



B. Tech. (2025-26)

● Reflections of — Artificial Intelligence & Data Science Excellence

Prepared by :

Department of AI and DS, DJSCE

PREFACE

It gives us immense pleasure to present **Project Prism: Reflections of Artificial Intelligence & Data Science Excellence**, a thoughtfully curated abstract booklet showcasing the innovative project work and research contributions of our final-year B.Tech. Students from the Department of Artificial Intelligence (AI) and Data Science.

Project Prism reflects the multifaceted brilliance of our students – like light dispersed through a prism, each project shines uniquely while emerging from a common foundation of knowledge, curiosity, and rigorous training. These abstracts illustrate the students' ability to apply advanced AI and data science techniques to address complex, real-world challenges with creativity, responsibility, and technical precision.

Spanning a wide range of themes including intelligent healthcare, predictive analytics, smart systems, cognitive computing, and ethical AI, these projects mirror the spirit of innovation and excellence that our program aims to instill. The diversity and depth of topics underscore the transformative potential of AI & DS in solving today's pressing problems and shaping tomorrow's technological landscape.

This booklet not only serves as a record of academic accomplishments but also as a source of inspiration for students, educators, researchers, and industry practitioners alike. We commend the dedication of our students and deeply appreciate the mentorship and support provided by our faculty members in guiding them through their research journey.

We hope that Project Prism will continue to be a beacon of the academic spirit, collaborative inquiry, and excellence that defines our department.

Department of Artificial Intelligence (AI) and Data Science

Dwarkadas J. Sanghvi College of Engineering

June, 2026.



Shri Vile Parle Kelavani Mandal's

DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING

(Autonomous College Affiliated to the University of Mumbai)

NAAC Accredited with "A" Grade (CGPA : 3.18)



GRAPHMIND: AI-Driven Semantic Knowledge Graph From Unstructured Data

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Ashna Aglawe	60018220109
Moksh Vora	60018220045
Parth Gala	60018220022
Vivek Vasani	60018220134

**Under the guidance of
Prof. Deepali Patil**

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

This work introduces the implementation of a complete SKG(Semantic Knowledge Graph) and GraphRAG pipeline, which automatically transforms heterogeneous document collections into a provenance-aware, queryable knowledge network while offering adaptive schema evolution and hybrid retrieval. The system ingests PDFs and plain text-with extensible interfaces to OCR/ASR-tools(Optical Character Recognition/Automatic Speech Recognition), implements multi-stage Natural Language Processing (sentence segmentation, Name Entity Recognition, entity linking, and relation extraction), and outputs reified RDF(Resource Description Framework) triples that capture relation attributes concerning type, confidence, and extraction method, along with fine-grained provenance links towards sentences and source documents. In order not to suffer from schema stagnation, the pipeline generates conservative ontology extensions from observed instances and safely merges them into a core OWL(Web Ontology Language) ontology while strictly maintaining Turtle serialization and version tracking. The knowledge base consists of two modalities: one is the standards-compliant RDF triple store supporting queries via SPARQL, and the other optional modality uses a Neo4j property graph for visualization purposes and graph traversal. Chunk embedding indexes are generated to support Sentence Transformers and FAISS(Facebook AI Similarity Search) speedups, which makes way for hybrid GraphRAG queries that can combine the best of both worlds regarding semantic graph patterns and high recall vector searches in one efficient model, aided by a fallback vectorizer. Provenance information and confidences are combined in a reranking framework that guarantees factual grounding and decreases hallucinations. SHACL(Shapes Constrained Language) validation and unit tests are performed to support the structural correctness of the overall framework along the entire processing pipeline, while performance analyses target the gains of parallelized knowledge extraction, conservative ontology updates, and hybrid search strategies. Finally, this work provides a sound scalable and standards-conformant framework in that the authors combine dynamic knowledge ontology evolution, reified provenance, and hybrid knowledge searches making a great starting point to support future work on LLM-powered relation extraction, enhanced temporal reasoning, cross-document coreferencing, and multimodal knowledge integration.

Keywords: Semantic Knowledge Graph, GraphRAG, Dynamic Ontology, RDF/OWL, SPARQL, Provenance Modeling, Reification, Entity Extraction, Relation Extraction, Apache Jena Fuseki, Neo4j, FAISS, Sentence Transformers, Hybrid Retrieval, SHACL Validation, Information Extraction, Knowledge Representation, Semantic Search



Shri Vile Parle Kelavani Mandal's
DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING
(Autonomous College Affiliated to the University of Mumbai)
NAAC Accredited with "A" Grade (CGPA : 3.18)



DataSentinel: A Whitebox LLMOPS Framework for Secure PII Redaction in RAG Pipelines

Submitted in partial fulfilment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Aayush Nisar	60018220021
Darshan Dihora	60018220093
Hemangini Patel	60018220018
Tejas Maroo	60018220052

Under the guidance of

Dr. Pratik Kanani

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT



Abstract

Large companies continue to quickly adopt large language models (LLM) into their systems, causing Retrieval-Augmented Generation (RAG) systems to gain popularity for the purpose of internal knowledge access. Nonetheless, there is a problem with this trend: sensitive data may be lost during the generation and retrieval processes. Current solutions for post-hoc redaction of sensitive data, and cryptographic embeddings, are still reactive and insufficient. However, when deploying artificial intelligence (AI) systems) in the real world, the security must span the entire operational lifespan. Moreover, many commercial artificial intelligence systems are designed as opaque "black-box" systems. Opaque systems prevent auditing, customization, and security; therefore, they are unsuitable for critical systems in highly regulated environments.

Data Sentinel is an enterprise-grade compliant RAG system LLMops white box framework with three features: (1) a proactive and data-centric Continuous Integration/Continuous Deployment (CI/CD) pipeline that verifies, updates, and scans knowledge bases for ingestion vulnerabilities; (2) a real-time protection layer that uses contextual, regex, and Named Entity Recognition (NER) models for the identification and redaction of PII (Personally Identifiable Information) within the generated response; and (3) a Continuous Integration (CI) automated red teaming framework for the assessment of security regressions, attack persistence, and data integrity all within the system's boundary. An observability stack using Prometheus and Grafana is tailored for monitoring latency, token consumption, security events, and overall system health and integrates with the aforementioned components for real-time monitoring.

The Data Sentinel reports a defense-in-depth architecture which goes after the full range of attacks in RAG systems which includes data poisoning, retrieval phase disclosure, embedding inversion, and indirect prompt injection via the use of intrinsic (data level) and extrinsic (runtime level) security elements. We see from empirical study that this multi-layer approach does in fact see large reduction in PII leakage, increased transparency, and also which provides proven output throughout the LLMops cycle. The Data Sentinel reports we have advanced the state of secure enterprise AI by putting forth a workable and repeatable approach for deployment of large-scale RAG based applications. It is put forth as an open, modular, and extensible framework instead of a commercial black box product.

Keywords: RAG, LLMs, Data Leakage, Personally Identifiable Information (PII) Redaction, White-Box, LLMops, CI/CD, Automated Red-Teaming, Vector Database Security, Prompt Injection, Data Poisoning, Enterprise AI Security.



Shri Vile Parle Kelavani Mandal's

DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING

(Autonomous College Affiliated to the University of Mumbai)

NAAC Accredited with "A" Grade (CGPA : 3.18)



SocialContentForge: Multi-Agent AI System for Automated Influencer and Brand Content Creation

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Miral Gopani	60018220097
Naman Rathod	60018220112
Rohit Parmar	60018220020
Sakshi Patel	60018220123

Under the guidance of

Prof. Sudhir Dhekane

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

The SocialContentForge is an advanced Multi-Agent AI framework that enables influencers and brands to simplify and improve their processes for generating and distributing content on Social Media Platforms. This framework addresses the most important pain points of producing and maintaining high-quality, consistent content across multiple platforms while also drastically reducing the amount of time and effort normally required to create and distribute content. It leverages cutting-edge Generative AI technology, including Google Gemini for Natural Language Generation and a Stable Diffusion v1-5 model installed locally for Image Synthesis, as well as OAuth-based connectivity to each of the supported platforms, allowing the creation and distribution of content automatically and in a way optimized for each of the supported platforms (currently Instagram, Facebook, LinkedIn, and Twitter / X). The architecture of the framework leverages multiple Coordinated Agents, with distinct Agents managing the generation of content, the distribution of content to a specific platform, and the tracking of content performance, leading to a reduction in the time to produce content of 70-75%. The framework also provides user assistance features, such as an intelligent Global AI Chatbot for navigation and content recommendations, custom instruction sets for maintaining the consistency of Brand Voice, Platform-Specific Tone controls, and Image Editing Modules for post-generation refinements. The Content Moderation System promotes the responsible use of AI-Generated content through the implementation of keyword-based filters for inappropriate prompts. Finally, Real-Time Analytics captures various engagement metrics (likes, comments, shares, reach, and impressions) across the connected platforms, allowing users to continually evaluate their content performance through an interactive dashboard. For example, a digital marketer can monitor a live campaign and instantly identify which posts are gaining higher engagement, enabling quick adjustments in content strategy, posting time, or targeting to maximize reach.

The system was implemented using MERN stack (MongoDB, Express.js, React.js, and Node.js). To enhance performance, a hybrid technique, combining GPU-accelerated computation, API-based processing, and edge computing was used to execute stable diffusion algorithms efficiently. We used both automatic and manual approaches to generate content using generative AI to support our marketing strategies on multiple social media channels. The results indicate that it is viable to automate the orchestration of contents across multiple platforms while maintaining quality, multi-user control and adherence to ethical practices of generating content.

Keywords: Multi-Agent AI Systems, Generative AI, Social Media Automation, Content Generation, Performance Analytics, OAuth Integration, Influencer Marketing, Stable Diffusion, Local Image Generation.



Shri Vile Parle Kelavani Mandal's
DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING
(Autonomous College Affiliated to the University of Mumbai)
NAAC Accredited with "A" Grade (CGPA : 3.18)



HyDRA-KV: Hybrid Deterministic Key-Value Cache Compression for Long-Context Reasoning

Submitted in partial fulfilment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Aryan More	60018220126
Himanshu Gohil	60018220014
Jenish Boricha	60018220051
Pranav Pawar	60018220032

**Under the guidance of
Dr. Pratik Kanani**

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

Long context language models face a "memory wall" during inference, with the key-value cache growing exponentially with sequence length. Existing key-value (KV) compression techniques have explored the dimensions of semantic grouping, removing redundancy, channel sparsity, query-agnostic reuse, and adaptivity. No previous method has integrated all of these into a single, numerically stable framework. We present HyDRA-KV, a new key-value compression technique. HyDRA-KV is a hybrid, deterministic, redundant, and adaptivity aware key-value compression technique. The technique is an end-to-end method for jointly optimizing sequence length, key-value cache structure, and numeric precision for long context reasoning.

This paper combines five key-value compression primitives: sentence-level semantic grouping, redundancy-aware key value eviction, channel-wise pruning, adaptive hyper tokenization, and mixed precision execution with deterministic operator ordering. HyDRA-KV is tested on the Qwen3.5-9B-Base model on symbolic equation discovery with LLM SRBench, long document question answering, and key-value reuse.

The results show a four-fold reduction in memory and latency with the new technique. The results are within $\pm 2\%$ of the baseline for symbolic accuracy. The results demonstrate the effectiveness of the new technique. The technique is reproducible with deterministic execution. The results are available for symbolic regression and long context benchmarks with the supplied evaluation harness.

Keywords: KV Cache Compression, Long-Context Large Language Models, Efficient LLM Inference, Adaptive Hyper Tokenization, Semantic-Aware Token Selection, Deterministic Decoding, Multi-Query KV Reuse, Transformer Memory Optimization.

Sentinel: Advanced Augmented layer

**Submitted in partial fulfillment of the requirements of the degree
of**

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Rishi Tiwari	60018220130
Shiv Prabhu	60018220116
Owais Ansari	60018230104
Rajvardhan Singh Chib	60018230106

Under the guidance of,

Prof. Dashrath Kale.

(Dwarkadas J. Sanghvi College of Engineering)



UNIVERSITY OF MUMBAI

Abstract

Conventional Enterprise Resource Planning systems are usually afflicted with numerous data silos and the inability to derive actionable insights out of the plethora of unstructured and structured data stored in them. This paper introduces a novel and enhanced building model that will transform ERP systems into intelligent and proactive decision-making centers. The system proposed is a complex of an advanced MoE AI router and a microservices-based backend, which allows interacting with complex enterprise data using natural language.

An AI query router is used to route users query to the right expert. For example, a query requiring access to SQL database is routed to SQL Expert. Similarly, query requiring access to data stored in documents is routed to RAG expert. That becomes the assurance of high accuracy: it is models that are explicitly trained to operate in their individual domains, and not the Large Language Model in general. It is founded on an extremely powerful tech stack- Node.js, NestJs, and Python microservices that are Docker-contained and run on Kubernetes to achieve enterprise-level scalability.

One of the biggest issues of AI-based ERPs is how to maintain data privacy and tenant isolation. This architecture meets such requirements in its two tier security paradigm. First, it has strict Tenant Data Isolation at the database level through different schemas and different vector namespaces. Second, a Federated Learning loop is evolved that enables the system to improve its global intelligence without transferring sensitive data of tenants. Fine-tuning occurs in a particular organizational situation with high level of compliance through Secure Aggregation Servers and per-tenant adapters.

Moreover, Explainable AI Layer inclusions imply that the results of the system are not black boxes. It offers reason, data provenance, and snippets of transparency hence gives an outline that brings about trust among the executive stakeholders. This is a smart ERP platform that is the next generation jump in enterprise technologies that will result in a scalable, secure and very interpretable solution that bridges the gap between raw data and strategic business implementation.

Keywords: Mixture Of Experts (MOE), RAG, machine learning, LLM, NLP, Federated learning, Explainable AI.



Shri Vile Parle Kelavani Mandal's
DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING
(Autonomous College Affiliated to the University of Mumbai)
NAAC Accredited with "A" Grade (CGPA : 3.18)



AI Research Lab in a Box: A Multi-Agent Collaboration System for Autonomous Scientific Discovery

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Anushri Venkitaramanan	60018220075
Arya Vaidya	60018220131

**Under the guidance of
Prof. Monali Deshmukh
Dwarkadas J. Sanghvi College of Engineering**



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

Meet AI Research Lab in a Box- another tool created on the open-source code and unites a group of agents to collaborate. It makes natural language processing, and machine learning the center of focus, with other forms of agents dealing with each aspect individually. The organization maintains order and components can be changed in and out easily. The tasks remain focused and would not overlap and the design can allow pieces to run independently should the need arise. The principle is to maintain flexibility and manage the complexity.

One of the parts of the system searches about scientific articles according to your question. It employs broad search terms, and meaning matching to locate relevant studies that are of interest. A second component goes down to those papers, extracting how experiments were carried out, what data was utilized, which measures were important and what results appeared. This work is based on the models of language models that are trained to comprehend a dense, complex text. The third agent transforms all the accumulated to clean write-ups or well organized reports and projects the facts plainly. The roles occur in a series in order to transform questions into organized knowledge.

At the same time, the ML system consists of five different agents, each of which is dedicated to one step of the working process, such as Researcher, Planner, Experimenter, Critic, and Writer. Instead of everything moving simultaneously, they do so sequentially as each one takes a step forward when it is his or her turn. The Researcher gets the background information and digs on previous findings in the data sets at hand. The next is the Planner, which is determining the way tests are run by deciding on the algorithms to use, the data to choose, and parameters to adjust. However, when prepared, the Experimenter introduces training activities, follows changes in performance, and gives the comparisons between different models. Thereafter the Critic will examine performance and point out areas that would benefit by being tweaked. Last but not least, the Writer transforms the results into readable and understandable summaries. With fewer manual labor needs, the researches flow through peer feedback, experiment design, simulations and write-ups with a simpler groove- more rapidly- pacing processes and an unconscious pen pushing the envelope. It is constantly changing, but development remains constant.

One of the functions that machines aid scientists is the execution of teams of smart programs. These groups deal with such tasks as the identification of research papers or data results studying. We may have one of those systems that have readers, test runners, and results checkers. One section conducts summaries and another strategizes the further actions of a project. Job-organizing ideas, spotting errors, and so forth are associated with each position. They collectively constitute an automation process that facilitates the process of discovery, not necessitating a human guiding process. Learning model-driven tools unite these functions into a single operation.

Keywords: Multi-Agent Systems, NLP Automation, Literature Retrieval, Experiment Analysis, Report Generation, Machine Learning Workflow, Researcher Agent, Planner Agent, Experimenter Agent, Critic Agent, Writer Agent, Autonomous Research Framework, Scientific Discovery Automation



Shri Vile Parle Kelavani Mandal's

DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING

(Autonomous College Affiliated to the University of Mumbai)

NAAC Accredited with "A" Grade (CGPA : 3.18)



Automated Compliance Check For Advertisement Approvals Using Agentic AI As Per NCDEX & SEBI Guidelines

**Submitted in partial fulfillment of the requirements of the degree
of
Bachelor of Technology**

In

Artificial Intelligence (AI) and Data Science

By,

Harsh Chitaliya	60018220029
Sameer Gupta	60018220007
Vinti Bhatia	60018220046
Yash Buddhadev	60018220122

**Under the guidance of
Prof. Nikita Raichada
Dwarkadas J. Sanghvi College of Engineering**



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT



Abstract

Advertisement compliance is a very important regulatory workflow in the commodity and securities markets in India, especially for trading members and brokerage firms regulated by NCDEX and SEBI, respectively. Approval of process for promotional content includes, but is not limited to, YouTube videos, audio advertisements, influencer promotional clips, social media posts, banners, brochures, and print PDFs, each is to be painstakingly checked by hand before publishing. Continuous re-evaluation of disclosures, disclaimers, claims, and financial messaging not only adds to the workload that compliance teams have to put in but also leads to delays, subjective decisions, and potential oversights resulting in investor misinformation and regulatory fines.

The project on hand proposes a fully automated Agentic AI-based Advertisement Compliance Platform that verifies promotional content across all media formats by aligning it with NCDEX circulars, SEBI advertisement guidelines, investor protection norms, and industry-standard disclosure requirements.

This is done through the use of multilayer automation pipelines, involving the following: the intake of banners, PDFs, audio-visual advertisements, or submitted YouTube URLs, followed by multimodal extraction using Whisper V3 for speech-to-text transcription, Pixtral-OCR for visual text recognition, and multilingual translation APIs for uniform English-language evaluation. A structured rule-mapping layer, implemented through LangChain Agents and LLM-based reasoning, evaluates the extracted information on regulatory compliance factors pertaining to mandatory risk disclosures, warnings regarding market volatility, product risk categorization, claims of return on investment, role of

"misrepresentation of guarantees, improper use of financial terminology, and incorrect delivery in financial instruments."

audio or visual form of disclaimers.

It does this in the following way by using the dynamic rule engine, which can recreate the SEBI/NCDEX circulars into

modular compliance checkpoints to supporting continuous regulatory updates, selective filtering of rules, and applicability in a specific context for a specific format, such as the rules that apply just to audio-visual media.

Detection of any violations is done via contextual descriptions after which the system automatically prepares an intelligent compliance report based on color codes (Red - Non-compliant; Yellow - Partly compliant; Green - fully compliant), along with exact timestamps and frame numbers, along with instructions for rectifying the same for re-submission. A real-time dashboard helps users evaluate multiple advertisements simultaneously and supports features like batch submissions, vendor/trading member level content evaluation, report archival and submission history for maintaining audit logs. Also, the system includes YouTube download automation, multilingual speech processing and a version-based circular comparison to synchronize past guidelines with current ones. Notifications and reminders can be sent through emails, SMS or WhatsApp messages in order to hasten approval process and maintenance of compliance records. Its scalable and modular structure makes this solution easily integrable with NCDEX and SEBI portals of its members, intermediary dashboards, CRM systems, Brokerage Back Office Systems and even cloud document management solutions. In



Shri Vile Parle Kelavani Mandal's
DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING
(Autonomous College Affiliated to the University of Mumbai)
NAAC Accredited with "A" Grade (CGPA : 3.18)



X-Tutor: Explainable Multi-Modal AI Teaching Assistant

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Aaron Mendonca	60018220028
Arya Gawde	60018220104
Nikki Mehta	60018220019
Nilay Koul	60018220023

**Under the guidance of
Prof. Shruti Dodani**

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

Modern education increasingly relies on AI-powered tutoring systems, yet most existing platforms remain limited in their ability to address diverse learner needs. These systems often generate generic, one-size-fits-all responses that fail to account for variations in learning style, prior knowledge, and individual student progress. Their reliance on text-only interaction further restricts effectiveness, as they are unable to interpret multimodal inputs such as handwritten solutions, diagrams, classroom notes, or nonverbal cues. The lack of transparent reasoning behind feedback creates a “black-box” effect that reduces student trust and prevents educators from understanding how conclusions are formed, while minimal analytical insight makes it difficult to track performance or identify learning gaps.

This project presents X-Tutor, a multimodal AI tutoring platform designed to overcome these limitations and deliver a scalable, adaptive, and explainable learning experience. The system processes both textual and visual inputs, enabling students to upload handwritten solutions or problem images for real-time interpretation. Using vision-language models, retrieval-augmented generation, and learner-profiling algorithms, X-Tutor adapts explanations, difficulty levels, and content flow to match each learner’s pace, strengths, and misconceptions. For educators, the platform automatically generates structured lesson plans, quizzes, worksheets, and teaching materials from classroom notes, while natural-language querying supports domain-wide, curriculum-aligned responses.

To ensure transparency, X-Tutor incorporates explainable AI techniques, offering step-by-step reasoning, source attribution, and interpretable justification behind each answer. A real-time analytics dashboard tracks concept mastery, error patterns, time spent per topic, and engagement metrics, providing teachers with actionable insights.

By unifying multimodal understanding, adaptive learning pathways, and explainable decision-making, X-Tutor enhances personalization, strengthens student confidence, and gives educators deeper visibility into how students learn. It establishes a foundation for future advancements in AI-driven pedagogy and has the potential to transform classroom teaching, remote learning, and individualized tutoring at scale.

Keywords: Multimodal AI, Intelligent Tutoring System, Adaptive Learning, Vision-Language Models, Explainable AI, Lesson Plan Generation, Educational Analytics, Personalized Learning.



PathFinder AI: An Agentic AI Career Guide and Technical Interview Assistant

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Akshad Jain	60018220107
Ami Khatri	60018220041
Bhavy Shah	60018220016
Shreya Singh	60018220128

**Under the guidance of
Prof. Nikita Raichada
Dwarkadas J. Sanghvi College of Engineering**



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

ABSTRACT

A highly fragmented and disjointed ecosystem of career resources, in which the candidates tend to experience a deficiency in personalized and tailored guidance, is currently defining the modern job marketplace. The classic job portals and inflexible resume parsing systems do not recognize the critical dimension of time, and leave applicants to navigate through unstructured study material, without a clear and scientifically-supported plan of future interviewing. All these systemic inefficiencies are solved with PathFinder AI: An AI-powered Career Guidance assistant, which brings an advanced Agentic AI system, a personalized virtual mentor. The solution's architecture is three-tier, decoupled and high-tech, leveraging a Node.js and Express.js backend to communicate with complex reasoning of Large Language Models (LLM) through the OpenAI API. The main innovation of the system is the Deadline-Aware Scheduling Algorithm developed by the system, which can be executed automatically on user PDF resume and set interview deadline, and will make a deep semantic skill gap evaluation. The system can automatically assign learning modules and curated materials dynamically in the rest of the time constraint for study to optimize study efficiency. This is to make sure that with 30 days or 3 days, the preparation strategy will be mathematically optimized (i.e., with 30 days, it will first focus on the high-yield topics before focusing on the low-yield topics).

In addition to its scheduling features, PathFinder AI includes a context-aware conversational chatbot that enables users to have real-time coaching and clarify complex technical doubts in the context of their particular roadmap. The ideas of algorithmic fairness and ethics are tightly coupled with the philosophy of system design. The architecture uses rigorous prompt engineering and bias-reduction measures, such as PII redaction, to make sure that the AI considers applicants based on technical merit alone, deliberately ignoring their demographic characteristics, such as gender, race, or age. The system has been proven by quantitative tests to be efficient with a high F1-score of 0.92 in detecting skill gaps and a median roadmap generation latency of only 15.2 seconds with state of the art models such as Gemini 3 Flash. This speech-to-text interview simulator further improves user experience as it provides a high-fidelity environment to practice behavioral and technical responses. Finally, PathFinder AI shows the potential of the Agentic AI in democratizing access to elite-level career coaching, turning a traditionally stressful and disorganized process of interview preparation into a data-driven, scientifically planned and highly empowering experience, both to students and to professionals.

Keywords: Agentic AI, Career Guidance System, Large Language Models (LLMs), Resume Parsing, Skill Gap Analysis, Generative AI, Natural Language Processing (NLP), Time-Bound Scheduling, Algorithmic Bias Mitigation, Prompt Engineering, Personalized Learning.



Shri Vile Parle Kelavani Mandal's

DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING

(Autonomous College Affiliated to the University of Mumbai)

NAAC Accredited with "A" Grade (CGPA : 3.18)



SaralGST: Multi-Agent XAI Framework for Automated GST Filing and Reconciliation

**Submitted in partial fulfillment of the requirements of the degree of
Bachelor of Technology**

In

Artificial Intelligence (AI) and Data Science

By,

Deep Jain

60018220118

Satyavrat Tiwari

60018220063

Morvi Panchal

60018230105

Gaurav Singh Khati

60018220057

Under the guidance of

Prof. Dashrath Kale

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT



Abstract

The complexity of GST compliance for businesses operating within India creates significant challenges due to repetitive data handling, provides many opportunities for errors during data reconciliation, and requires timely filing of relevant forms. To help solve this problem, a project has been developed that enables businesses to manage GST compliance using a completely automated, AI-based GST Management and Compliance Platform that provides automation at every stage of the GST lifecycle and reduces human involvement wherever possible. Combining intelligent data processing with predictive analytics, fuzzy matching algorithms, and n8n-based automated backend orchestration, the system is designed to provide businesses with the ability to achieve greater compliance by leveraging more accurate data with significantly less manual input. The workflow of this process involves multiple automated layers, from the ingestion of invoices, bank statements, GSTR-2B files, and accounting software exports, to real-time validation, data cleaning, data standardization, and structured transformation. A sophisticated reconciliation engine is used to compare a business's internal purchase and sales records with GST portal data using semantic and fuzzy logic to identify discrepancies and missing transactions. Other modules have been created to automatically generate GSTR-9, create e-way bills in real time, optimize ITC for capital goods, simulate capital goods ITC, enable zero-click filing, and provide aggregated real-time validated transactional data from multiple sources. AI components that provide core functionality for the GST platform include expense categorization, separation of business and personal expenses, and prediction of ITC eligibility. The voice-enabled assistant provides multilingual support for queries, allowing businesses to comply with GST requirements without needing extensive technical knowledge.

Keywords:

GST automation, invoice reconciliation, AI compliance systems, ITC optimization, semantic matching, AI Automation workflows, GSTR-9 automation, e-way bill generation, predictive tax planning, voice-based GST assistant, diagnostic automation for taxation



Shri Vile Parle Kelavani Mandal's
DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING
(Autonomous College Affiliated to the University of Mumbai)
NAAC Accredited with "A" Grade (CGPA : 3.18)



SkillTrace : Model Context Protocol (MCP) for Adaptive, Cross-Platform Learning

Submitted in partial fulfilment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Ayaan Zaveri- 60018220061

Pratham Malavia- 60018220060

Rachit Doshi - 60018220095

Swayam Shah - 60018220124

Under the guidance of

Dr. Mrunal Rane

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

However, contemporary self-directed learning becomes increasingly fragmented in terms of digital technologies usage, contributing to cognitive overload, inconsistent progress monitoring, and vague suggestions from generic AIs. SkillTrace – a proposed context-aware solution – unifies online learning activities by integrating browser extensions for YouTube and LeetCode into a unified platform through a Model Context Protocol (MCP) server. Browser extensions automatically capture rich context (timestamp, transcript fragments, and coding exercises) into a shared MongoDB database. In turn, an MCP server provides a querying mechanism to provide the collected context to Gemini AI, allowing for deterministic and retrieval-augmented generation based on the learner’s past behavior. Evaluation shows that SkillTrace reaches 100% Tool Selection Accuracy (TSA) for complex user queries and decreases the average latency of response by 78% (17.48s versus 3.79s), as opposed to monolithic context injection. Moreover, the use of structured database querying through an MCP allows to mitigate state hallucination completely and increase token context efficiency by 95%.

Keywords: Adaptive learning, Model Context Protocol, browser plug-in, LLMs, cross-platform analytics, personalized education.



Shri Vile Parle Kelavani Mandal's
DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING
(Autonomous College Affiliated to the University of Mumbai)
NAAC Accredited with "A" Grade (CGPA : 3.18)



The AI Autopsy: An XAI Debugging Assistant

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

**Esha Patel
Siddhi Shetty**

**60018210106
60018210081**

**Under the guidance of
Dr. Mrunal Rane**

Dwarkadas J. Sanghvi College of Engineering



ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

UNIVERSITY OF MUMBAI





Abstract

As machine learning models get more complex, decisions made by machine learning models are difficult to interpret, and therefore, they turn into black box models. Even though there are methods of XAI, for instance, SHAP, Grad-CAM and LIME, that give clear information regarding individual predictions, generally, they lack a way of systematically archiving, comparison, or accessing of information in relation to model explanations over time.

The problem this project tackles, 'The AI Autopsy: An XAI Debugging Assistant,' provides a compatible structure which can assist in explaining and debugging the AI by providing a queryable repository of model explanations. The project can process explainability results produced by Grad-CAM, SHAP and LIME in a number of ways, including in the form of JSON files, numerical arrays, and plots (heatmaps). These results are then analyzed to convert them into a standard format of explanation statements containing prediction context information, feature contributions, text summaries, and image support.

The retrieved text data is embedded and stored in a vector-oriented storage infrastructure in a way that allows for efficient project-level access. By natural-language query by a developer in the system, patterns of past explanations are determined and formulated into responses based on past explanations accessed due to their relevance to a query, allowing a developer to track patterns of repetitions, analyze explanations based on experiments, and determine areas of possible errors in model reasoning more efficiently.

AI Autopsy is better at providing explanations since it organizes explainability results into a systematic framework. It enhances the development of AI since the results reduce uncertainty associated with debugging and provide a valuable resource for developers when they are faced with complex machine learning systems.

Keywords: Explainable AI (XAI), Model Debugging, Retrieval-Augmented Generation (RAG), Vector Indexing, LIME, SHAP, Grad-CAM, System Design.



LapTempo: AI Driven Competitive Swimming Pacing, Prediction, and Performance Diagnostics.

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Hrishikesh Karmalkar

60018210096

Under the guidance of

Dr. Mrunal Rane & Prof. Purva Badhe

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT



ABSTRACT

LapTempo is an AI-powered pacing intelligence system to revolutionize the experience of competitive swimmers and coaches as it interprets, forecasts, and optimizes race performance through its pacing intelligence. From scientifically validated pacing understandability and outcome perspective, stroke information combined with modular machine learning approaches for pacing intelligibility and result from two core capabilities of the platform - Pre-Race Ideal Stroke Split Prediction (with stroke-specific regression model predicting individualized pacing strategy) and Post-Race Diagnostic Assessment, where Δ deviations are used to reason biodynamic inefficiency and physiological disintegration including lactate formation, neuromuscular fatigue, under-water deterioration, and technical instability. With combination of biodynamics and physiological principles of fatigue, together with explainable AI principles, LapTempo addresses the gap between timing information and strategy insight. Modular architecture of the LapTempo platform is extensible to accommodate advanced modules like HRV based training readiness, lactate modeling, stroke number and stroke rate, and IMU underwater detection. Outcomes from exhaustive testing on Worlds, Reg, and Comp data sets conclusively demonstrate that LapTempo can successfully replicate pacing patterns, identify fatigue charts, and so forth, and that diagnostic components are coach-friendly and aligned with current sports science principles. This tooling platform offers an enabling technology to leap to the next level of AI-driven swimming analytics to mark the milestone to an integrated sports analytics platform for an integrated and comprehensive data-intelligent racing strategy for athletes and coaches.



Shri Vile Parle Kelavani Mandal's
DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING
(Autonomous College Affiliated to the University of Mumbai)
NAAC Accredited with "A" Grade (CGPA : 3.18)



CTDE-Driven MAPPO Framework for Multi-Agent Race Strategy Optimization in Formula One

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Hitansh Mehta 60018220094

Kunj Modi 60018220017

Ronit Jariwala 60018220133

Veer Raje 60018220110

Under the guidance of

Prof. Deepali Patil

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

The pinnacle of motorsports, Formula One (F1), is known to be sustained in pursuit of performance in which the execution of strategy in racing becomes the deciding factor. The proposed project is the development of an intelligent framework using Artificial Intelligence to optimize strategies, focusing on tire management strategy and pit stops, with the use of Multi-Agent Reinforcement Learning (MARL). Unlike classical deterministic systems, the developed system describes the adversarial and cooperation activities involved in a racing situation. The main focus of this project involves a fully automated simulation system mimicking the stochastic situation of a Grand Prix event, more specifically the 2023 Azerbaijan Grand Prix. Real-time data regarding the performance of Pirelli tire materials and the influence of a Grand Prix event like a Safety Car system are incorporated into the system to analyze these stochastic activities. The solution incorporates a data ingestion process performed using a custom physics system and a Multi-Agent Proximal Policy Optimization (MAPPO) algorithm with a Centralized Training Decentralized Execution (CTDE) architecture. The innovations are in the Targeted Selective Broadcasting (TSB) process for threat perception, which facilitates 'smart defense' actions. The policies learned do show emergent properties like under cuts, over cuts, and Safety Car insights, thus justifying that Deep Learning can complement human intelligence on the pit wall.

Keywords: Multi-Agent Reinforcement Learning (MARL), Formula One Strategy, Pit Stop Optimization, Centralized Training Decentralized Execution (CTDE), Targeted Selective Broadcasting (TSB), Tire Management, Deep Reinforcement Learning, Real-Time Decision Making, Threat-Aware Defense, Race Simulation, FastF1 Integration, Proximal Policy Optimization (PPO), Strategic Decision-Making, Motorsport Analytics.



Beyond Code Generation: A Multi-Tool Agentic AI Framework for Autonomous End-to-End DevOps Automation

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Aditi Ambasta	60018220113
Aditya Dubey	60018220132
Akshansh Dwivedi	60018220103
Bhavvya Jain	60018220108

**Under the guidance of
Prof. Pratik Kanani**

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

The rapid evolution of the microservices and cloud native architectures and technologies have greatly increased the complexity of DevOps based workflows, which has presented a steep learning curve for developers. While at scale, many companies have a separate DevOps team that handles the entire deployment, operations and monitoring for the product or services, at medium to smaller scale, there's a need to automate the development and operations process for faster, proper, ready to scale, and flexible delivery and release of application with minimal human intervention. This project introduces an Intelligent LLM powered Multi-Agent System which is designed to automate the entire complex DevOps processes pipeline directly within an Integrated Development Environment (IDE), which has been implemented as a dynamic Visual Studio (VS) Code extension. The novel multi agentic AI framework specializes in 82 tools which are organized across 16 functional categories which are capable of autonomously executing complex software lifecycle tasks through natural language interface. The system integrates code review, unit test generation, Github repository management, Docker containerization, Kubernetes orchestration, CI/CD pipeline monitoring and multi platform cloud deployment within a unified architecture.

Keywords: Agentic AI, DevOps Automation, Large Language Models, CI/CD, Autonomous Deployment, Code Review, Software Engineering, WebSocket, Cloud Deployment.



Shri Vile Parle Kelavani Mandal's
DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING
(Autonomous College Affiliated to the University of Mumbai)
NAAC Accredited with "A" Grade (CGPA : 3.18)



Assessly – An End-To-End Online Assessment and Plagiarism Detection System

Submitted in partial fulfillment of the requirements of the degree of

**Bachelor of Technology
In
Artificial Intelligence (AI) and Data Science**
By,

Aaryan Chodankar 60018220064

Divye Shah 60018220027

Dnyanesh Sawant 60018220114

Vidit Shah 60018220058

Under the guidance of
Prof. Sudhir Dhekane
Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

The abrupt switch to online education systems makes the struggle against plagiarism even more difficult since current methods of combating plagiarism depend heavily on surface-level similarity in texts and can be simply bypassed by changing the names of variables and other coding practices. Moreover, there is no way of preventing cheating, identity fraud and access to external resources while testing in the untethered remote setting. In this work we will introduce a novel system called Assessly, which aims at providing end-to-end administration of online tests along with detecting plagiarism with a combination of artificial intelligence proctoring and code comparison from multiple angles. The proposed architecture is based on a tiered system that includes React front-end web portal for admin tasks, Tkinter desktop application for secure test distribution, Node.js & Express as an orchestration layer, Flask microservices with machine learning similarity algorithms and MongoDB storage. The code similarity engine consists of five measures, such as Cosine Similarity (CountVectorizer bigram mode), TF-IDF Cosine Similarity, Jaccard Similarity, Levenshtein Ratio, and AST similarity and their combination via custom weighting and SVM classification.

Following the performance test of the classifier implemented using SVM in the classification of 220 sets of the encoded text data, it was found that the model obtained an accuracy rate of 84.09%, with the precision of 0.81, recall of 0.70, and F1 of 0.75, which outperforms the others, including XGBoost, Random Forest, Logistic Regression, Gradient Boosting, AdaBoost, and KNN.

From the performance test conducted on the algorithm employed in this study, it can be seen that the algorithm implemented in this study was able to generate the plagiarism graph for 30 individuals in less than five seconds. The rapid shift to online learning has amplified the challenges of maintaining academic integrity in programming assessments.

Keywords: Plagiarism Detection, Code Similarity Analysis, Graph-based Clustering, AI Proctoring, Machine Learning (SVM), Abstract Syntax Tree, Academic Integrity, Online Coding Assessments.



Multi-Agent AI powered social media agency: Automating content, Engagement, Analytics

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Mahek Shah

60018230107

Parth Karia

60018230108

Dev Parekh

60018230109

**Under the guidance of
Prof. Sudhir Dhekane**

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT



Abstract

The rapid growth of digital platforms has made social media management increasingly complex and time consuming. Traditional approaches that rely on manual effort and basic tools are no longer sufficient to meet the demands of real time engagement, strategic planning, and large scale content coordination. This gap highlights the need for a more intelligent and automated solution. This project proposes the development of an Agentic Social Media Agency, a system that combines the capabilities of Large Language Models and multi agent systems to create a collaborative and autonomous digital workforce. Instead of focusing only on content scheduling, the system introduces specialized AI agents such as content creators, orchestrators, audience engagement handlers, and real time analysts. These agents work together seamlessly across the entire social media lifecycle to manage, analyze, and optimize digital presence. The framework is supported by recent advancements in multi agent architectures, task specific AI applications, orchestration systems, and reinforcement learning techniques. By continuously learning from data and adapting strategies, the system is designed to deliver faster, more personalized, and context aware interactions at scale. In addition to improving efficiency, this approach allows human marketers to shift their focus from repetitive tasks to strategic decision making and creative planning. The project also emphasizes the importance of ethical considerations such as bias reduction, transparency, and accountability to ensure responsible use of autonomous systems in digital marketing.

Overall, this work presents a forward looking solution that redefines how social media can be managed using intelligent, adaptive, and collaborative AI systems.

Keywords: Agentic AI, Social Media Management, Large Language Models (LLMs), Multi-Agent Systems, Autonomous Agents, AI-Driven Marketing, Content Automation, Audience Engagement, Real-Time Analytics, Digital Marketing, Reinforcement Learning, AI Orchestration, Intelligent Automation.



Shri Vile Parle Kelavani Mandal's

DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING

(Autonomous College Affiliated to the University of Mumbai)

NAAC Accredited with "A" Grade (CGPA : 3.18)



VisionSafe: Real-Time CCTV Suspicious Activity Detection

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Armaan Ajani 60018220053

Deep Doshi 60018220096

Udit Prabhu 60018220127

Vedant Gavankar 60018220106

Under the guidance of

Prof. Shruti Dodani

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

The intelligent video surveillance system, VisionSafe, is designed to automatically detect suspicious and anomalous activities from large volumes of CCTV footage. This will be made possible through its novel two-stage deep learning architecture. Traditional surveillance relies on manual monitoring of feeds, which is prone to attention fatigue, highly costly, and insufficiently scalable. High costs and excessive false alarms resulting from context insensitive and irrelevant anomalies detection have prevented existing automated anomaly detecting solutions from achieving widespread deployment. To address these challenges, VisionSafe relies on a hybrid deep learning framework that combines 3D convolutional neural networks with LSTM layers to learn spatial and temporal patterns in surveillance video.

The proposed system operates in a stagewise manner : Stage 1 performs binary classification, Stage 2 further classifies detected abnormal events into 13 threat categories. One novelty in our product is a differential frame-sampling strategy that efficiently captures motion-rich segments by extracting normal activities at 1.5 FPS and anomalous events at 4.5 FPS for our custom dataset, while reducing redundant frames. Optimized computational efficiency by 40% compared to uniform sampling.

VisionSafe provides real-time notifications and alerts the security personnel within seconds upon detecting an anomaly, enabling a quick security response. In this way, the security system turns passive surveillance into proactive threat management, and compared with the current solutions, it reduces false alarms while maintaining high target-detection sensitivity and being the most accurate in the market. It has been tested across varied lighting conditions, weather, and occlusion levels, VisionSafe shows strong generalization across residential, commercial, and public surveillance environments.

Keywords: Video Anomaly Detection, 3D CNN, LSTM, Real-time Surveillance, CCTV Analysis, Deep Learning, Security System



Shri Vile Parle Kelavani Mandal's

DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING

(Autonomous College Affiliated to the University of Mumbai)

NAAC Accredited with "A" Grade (CGPA : 3.18)



ATELIER: Auditable Tensor-based Engine for Large-scale Integrated Emotional Response

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Atharva Nagarsekar 60018220105

Atharva Patkar 60018220129

Darshan Jain 60018220117

Vinayak Mestry 60018220039

Under the guidance of

Prof. Deepali Patil

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

ATELIER reframes the social issue of sentiment analysis as an ex ante prediction engine. It successfully utilizes a structured social simulation system rather than simple label assignment. Large Language Models (LLMs) often exhibit hindsight bias. Unlike these models, ATELIER maps a specific natural language narrative into a compact symbolic world tensor. The engine then propagates this signal through a deterministic tensor pipeline over a synthetic population. This population features heterogeneous exposures, personalities, wealth, influence, and memory.

This approach ensures that the system bases predictions solely on the immediate event signal and the pre-existing societal state. It prevents reports released after the event from influencing the outcomes. The resulting modularity provides an auditable environment for controlled scenario analysis, ensuring seed-based deterministic reproducibility while enabling diverse scenario exploration. The system outputs a population-level distribution over emotions, engagement, and collective action. This multidimensional output demonstrates the efficacy of ATELIER as a rigorous, predictive computational sociology engine.

Keywords: Ex ante predictions, neuro symbolic reasoning, social simulation, sentiment analysis, agent-based modeling, Explainable AI.



Shri Vile Parle Kelavani Mandal's
DWARKADAS J. SANGHVI COLLEGE OF ENGINEERING
(Autonomous College Affiliated to the University of Mumbai)
NAAC Accredited with "A" Grade (CGPA : 3.18)



Integrating Language Models in QGIS for Reproducible Geospatial Workflows from Natural Language

**Submitted in partial fulfillment of the requirements of the degree of
Bachelor of Technology**

**In
Artificial Intelligence (AI) and Data Science**

By,

Parshwa Nandu 60018220119

Under the guidance of

Prof. Dashrath Kale

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

The high rate of high resolution geospatial information requires more sophisticated tools to analyze it, but new Geographic Information Systems. Steep learning curves and manual, error prone workflows are needed to run (GIS) based on QGIS. Although independent Large Language Model (LLM) agents provide automation, their non deterministic aspects tend to conflict with strict procedural accuracy of spatial analysis, resulting in problems of reproducibility and hallucinations in domains.

To fill the gap between the conceptualization of an analytical concept by humans and software specific execution, this paper presents an LLM run QGIS plugin that translates natural language queries into python based workflows structured in a way that they are reproducible. The main three methodologies used in the system to preserve this deterministic nature are, Retrieve Augmented Generation (RAG) to interactively retrieve curated knowledge base from authenticated GIS workflow, Chain of Thought (CoT) reasoning of complex multi step workflow and semantic functional. Mapping to consistently trigger the exact sequence of operations despite differences in language.

The main contribution of this project is a production ready QGIS plugin that converts natural language queries to structured geospatial operations using Python GIS libraries. The system bridges the gap between the human's intent and the machine's level of GIS operations, minimizing human error, increasing efficiency while increasing accuracy and making complex geospatial analysis available to everyone. The plugin is able to accomplish advanced tasks including (but not limited to) network analysis, data ingestion and vector processing within the QGIS user interface. By implementing a smart reasoning layer into the existing GIS framework, the solution democratizes up the opportunities of challenging procedural positions smartly and demonstrates the fact that the main value of AI in GIS is to intelligently navigate between procedural positions software intricacy as opposed to conducting spatial mathematics.

Keywords: Large Language Models, PyQGIS, Automation, Reproducibility, Spatial Analysis, Natural Language Interface, Geospatial Information Systems



Bias Radar: Real-Time Quantification of Media Bias using NLP and Machine Learning

Submitted in partial fulfillment of the requirements of the degree of

Bachelor of Technology

In

Artificial Intelligence (AI) and Data Science

By,

Aditya Tarsaria

60018220121

Under the guidance of

Prof. Sudhir Dhekane

Dwarkadas J. Sanghvi College of Engineering



UNIVERSITY OF MUMBAI

ARTIFICIAL INTELLIGENCE (AI) AND DATA SCIENCE DEPARTMENT

Abstract

For the most part every individual depends on digital news media sites and social media conversations to get information on what is happening around the world. However, there are underlying ideological frames and rhetorical devices such as sarcasm embedded in news media and online conversations. It is hard to manually identify the bias since a lot of information is being produced every day. Bias Radar is a machine learning-based tool that has been developed to detect political biases as well as the presence of sarcasm in the content of the text. The tool has been developed based on Natural Language Processing (NLP) concepts to read news articles as well as social media posts to determine the type of bias in the content of the articles. The tool has been developed based on the Distil-BERT model that helps to understand the context of the language. The system fetches live news headlines from news APIs and gathers public discussions from online platforms. Hence, the system can perform an analysis of both professional and user-generated content. The interactive dashboard created using the gradio library allows users to input their own text, analyze live news articles, and social media content in real-time. The goal bias radar is not to censor content but to encourage critical thinking and media literacy by helping users comprehend the ideological perspective in the information they consume.

The performance of the proposed system proved satisfactory, as the accuracy, precision, recall, and F1-score obtained from the test data were close to 0.93. The aim of Bias Radar is not to ban or censor any form of media content but to make individuals conscious consumers of media through a deeper understanding of the ideology hidden behind the everyday consumption of media information. Future works may focus on supporting multiple languages, explainable AI, and large-scale training data sets.

Keywords: Political Bias Detection, Natural Language Processing (NLP), DistilBERT, Transformer Models, Large Language Models (LLMs), Media Analysis, Real-Time News Analysis

Batch 2025-26

