APTINOVA

Project Reference No.: *49S_BE_0226*

College : Mangalore Institute of Technology & Engineering, Moodabidri
Branch : Department Of Computer Science & Engineering (IOT & Cyber Security with Blockchain Technology)
Guide(s) : Mr. Sandeep Seetaram Naik
Student(s) : Mr. Abdul Khadar Jilani K K
Mr. Anujith P Shetty
Mr. Makara
Mr. Poornachandra S Gowda

Keywords:

Hybrid AI, Symbolic Reasoning, Neural Language Models, LoRA Fine-Tuning, Aptitude Problem Solving, Educational Chatbot, OCR Integration, Multimodal Learning

Introduction:

The rapid expansion of competitive examinations—campus placements, banking assessments, government recruitment tests, and engineering entrance evaluations—has generated an urgent need for intelligent, domain-specific learning assistants capable of solving aptitude problems with both mathematical precision and pedagogical clarity. Aptitude questions span arithmetic reasoning, ratio analysis, profit-loss computations, time-work problems, and logical deduction; these domains demand a structured, step-by-step reasoning approach that general-purpose AI tools consistently fail to replicate. Traditional rule-based systems provide deterministic accuracy but are rigid, brittle to phrasing variations, and incapable of generating natural language explanations. Conversely, large language models (LLMs) produce fluent text but frequently hallucinate intermediate steps, exhibit imprecise arithmetic, and offer limited mathematical transparency.

AptiNova addresses this dual limitation through a novel hybrid symbolic-neural architecture that integrates deterministic symbolic solvers with a LoRA-adapted DistilGPT-2 language model. The symbolic layer handles mathematical computation through domain-specific engines covering profit-loss, averages, ratios, and time-work problems, ensuring correctness via transparent algorithmic procedures. The neural layer converts the structured symbolic output into coherent, human-readable step-by-step explanations, delivering pedagogical value beyond mere answer provision. An EasyOCR-based multimodal input pipeline allows students to photograph textbook problems and submit them directly to the system, extending accessibility across diverse learning contexts. The system is deployed as a lightweight Streamlit web application operable on standard consumer hardware without GPU requirements, making it suitable for resource-constrained educational environments.

Experimental evaluation across 150 unseen aptitude problems demonstrates 94.7% overall accuracy, average response times under 0.5 seconds, and a peak memory footprint below 1 GB. Human evaluators rated explanation quality at 4.5 out of 5 across clarity, completeness, and pedagogical usefulness—outperforming standalone symbolic or neural approaches in all dimensions. A real-world pilot with 45 competitive-exam aspirants confirmed improved conceptual understanding and practical usability, validating AptiNova's potential as a production-grade intelligent tutoring system.

Objectives:

1. Develop a hybrid symbolic-neural reasoning engine that guarantees mathematical correctness through deterministic domain-specific solvers while simultaneously generating coherent, step-by-step natural language explanations via a LoRA fine-tuned language model.
2. Implement an EasyOCR-powered multimodal input pipeline capable of extracting aptitude questions from photographs of textbooks or examination papers, enabling seamless image-to-solution workflow without manual transcription.
3. Train the neural explanation module on a curated hybrid dataset combining real-world extracted aptitude problems and synthetically generated samples to ensure broad category coverage, linguistic diversity, and robust generalization.
4. Deploy a resource-efficient, production-ready web application using the Streamlit framework that operates within a 1 GB memory footprint on standard CPU-based hardware, ensuring accessibility for institutions without high-performance computing infrastructure.
5. Validate the integrated system through quantitative benchmarking (accuracy, response latency, memory utilization) and qualitative human evaluation (explanation clarity, completeness, and pedagogical usefulness) against standalone symbolic and neural baseline approaches.

Methodology:

AptiNova is organized around a three-layer pipeline: (i) multimodal input processing, (ii) dual solving pathways, and (iii) natural language output generation.

Input Processing and Multimodal Integration

The input layer accepts both text queries and image uploads. For images, EasyOCR employs the CRAFT text-detection network for character localization followed by an attention-based recognition network for transcription. The extracted text passes through a post-processing pipeline that normalizes whitespace, corrects capitalization, and removes OCR artifacts before entering the main processing flow.

Topic Detection and Routing

A lightweight rule-driven classifier performs keyword matching against curated domain lexicons to categorize each query. Profit-Loss-Discount problems are detected via terms such as profit, loss, discount, cost price, and selling price; Average problems through average, mean, and sum; Ratio-Proportion problems through ratio and proportion; and Time-Work problems through work, days, together, alone, and complete. Regular expressions simultaneously extract numerical values. Unmatched queries are routed to the neural fallback pathway.

Symbolic Solver Engine

Matched queries are processed by domain-specific deterministic modules. Each module implements the canonical mathematical formulas for its category—for example, the Profit-Loss engine computes Profit% = ((SP − CP) / CP) × 100—and produces both a numerical result and a structured intermediate computation trace. This trace is directly consumed by the explanation generation module, ensuring transparency and auditability of every arithmetic step.

Neural Language Model Architecture and Training

The neural component is built on DistilGPT-2 (≈82 M parameters), adapted using LoRA with rank $r = 8$ and scaling factor $\alpha = 32$. Training utilized a hybrid dataset of real-world extracted problems (via PDFPlumber and Tesseract OCR) supplemented by 800 synthetically generated samples across four categories (200 per domain). The fine-tuning configuration employed AdamW optimizer ($lr = 2 \times 10^{-4}$), batch size 8 with gradient accumulation to effective batch size 32, 3 epochs, cosine learning rate scheduling, and mixed-precision (FP16) training on a T4 GPU. At inference, top-p sampling ($p = 0.9$) at temperature 0.7 generates explanations of up to 200 tokens.

Hybrid Decision Logic

The orchestration layer first attempts symbolic resolution. On success (occurring in 92% of evaluated queries), the neural model generates an explanation from the symbolic trace. On failure—typically for ambiguous or novel problem types—the neural model independently performs both reasoning and answer generation. This conditional routing achieves the system's combined accuracy and speed profile: symbolic paths complete in ~0.03 s while neural paths require ~0.7 s with GPU acceleration.

Result and Conclusion:

AptiNova was evaluated on 150 unseen aptitude problems drawn from standard competitive examination materials across four problem categories. The hybrid system achieved 94.7% overall accuracy—100% for Average and Ratio problems, 98.3% for Profit-Loss, and 91.4% for Time-Work—representing a 14.2 percentage-point improvement over the standalone neural baseline (82.9%). Explanation quality, assessed by five education professionals on a five-point scale, averaged 4.6 for clarity, 4.5 for completeness, and 4.4 for pedagogical usefulness, with BLEU score analysis confirming 23% higher coherence than purely neural outputs. Average hybrid response time was 0.47 s with GPU and 0.89 s on CPU-only configurations, representing a 67% latency reduction over neural-only processing. The total system memory footprint remained below 1 GB, confirming suitability for deployment on standard educational hardware.

Comparative benchmarking demonstrates that AptiNova outperforms Mathway (91.5% accuracy) and Pure GPT-3.5 (84.2%) while matching or exceeding PhotoMath (88.3%) in accuracy—and surpasses all three in response speed with integrated OCR support. A pilot deployment with 45 students over one month showed that approximately 90% reported improved problem-solving understanding and 75% preferred AptiNova over video tutorials for step-by-step guidance.

In conclusion, AptiNova successfully demonstrates that a hybrid symbolic-neural architecture can overcome the complementary weaknesses of rule-based and data-driven paradigms in educational AI. By pairing deterministic mathematical correctness with fluent natural language explanation generation, and extending accessibility through multimodal OCR input, the system offers a technically rigorous and pedagogically effective solution for aptitude learning that is deployable in resource-constrained real-world environments.

References:

1. [1] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, "LoRA: Low-rank adaptation of large language models," in Proc. Int. Conf. Learn. Representations (ICLR), 2022.
2. [2] J. Wei, X. Wang, D. Schuurmans, M. Bosma, B. Ichter, F. Xia, E. Chi, Q. Le, and D. Zhou, "Chain-of-thought prompting elicits reasoning in large language models," in Advances in Neural Information Processing Systems, vol. 35, pp. 24824–37, 2022.
3. [3] A. S. d'Avila Garcez and L. C. Lamb, "Neurosymbolic AI: The 3rd wave," Artificial Intelligence Review, vol. 56, no. 11, pp. 12387–406, 2023.
4. [4] T. Dettmers, A. Pagnoni, A. Holtzman, and L. Zettlemoyer, "QLoRA: Efficient finetuning of quantized LLMs," in Advances in Neural Information Processing Systems, vol. 36, 2023.
5. [5] K. Cobbe et al., "Training verifiers to solve math word problems," arXiv preprint arXiv:2110.14168, 2021.
6. [6] A. Patel, S. Bhattamishra, and N. Goyal, "Are NLP models really able to solve simple math word problems?" in Proc. NAACL-HLT, pp. 2080–2094, 2021.

Future Scope:

The future scope of this project includes:

1. Domain Expansion — The modular solver architecture is designed for incremental extension. Subsequent phases will incorporate additional aptitude categories including speed-distance-time, simple and compound interest, permutations and combinations, probability, and data interpretation, without requiring changes to the core hybrid decision logic.
2. Multilingual and Voice-Enabled Interface — Integrating multilingual NLP pipelines and text-to-speech/speech-to-text modules will broaden accessibility for non-English-speaking learners and students with reading difficulties, aligning with inclusive education frameworks.
3. Adaptive Personalized Learning — Incorporating a user profiling module that tracks historical performance, identifies conceptual weak areas, and dynamically adjusts question difficulty and explanation depth will transform AptiNova from a static solver into a fully adaptive intelligent tutoring system.
4. Mobile and Offline Deployment — Quantizing the LoRA-adapted model to INT4/INT8 precision and packaging it as a Progressive Web App (PWA) will enable on-device inference without internet connectivity, addressing infrastructure limitations in rural and semi-urban educational institutions.
5. Cloud-Integrated Analytics and LMS Compatibility — Connecting AptiNova to cloud databases for performance analytics, integrating with Learning Management Systems (LMS) such as Moodle, and providing educator dashboards for tracking cohort progress will enhance institutional utility and support data-driven pedagogical decisions.