Suprabhat Rapolu

Hyderabad, India · suprabhat@dirmacs.com · github.com/suprabhatrapolu · linkedin.com/in/suprabhatrapolu · suprabhat.dirmacs.com

RESEARCH INTERESTS

Information governance; misinformation, agenda-driven information and the selective use of evidence; claim verification, provenance and traceability; versioned and time-indexed records, including revisions to official statistics; evaluation and reliability of NLP systems and automated fact-checking; governance of AI agents and the knowledge they act on; how the evidence people are shown shapes what they believe and decide, and who bears the cost of checking it.

I study whether a claim can be traced to what it rests on, and whether that basis is still true, for claims made by people, institutions and AI systems. My work so far measures this in records, combining claim-level audits, NLP evaluation, statistical methodology and provenance-carrying system design. Next, I want to study how what people are shown about evidence changes their beliefs and decisions.

EDUCATION

Postgraduate Diploma in Applied Statistics 2025 – present
Indira Gandhi National Open University, India.

Postgraduate Diploma in Artificial Intelligence and Machine Learning 2025 – 2026
University of Hyderabad, India.

- Capstone: automated fact verification of financial claims, extended into publication [1].

Bachelor of Business Administration, Business Analytics 2021 – 2024
Manipal Academy of Higher Education, India.

RESEARCH EXPERIENCE

A controlled benchmark of financial claim verification 2025 – 2026
University of Hyderabad, as my diploma capstone; extended at Dirmacs, Hyderabad. Completed. Output: [1].

- Question: which factual edits can verifiers trained on FEVER detect in financial text, and does their accuracy come from checking the evidence?
- Methods: 1,489 claim–evidence pairs from SEC 10-K filings with documented single-edit perturbations (number change, negation flip, entity swap, evidence mismatch), each tagged with its generation rule; eleven models across three generations under one protocol (sparse linear: Naive Bayes, Logistic Regression, Linear SVM; tree and voting ensembles incl. Random Forest and XGBoost; neural: BiLSTM with GloVe, fine-tuned DistilBERT; plus a tuned lexical-overlap baseline); a hypothesis-only probe for annotation artifacts; transfer to real financial claims (Fin-Fact); Wilson intervals throughout.
- Finding: the strongest verifier scores 98.8% on verbatim claims but 33% or below on single-token falsifications, with entity swaps hardest at 2.6%; on real financial claims no learned model beats a tuned lexical-overlap baseline. Benchmark, code and harness released at github.com/suprabhatrapolu/factverify-benchmark; every reported figure reproduces from it.

A vintage-resolved audit of claims about official statistics 2026
Dirmacs, Hyderabad. Completed. Research note and extended abstract under review. Outputs: [3], [4].

- Question: are published claims about official statistics grounded in the record they cite, and does checking them against today's record misjudge them?
- Methods: 4,669 claims from 334 news articles on 19 US statistical releases, extracted, typed and linked to FRED series, then resolved against the ALFRED vintage current on the day each was published; preregistered hypotheses; comparison across outlet tiers; an archival analysis of payroll revisions; validation by blind adjudication, including two held-out releases, and a cross-model replication.
- Finding: outright contradictions are rare (about 2.6% after correction), but of claims supported by the record at publication, 17.8% are contradicted by the same record today, rising to 43.3% for monthly payroll; outlet tier did not predict groundedness. The first published monthly payroll change later moved beyond its published precision in 97% of 33 months (largest move 194,000 jobs), showing that treating the latest record as static ground truth can misclassify claims that were correct when made.

Tool design for time-indexed claim checking with language models 2026
Dirmacs, Hyderabad. Completed; accepted at the OE-Agent Workshop at ACML 2026. Output: [2].

- Question: when a language-model agent checks claims about revised statistics, does the time index of its data tool decide whether accurate reporting is judged correctly, and can a small model with the right tool match a much larger one without it?
- Methods: a preregistered census evaluation of 2,132 claims from the audit; three conditions (no tool, a tool returning the latest vintage, a tool returning the vintage in force on the claim's date); four models from 8B to 1.6T parameters, including a closed frontier model; exact McNemar tests with Holm correction; document-level bootstrap intervals.

- Finding: with the as-of tool, Llama-3.1-8B was correct on 64.5% of claims that were accurate when published but conflict with the latest data, against 51.4% for a 1.6T model without tools (Holm p = 0.0009). A latest-value tool made models reject accurate reporting more often, though for Llama this depended on the serving host; remaining errors concentrate in revision claims, which need two vintages.

An open archive of India's official statistics as they stood on each date 2026 – present

Dirmacs, Hyderabad. In progress; first version public at github.com/suprabhatrapolu/india-record-versions. Output: [8].

- Question: after India's 2026 rebasing of GDP, consumer prices and industrial production, which version of the record does a claim rest on, and does it still hold?
- Methods: dated versions of headline statistics, each read from a saved government release with a SHA-256 manifest; claims linked to the version in force then and now. Next: a small-open-model checker, tested on hand-checked Indian claims.
- Status: 2023-24 GDP growth has had four official values (8.2%, 9.2%, 7.2%, 7.3%); every claim recorded so far matched the record on its date but not today.

Study of governance in a live multi-agent deployment (under double-blind review) 2026 – present

Output: [5].

Bilingual hate speech detection: consistency across English and Telugu 2026 – present

School of Computer and Information Sciences, University of Hyderabad. With Harshith Patnala (M.Tech, Artificial Intelligence and Machine Learning); supervised by Prof. Atul Negi. In progress; paper in preparation for RANLP 2027. Output: [9].

- Question: do multilingual hate-speech models give the same verdict to the same meaning written in English, Telugu script, Romanized Telugu and code-mixed text, in label, highlighted evidence and representations?
- Methods: about 300 X-MuTeST seed items in all four forms plus about 20 HateCheck-style tests adapted to Telugu (about 3,200 items); meaning-preserving edits should keep the label, meaning-changing ones should flip it. My part: functional tests, invariance and directional evaluation, statistics.

Time-indexed verification of restated company figures 2026 – present

Dirmacs, Hyderabad. Data, design and analysis plan complete. Output: [6].

- Question: do language-model verifiers judge a claim against the record as it stood when the claim was made, or against the latest figure?
- Methods: 300 figures that companies restated in later 10-K filings, from 86 companies, each version confirmed against the filing's XBRL document; each presented in 12 controlled variants crossing the claimed figure with the claim's date; nine frontier models, a reasoning model and a prompted-rationale arm; memory tested with counterfactual evidence.
- Design: five hypotheses, to be preregistered before the runs; gold labels balanced within each item, so a fixed answer shows no effect; a wild cluster bootstrap whose false-positive rate was checked by simulation before any model run.

Decision models as write-time gates for AI knowledge stores 2026 – present

Dirmacs, Hyderabad. In progress, with Baalateja Kataru (co-author); benchmark build from November 2026. Output: [7].

- Whether fast, typed decision models can make the judgments a knowledge store needs at write time (does a source support a fact; does new evidence contradict a stored one; has a revision changed it), compared with NLI models, language-model judges and rule-based checks on accuracy and robustness to planted instructions, under a preregistered protocol.

PUBLICATIONS

Peer-reviewed

[1] Suprabhat Rapolu. 2026. When FEVER Meets Finance: A Controlled Benchmark Reveals Overlap-Driven Financial Claim Verification. In Proceedings of the 11th Workshop on Financial Technology and Natural Language Processing, pages 518–523, Budapest, Hungary. Association for Computational Linguistics. Poster presentation (online), October 28, 2026. Code: github.com/suprabhatrapolu/factverify-benchmark. ACL Anthology: aclanthology.org/2026.finnlp-1.42.

[2] Suprabhat Rapolu. 2026. Today's Number, Yesterday's Claim: How Tool Design Shapes LLM Checks of Revised Statistics. Accepted, OE-Agent Workshop at ACML 2026 (non-archival), Melbourne, Australia, December 1, 2026.

Under review

[3] Suprabhat Rapolu. Misinformation About Economic Numbers Is Rare; Mis-Verification Is Common. Research note, under review, Harvard Kennedy School (HKS) Misinformation Review.

[4] Suprabhat Rapolu. Checked Against Which Vintage? Data Revisions and the Automated Verification of Claims about Official Statistics. Extended abstract, under review, New Techniques and Technologies for Statistics (NTTS 2027), Eurostat; decision expected November 30, 2026.

[5] Study of governance in a live multi-agent deployment (under double-blind review).

In preparation

[6] Suprabhat Rapolu. Checked Against Which Record? Time-Indexed Verification of Restated Financial Figures (working title). Target venue: ACL Rolling Review, 2027.

[7] Suprabhat Rapolu and Baalateja Kataru. Decision Models as Write-Time Gates for AI Knowledge Stores (working title). Target venue: ACL Rolling Review, 2027.

[8] Suprabhat Rapolu. Checking Claims about India's Revised Statistics (working title). Target venue: RANLP 2027.

[9] Harshith Patnala, Suprabhat Rapolu and Atul Negi. English–Telugu Hate Speech Detection (working title). Target venue: RANLP 2027.

PATENT, SERVICE AND AWARDS

- Co-inventor, Governed epistemic store with per-claim state assignment and permanent evidence. Provisional application filed March 2026, India; complete specification in preparation.
- Reviewer, ACM FAccT 2027 (invited).
- Reviewer, FinNLP Workshop at EMNLP 2026 (two submissions).
- EMNLP 2026 Diversity and Inclusion Subsidy, awarded for presentation of the FinNLP paper.
- Google for Startups: Hub, Hyderabad (T-Hub), member for Dirmacs since August 2026; Immersion (Antler India).

PROFESSIONAL EXPERIENCE

Founder and Chief Executive Officer, Dirmacs 2024 – present

Dirmacs Private Limited, Hyderabad, India (dirmacs.com). Incorporated March 2026; DPIIT-recognized; three co-founders, team of four; bootstrapped.

- Lead the company's systems (below) and its research, with a standing rule that no performance figure is published until it reproduces on a pinned test harness.
- Deliver the hosted intelligence backend for a mental-health platform, returning provenance-stamped answers while holding no user data.
- Started Dirmacs in 2024 as a strategy and software consultancy for more than fifteen businesses; moved the firm from services to infrastructure in 2025.

Private Tutor, Mathematics, Economics and Computer Science 2020 – present

IB Diploma students, Hyderabad, India.

- Teach programming, quantitative methods and economics to IB Diploma students, and coach them through the research component of their coursework; a smaller group since 2024.

Equity Research Analyst Intern, 3C Capitals Jul 2023 – Aug 2023

Manipal, India.

- Research reports on six listed companies: profitability, cash flow and valuation multiples, from the filings.

Strategy Consulting Intern, Cushman & Wakefield Mar 2023 – May 2023

Hyderabad, India.

- Built the market databases used as the evidence base for feasibility reports.

Customer Relationship Associate, Sutherland Global Services: client queries for international clients Dec 2022 – Feb 2023

Business Development Intern, Think Evolve Consultancy: TiE Global Summit pitch; 6 sales, 30+ leads Dec 2022

Customer Relations and Product Analyst Intern, Decathlon: customer surveys to shape merchandising Apr 2018 – Jul 2018

SYSTEMS

Built at Dirmacs; open-source components at github.com/dirmacs.

- Eruka (dirmacs.com/eruka): designed and lead a governed knowledge store. Every fact enters through one write gate carrying one of four epistemic states and a timestamp; evidence is append-only, and conflicting writes are logged for review. Open components: openeruka, eruka-mcp. Subject of the patent above.
- Eekvul: designed an analysis engine for public data in which every finding carries its sources and their versions, its method, a confidence label with written reasons and its gaps, computed from daily sealed snapshots so any figure can be recomputed. In development; first brief November 2026. Next: a library of analysis methods across domains, with an AI layer that proposes the data and analysis for a research question, open to anyone.
- ARES (github.com/dirmacs/ares): architecture contributor to an open-source, multi-provider agent harness in Rust, with tool calling, retrieval and MCP. Principal architect: Baalateja Kataru.
- Haimdall (ai.dirmacs.com): co-built a model gateway offering one API across model providers with automatic failover, priced in rupees.

TECHNICAL

- Research methods: benchmark and contrast-set design; annotation-artifact diagnostics; LLM evaluation harnesses; vintage-resolved record matching; preregistration; bootstrap and Wilson intervals; wild cluster bootstrap with simulation-based calibration; exact tests; regression.
- Tools and data: Python, R, SQL, Java, JavaScript; Tableau, Power BI; PyTorch, Hugging Face Transformers, scikit-learn, XGBoost; AWS Bedrock; SEC EDGAR and XBRL, FRED and ALFRED archival vintages, FEVER, Fin-Fact.
- Engineering: Rust and Python; PostgreSQL, Redis, Docker, Azure; MCP and REST; built with AI coding agents I direct.
- Coursework (NPTEL, IIT, proctored): Calculus of One and of Several Real Variables, IIT Kanpur, 2025 (both Elite); Applied Multivariate Statistical Modeling, IIT Kharagpur, and Applied Econometrics, IIT Madras, 2026.