIELD	RESULT	BUILT
2023	SIH 2023, college round		Runner-up, cash prize	Disaster-response platform; led the AI layer
2023	NASA Space Apps			Unity lunar exploration module
2024	SIH 2024, college round		Top of the pitch round	Deepfake image detector
2024	NSHM Innovathon	100+ teams	3rd, behind IIT Kharagpur and IIT Kanpur	University advisor; owned the AI
2024	MLX (GDG On Campus, inter-college, Kaggle)	~90	8th; led both tracks for long stretches	Classification and regression
2025	MLX 2.0 (GDG On Campus, Kaggle)	~50	1st	Classification and regression
2025	National hackathon with investor scouting (up to ₹1 crore)		Cleared the pitch round; 23rd, top 21 advanced	AI music critic, built solo
2025	Smart Bengal Hackathon	~1,100 teams	Finalist (top ~80); about top 20	Speaking coach; led the AI
2025	ISRO Bharatiya Antariksh Hackathon		Did not reach the final	Flood and drought spread
2025	SIH 2025, college qualifier		Won the pitch round; qualified for nationals	City traffic platform; led the team
2025	Hack Heritage 2025	~300 teams	Won the pitch round; selected for SIH nationals	Ecology map platform, built solo in ~12 h
2026	SIH 2026			Blindfold; co-led, most commits
2026	Build with Gemma Hackathon (virtual, 4 Jul)			ParaReport
2026	ML Challenge 2026, Business Entity Resolution	top team 0.991483; 10th 0.990282	0.987978 public macro F0.5 (single run; ensemble best)	Entity matching across 2.2M businesses
2026	ARC Prize 2026 (Kaggle)		Account best 0.23	Blackstar

HONORS AND SELECTIONS

Google × Hugging Face Gemma meetup, 1 of 50.	Apr 2026
DisCo Cohort 001.	2026
PrettiFlow Cohort 2, selected 1 of 6, declined.	2026
Thinking About Thinking Ambassador Programme.	Aug 2026
F Inc Blueprint 2, interviewed.	Sep 2026
Entrepreneurs First Fall 2026, interviewed.	May 2026

Invited as a design partner by a startup from DeepMind founders.

Aug 2026

OPEN SOURCE

unifyai/unify-agent#171, merged.

Aug 2026

nossa-y/activity-frames#47, open.

SKILLS

LANGUAGES: Python, TypeScript, JavaScript, Rust, C++, SQL, Java.

ML AND AI: PyTorch, TensorFlow, Hugging Face Transformers and Hub, scikit-learn, LLMs (inference, fine-tuning, evaluation), RAG, prompt engineering, machine translation (IndicTrans2, NLLB-200, FLORES-200), 4-bit quantization, state-space and post-transformer models (Mamba, RWKV, HRMs), neurosymbolic and symbolic AI, reinforcement learning, computer vision, multi-agent systems, multilingual and low-resource NLP, MLX, llama.cpp, ONNX Runtime.

DATA: NumPy, pandas, SciPy, NetworkX, linear programming, benchmarking, data pipelines.

ENGINEERING: React, Next.js, Tailwind CSS, shadcn/ui, FastAPI, Flask, Django, Node.js, Supabase and PostgreSQL, Tauri, Playwright, Streamlit, Docker, REST APIs, CI/CD, MLOps, AWS and GCP, Git, Linux, pytest, Vitest.

PRACTICE: experimental design, preregistration, benchmark construction, paper reproduction.

CONTACT

kanishkpaul1729@gmail.com, GitHub, LinkedIn.