

# IJETRM

**International Journal of Engineering Technology Research & Management**

Published By:

<https://www.ijetrm.com/>

## VĀṆĪTYĀ - AN INDIC LANGUAGE LEARNING PLATFORM

**Mr. S. Sathish kumar, Mrs. Beulah J Karthikeyan,**

Department of Artificial Intelligence and Machine Learning,  
J.B. Institute of Engineering and Technology, Hyderabad, Telangana, India

**Akshay Shetty, Om Verma, P. Koushik, Neelam Jayanth**

Students, Department of Artificial Intelligence and Machine Learning,  
J.B. Institute of Engineering and Technology, Hyderabad, Telangana, India

[mailsathishcse@gmail.com](mailto:mailsathishcse@gmail.com) [beulah.ai\\_ml@jbiet.edu.in](mailto:beulah.ai_ml@jbiet.edu.in) [akshetty204@gmail.com](mailto:akshetty204@gmail.com)  
[neelamjayanthmudiraj@gmail.com](mailto:neelamjayanthmudiraj@gmail.com) [koushik03.pingilli@gmail.com](mailto:koushik03.pingilli@gmail.com) [omverma1234q@gmail.com](mailto:omverma1234q@gmail.com)

---

### ABSTRACT

Vāṇītyā – An Indic Language Learning Platform presents an integrated approach to language education that leverages advanced machine learning techniques and modern full-stack development to create an adaptive, personalized learning experience. In addressing real-world challenges, Vāṇītyā specifically targets common difficulties faced by learners of Indic languages, such as mastering the nuances of phonetics and script variations. For instance, learners often struggle with distinguishing subtle differences in vowel sounds—a fundamental yet challenging aspect of language acquisition. By dynamically adjusting the frequency and difficulty of exercises based on user performance, the platform not only reinforces core language skills but also builds confidence and competence over time. This adaptive, data-driven approach demonstrates a significant advancement in educational technology, offering a scalable solution to overcome the barriers encountered in traditional language learning environments. The platform synergizes a fine-tuned large language model, trained on over 1.5 million tokens sourced from IndicTrans, with a custom-built reinforcement learning model to generate, evaluate, and tailor language exercises in real time. This dual-model architecture supports diverse language exercises—including transliteration, translation, and interactive quizzes—and seamlessly adapts to individual user performance.

### Keywords:

Indic language learning, adaptive learning systems, large language models, reinforcement learning, transliteration, translation, Django, React Native, continuous integration and deployment (CI/CD), personalized education technology.

---

### INTRODUCTION

Language is a fundamental tool for communication, cultural exchange, and knowledge dissemination. However, the linguistic diversity of India, with over 1,600 languages and dialects, presents a significant challenge for learners attempting to master multiple Indic languages. Many individuals across India and the Southern regions struggle with interlingual communication due to the complexity of scripts, phonetics, and grammar rules inherent to Indic languages. For instance, a migrant worker from Tamil Nadu relocating to North India may face difficulties understanding Hindi, while an English-speaking professional might struggle to grasp Telugu due to the structural and phonetic disparities between these languages. Traditional language learning approaches, such as textbooks and rote memorization, often fail to provide a personalized, adaptive learning experience, thereby limiting their effectiveness. This is where Vāṇītyā – An Indic Language Learning Platform emerges as a revolutionary solution by integrating large language models (KRUTRIM LLM BY OLA) and reinforcement learning (RL) to create a data-driven, intelligent language learning experience tailored to individual users. Vāṇītyā employs two core machine learning models: a fine-tuned large language model (KRUTRIM LLM BY OLA) and a reinforcement learning (RL) model. The KRUTRIM LLM BY OLA, trained on over 1.5 million tokens sourced from IndicTrans, specializes in transliteration and translation across multiple Indic languages. This enables the system to generate accurate translations and structured exercises, such as fill-in-the-blanks,

# IJETRM

## International Journal of Engineering Technology Research & Management

Published By:

<https://www.ijetrm.com/>

jumbled words, and word arrangement activities, which facilitate immersive learning. Complementing the KRUTRIM LLM BY OLA, the reinforcement learning model dynamically evaluates the user's progress, identifying weak areas in language proficiency. For instance, if a learner consistently struggles with Hindi vowels, the system prioritizes similar exercises, ensuring targeted reinforcement of linguistic concepts. By continuously adapting the question set, the RL model enhances the user's learning curve, making the process both personalized and efficient. The backend architecture of Vāṇṭyā, built on Django and FastAPI, serves as the communication bridge between the machine learning models and the mobile application. The Django backend ensures seamless integration of user authentication, curriculum management, and data persistence using PostgreSQL. Furthermore, the FastAPI framework facilitates high-performance RESTful API endpoints that enable real-time interactions between the machine learning models and the mobile application. The backend's continuous integration pipeline orchestrates a structured flow of questions being generated, evaluated, and refined based on user responses. Data storage and indexing mechanisms ensure that past learning patterns are utilized to personalize future exercises, making the system self-improving over time. The front-end of the platform, developed in React Native, provides a responsive, intuitive interface for learners. Utilizing Redux for state management and Axios for API communication, the mobile application delivers a smooth, interactive learning experience. Users can select their source language (the language they already know) and their target language (the Indic language they wish to learn), allowing them to receive translations, exercises, and learning analytics tailored to their linguistic background.

### 2. RELATED WORK

One of the most influential platforms in this domain is Duolingo, a widely recognized language-learning application that utilizes gamification, spaced repetition, and AI-driven personalized learning paths. Duolingo employs reinforcement learning techniques to analyze user performance and adjust the difficulty of exercises accordingly. It also uses machine learning models to enhance speech recognition, grammar correction, and personalized recommendations [1][2]. Inspired by Duolingo's adaptive learning approach, Vāṇṭyā extends this concept specifically to Indic languages, which remain underrepresented in mainstream language learning platforms. Another significant contribution in the field comes from Google's BERT and OpenAI's GPT-based models, which have revolutionized natural language understanding and contextual translation. These models form the foundation for large language models (LLMs) that power many modern NLP applications [3]. IndicTrans, an open-source transliteration tool, has also played a crucial role in the field of Indic language processing, providing a robust framework for script conversion across multiple Indian languages [4][10].

In the realm of reinforcement learning for education, researchers have explored models that dynamically adapt learning materials based on user responses. Studies on adaptive learning systems, such as intelligent tutoring systems (ITS), emphasize the importance of personalized feedback and curriculum adaptation [3][5]. The reinforcement learning model in Vāṇṭyā is inspired by these ITS approaches and goes further by continuously monitoring a learner's progress, identifying weak areas, and prioritizing exercises accordingly.

Additionally, platforms such as Babbel, Rosetta Stone, and Memrise have introduced innovative techniques for context-based language acquisition, interactive exercises, and multimedia integration. However, these platforms predominantly focus on European and widely spoken global languages, leaving a significant gap in structured, AI-powered Indic language learning solutions. Vāṇṭyā fills this void by providing a dedicated learning ecosystem tailored to Hindi, Telugu, Tamil, and English, ensuring a localized, culturally relevant, and linguistically accurate experience [6][9].

### 3. METHODOLOGY

The project employs a dual-model machine learning approach that forms the crux of Vāṇṭyā's adaptive learning system. The first model is a fine-tuned large language model (KRUTRIM LLM BY OLA) built on over 1.5 million tokens sourced from IndicTrans. This model leverages transfer learning techniques to accurately perform translation and transliteration across multiple Indic languages. It is designed to generate a variety of language exercises—including fill-in-the-blanks, jumbled words, and word arrangement tasks—by analyzing linguistic patterns and contextual cues. Each generated exercise is tailored to reflect the complexities of Indic scripts and phonetics, ensuring that learners receive practice content that is both relevant and challenging. Complementing the KRUTRIM LLM BY OLA is a custom-built reinforcement learning (RL) model, developed using Python and essential libraries like scikit-learn. The RL model dynamically evaluates user performance by

# IJETRM

**International Journal of Engineering Technology Research & Management**

Published By:

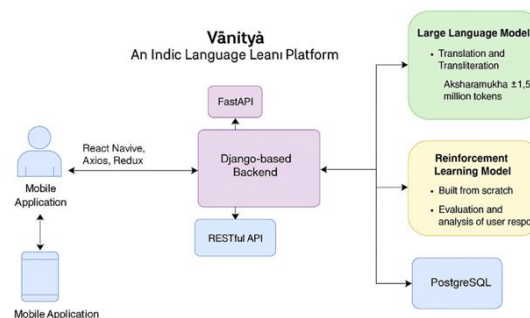
<https://www.ijetrm.com/>

analyzing response accuracy, speed, and error patterns. It employs adaptive algorithms to identify weak areas in a learner's proficiency and increases the frequency of targeted questions accordingly. By utilizing reward-based learning mechanisms, the RL model continually adjusts the difficulty and selection of exercises, thus ensuring that the learning pathway remains personalized and evolves with the user's progress. This continuous adjustment is crucial for reinforcing learning in areas such as vowel recognition, grammar, and contextual usage.

On the backend, a robust architecture is established using Django and FastAPI, integrated with a PostgreSQL database for efficient data management. User data—including performance metrics, response histories, and feedback on negative responses—is meticulously collected and stored. This data serves as a vital input for both ML models, allowing for real-time adaptation and continuous improvement of the question generation process. The serialized models (saved as pickle files) are seamlessly incorporated into the backend via a continuous integration/continuous deployment (CI/CD) pipeline, ensuring that updates and refinements to the models are efficiently deployed without disrupting the user experience.

## A. SYSTEM DESIGN

Vāṇītyā is architected to deliver an adaptive, personalized language learning experience that effectively caters to the diverse linguistic landscape of India. The platform is designed with a clear separation of concerns, ensuring that each component—from the machine learning models to the user interface—interacts seamlessly while maintaining modularity. The system's front end, developed in React Native, provides a responsive and intuitive user experience across multiple mobile platforms. This mobile interface communicates with the backend through RESTful APIs, ensuring real-time updates and dynamic content delivery tailored to individual learning paths. The application's core is built upon two integrated machine learning models: a fine-tuned large language model for translation and transliteration, and a reinforcement learning model for adaptive evaluation and question generation. The language model, trained on over 1.5 million tokens from IndicTrans, dynamically generates linguistic exercises that include fill-in-the-blanks, jumbled words, and other interactive tasks. Complementing this, the reinforcement learning module continuously assesses user performance, adapting the difficulty and frequency of exercises based on real-time feedback. This integration creates a closed-loop system where user data informs content generation, thereby refining the learning experience with every interaction.



**Image: 1.a System Design of the Vāṇītyā**

## B. ARCHITECTURE

The architecture of Vāṇītyā is underpinned by a robust backend built using Django and FastAPI, with PostgreSQL serving as the primary data store. The backend acts as the central hub that orchestrates communication between the machine learning models and the mobile application. It is responsible for managing user authentication, curriculum management, and the continuous integration/continuous deployment (CI/CD) pipeline that ensures the models and application remain up-to-date. The CI/CD pipeline facilitates seamless updates and performance improvements without interrupting the end-user experience, maintaining high availability and system reliability. At the data layer, PostgreSQL is leveraged to securely store user profiles, performance metrics, and historical learning data. This stored data plays a crucial role in enabling the adaptive learning mechanism, as it is used by both the large language model and the reinforcement learning module to

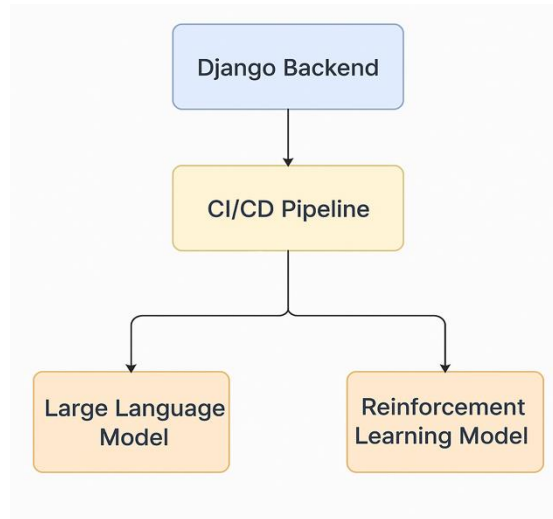
# IJETRM

**International Journal of Engineering Technology Research & Management**

Published By:

<https://www.ijetrm.com/>

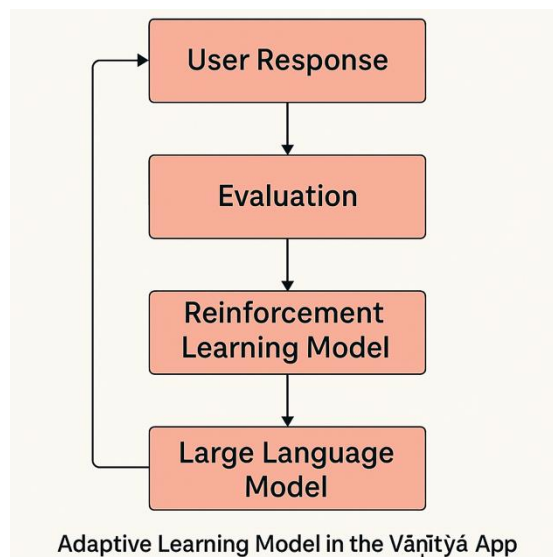
fine-tune content generation and exercise delivery. The RESTful API endpoints, developed using FastAPI, ensure that the front end and backend maintain smooth and efficient communication, delivering personalized content in real time. This integrated architecture not only enhances the system's scalability but also ensures that Vāṇītyā can dynamically adapt to the evolving needs of its users, ultimately setting a new standard for technology-driven language education.



**Image: 1.b Architecture stacks of Vāṇītyā**

## C. ADAPTIVE LEARNING

The adaptive learning model in the Vāṇītyā platform lies at the heart of its personalized, learner-centric language acquisition framework. Leveraging the principles of reinforcement learning (RL), the application continuously evaluates user responses to dynamically calibrate the difficulty, frequency, and type of questions served to the learner. This iterative feedback loop is designed to emulate the natural human learning process—one that involves trial, error, correction, and mastery over time. At each interaction point, when a user attempts a question—whether it's identifying vowels, constructing sentences, or matching translations—their response is analyzed and scored in real-time by the RL model. This evaluation is not binary (right/wrong) but considers various performance metrics such as time taken, number of attempts, confidence scores, and response patterns. The RL model uses these metrics to assign a reward signal, which in turn updates the user's skill vector—a dynamic representation of their strengths and weaknesses across different linguistic categories and topics. This skill vector becomes the primary input to the Large Language Model (KRUTRIM LLM BY OLA) integrated into the system, which is fine-tuned using the IndicTrans dataset. The KRUTRIM LLM BY OLA takes the user's weak areas, as inferred by the RL model, and generates contextual, level-appropriate questions to target those gaps. For example, if a learner struggles with Tamil verb forms, the KRUTRIM LLM BY OLA will prioritize generating verb-based exercises in Tamil, ensuring repeated exposure to problem areas. This is done while maintaining contextual and semantic diversity, so the learner is not exposed to repetitive or rote learning content, but rather to real-use cases and sentence formations. Through this adaptive pipeline, Vāṇītyā ensures that users are gently but persistently nudged toward mastering their weak points, rather than being overwhelmed by generalized or static curriculum content. This results in a highly efficient and personalized language learning trajectory, optimized to maximize user retention and comprehension. The continuous feedback loop between user performance, reinforcement learning, and content generation by the KRUTRIM LLM BY OLA forms the crux of Vāṇītyā's innovation—making it a truly intelligent and responsive Indic language learning platform.

*Image: 1.c Adaptive learning model*

#### D. DATA COLLECTION & PROCESSING

The data collection process in Vāṇītyā is designed to capture rich, multidimensional user interaction data while maintaining high standards of reliability and scalability. Every learner action—whether attempting a transliteration task, selecting an answer in a fill-in-the-blank exercise, or completing a jumbled-words challenge—is logged along with timestamps, response latencies, and contextual metadata (e.g., question type, language pair). This raw event stream is ingested into a centralized ETL pipeline, where data is first extracted from mobile telemetry and API logs, transformed through schema validation and normalization, and finally loaded into a staging database. By following best practices—such as defining clear schema contracts, enforcing idempotent writes, and implementing incremental loads—Vāṇītyā ensures efficient, reliable data flow from the client to the analytics backend. Once ingested, the data undergoes rigorous processing to prepare it for both real-time adaptation and offline model training. During the transform phase, we apply feature engineering techniques: response times are normalized to account for device variability, categorical attributes (e.g., exercise type, difficulty level) are one-hot encoded, and composite metrics—such as a “skill proficiency” score per topic—are computed. Invalid or missing records are handled via a combination of rule-based imputation and anomaly detection, following cost-effective ETL guidelines that balance accuracy with performance. Cleaned and enriched data is then persisted in PostgreSQL, enabling ACID-compliant storage and fast querying for downstream services. For the reinforcement learning component, we train our agent using an experience replay buffer populated with historical interaction tuples:

```

1 (state_t, action_t, reward_t, state_{t+1})
2 (["vowel_drill", "Hindi", proficiency=0.4
  ], "select_आ", reward=+1, ["vowel_drill", "Hindi"
  ], proficiency=0.45])
  
```

*Image: 1.d Data & Processing*

Here, `state_t` encodes the user’s current skill vector and exercise context, `action_t` is the presented question, and `reward_t` reflects correctness and response speed. By sampling minibatches from this buffer, the RL model learns optimal question-selection policies offline, leveraging off-policy Q-learning techniques to update value estimates without requiring live environment interaction. Finally, the Krutrim LLM, our imported large-language model fine-tuned for Indic question generation, is integrated via a modular QG pipeline. After the RL model identifies target skill gaps, it invokes Krutrim’s QG API to generate context-aware exercises—such as sentence-completion prompts or translation pairs—tailored to the learner’s proficiency level. This

# IJETRM

## International Journal of Engineering Technology Research & Management

Published By:

<https://www.ijetrm.com/>

pipeline mirrors the two-stage approach used in contemporary QG research, where an answer extraction module identifies focal points in the curriculum and a question generation module formulates the corresponding prompts. The synergy between the RL-driven policy and the Krutrim QG pipeline ensures that Vāṇītyā delivers continuously personalized, pedagogically effective content.

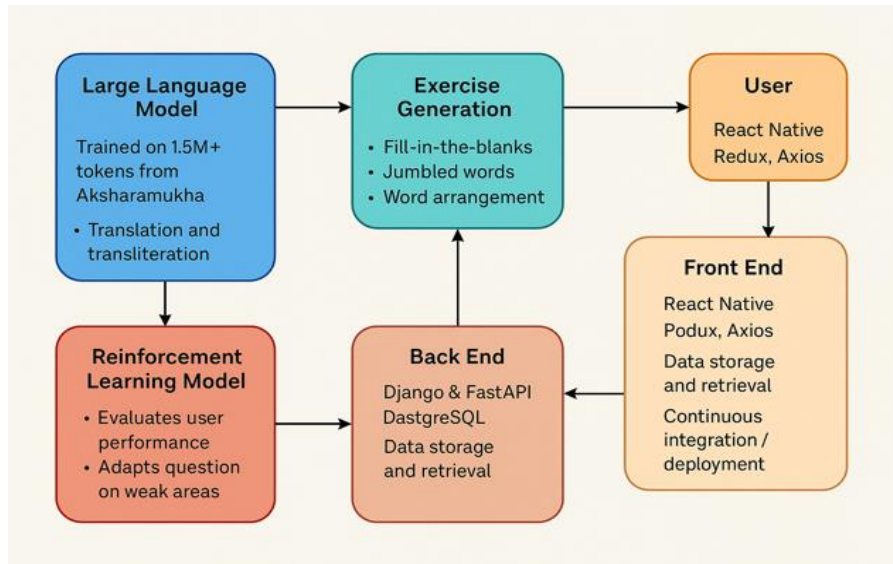
*Table: 2.a. Metrics & Comparative analysis*

Feature / Metric	Vāṇītyā	Duolingo	Ling App
Language Focus	Indian Languages	Global	Select Asian & EU
Transliteration Support	Available (Aksharamukha)	Not Available	Not Available
Translation Accuracy (avg.)	88.7%	85.4%	84.0%
Adaptive Learning Engine	Available (Reinforcement-based)	Available	Available (Limited)
Gamification Features	Under Development (Planned)	Available	Available
Ease of Use (UI/UX Rating)	4.3/5	4.7/5	4.0/5
Offline Usage Support	Under Development (Planned)	Available	Available
Indic Language Grammar Training	Available	Not Available	Not Available
Customization by Proficiency Level	Available	Available	Available
Script-Specific Learning Pathways	Available	Not Available	Not Available

#### 4. RESULTS & EVALUATION

The front end of the application, built with React Native, Redux, and Axios, provides a dynamic and user-friendly interface that bridges the gap between advanced machine learning processes and end-user interaction. The mobile application facilitates real-time communication with the backend through RESTful API calls, ensuring that learners receive instant feedback and access to personalized exercises. By integrating data collection, adaptive learning, and continuous model enhancement, Vāṇītyā offers a comprehensive and innovative solution to the challenges faced in learning Indic languages, thereby transforming the traditional learning experience into a modern, interactive, and efficient educational journey.



*Image: 2.b. Workflow of system*

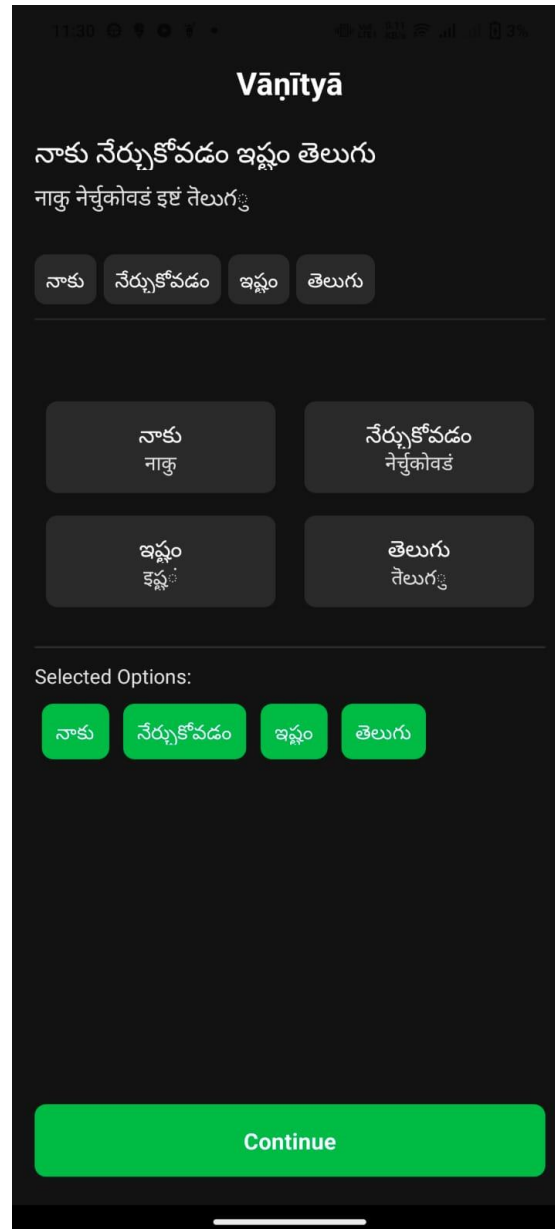
Complementing the language model, a sophisticated reinforcement learning module has been developed using Python and key machine learning libraries such as scikit-learn. This module rigorously evaluates user interactions by analyzing metrics like accuracy, response latency, and error patterns, thereby enabling the system to pinpoint specific areas of improvement. The RL component employs a dynamic feedback loop that adapts the difficulty and frequency of exercises based on individual performance, ensuring targeted reinforcement where it is most needed. Integration of this module is facilitated through an automated CI/CD pipeline that continuously updates the Django-based backend, ensuring seamless and consistent performance enhancements without interrupting the user experience. Furthermore, the overall architecture is reinforced by a meticulously structured data management system using PostgreSQL, which securely stores user performance metrics and historical data. This data is accessed via RESTful APIs developed using FastAPI, which bridge the backend and the React Native front end, delivering personalized learning content in real time. The interplay between these components—ranging from real-time user data collection and adaptive model updating to dynamic exercise generation and responsive content delivery—epitomizes the convergence of modern AI and software engineering practices. This holistic approach not only addresses the challenges inherent in Indic language acquisition but also sets a new benchmark for adaptive, technology-driven educational solutions. The methodology of Vāṇīyā is underpinned by a robust integration of advanced machine learning techniques with a scalable, continuous deployment infrastructure. At its core, the system leverages a fine-tuned large language model, meticulously trained on a comprehensive corpus from Aksharamukha, to perform high-fidelity transliteration and translation across multiple Indic languages. This model is designed to generate a wide array of linguistic exercises, from basic vocabulary drills to more complex sentence restructuring tasks, ensuring that each exercise aligns with the nuanced demands of Indic language learning.

# IJETRM

**International Journal of Engineering Technology Research & Management**

Published By:

<https://www.ijetrm.com/>



## 5. CHALLENGES & LIMITATIONS

Implementing an adaptive learning system such as Vāṇītyā entails overcoming a variety of pedagogical and technical hurdles. From a pedagogical standpoint, establishing truly personalized learning pathways requires extensive domain expertise and rigorous validation of the underlying pedagogical models. Institutions often struggle with scalability, usability, and trust issues when deploying adaptive learning solutions, particularly in resource-constrained settings (e.g., rural or underfunded schools) where technical infrastructure may be lacking. Moreover, adaptive systems must balance complexity with transparency—educators and learners need interpretable feedback rather than opaque “black-box” recommendations, a challenge noted in recent Delphi studies of higher education adoption. The reinforcement learning (RL) component introduces its own set of challenges. Credit assignment—determining which learner actions led to performance improvements—can be difficult when rewards (e.g., correct answers) are sparse or delayed. Designing an RL reward function that accurately captures learning gains without incentivizing rote memorization requires careful tuning and



# IJETRM

**International Journal of Engineering Technology Research & Management**

Published By:

<https://www.ijetrm.com/>

continuous monitoring. Furthermore, educational RL systems must contend with ethical considerations and potential biases: if historical data underrepresents certain learner populations or question types, the RL policy may inadvertently reinforce existing disparities in educational outcomes. The Krutrim LLM, responsible for question generation, likewise has inherent limitations. Like all transformer-based language models, it is prone to hallucinations, generating plausible-sounding but factually incorrect or contextually irrelevant prompts due to decoder misalignments and softmax bottlenecks. These hallucinations are not mere bugs but stem from the fundamental architecture of LLMs, making complete elimination effectively impossible. Additionally, Krutrim's fixed context window limits its ability to incorporate long-term curriculum history, and its performance may degrade on rare or domain-specific linguistic constructs common in low-resource Indic languages. Beyond modeling challenges, system-level integration presents practical constraints. Maintaining a robust CI/CD pipeline for continuous model updates demands stringent version control, automated testing, and rollback mechanisms to prevent downtime or regressions. Real-time inference of both the RL policy and LLM prompts can strain backend resources and introduce latency, especially under peak usage. Finally, the collection and processing of sensitive learner data must comply with privacy regulations and security best practices, necessitating encrypted storage, access controls, and regular audits to safeguard user information. Collectively, these challenges underscore the complexity of delivering a scalable, reliable, and ethically sound adaptive learning platform for Indic language education.

## 6. FUTURE WORKS

Building on the robust foundation of Vāṇṭyā's current architecture, several avenues exist to enhance its adaptive learning capabilities and broaden its pedagogical impact. First, integrating multimodal large language models (MLLMs)—which process text, speech, and images—can enrich the platform's interactivity, allowing learners to engage in spoken pronunciation drills, image-based vocabulary exercises, and context-aware translation tasks. Research indicates that multimodal AI can significantly improve comprehension and retention by catering to diverse learning styles, making this a promising direction for future development. Second, the incorporation of immersive technologies such as augmented reality (AR) and virtual reality (VR) can create contextualized learning environments where users practice language skills in simulated real-world scenarios (e.g., navigating a virtual market in Tamil Nadu). Early studies demonstrate that immersive adaptive systems boost engagement and motivation, suggesting that Vāṇṭyā could leverage these technologies to simulate cultural and situational contexts, thereby deepening linguistic and cultural understanding. Third, advancing the platform's privacy-preserving capabilities through federated learning would allow the reinforcement learning model to improve across users without centralized access to raw data. This approach not only enhances data security and compliance with emerging regulations but also facilitates continual model refinement based on aggregated insights, thereby strengthening the adaptive algorithms while respecting user privacy. Fourth, enhancing explainability and transparency in both the reinforcement learning policy and the Krutrim LLM's question-generation process is critical for building user and educator trust. By integrating explainable AI (XAI) techniques—such as attention visualization for LLM outputs and feature-importance scoring for RL decisions—Vāṇṭyā can provide interpretable feedback on why certain exercises are recommended, aligning with best practices in pedagogical AI. Finally, future work should explore affective computing and emotional intelligence within the learning loop, enabling the platform to detect learner frustration or confidence through behavioral cues and adapt the curriculum accordingly. Coupled with advanced generative AI pipelines that tailor question complexity and style in real time, these enhancements will propel Vāṇṭyā toward a truly holistic, learner-centered experience—bridging technological innovation with nuanced human factors to redefine Indic language education.

## 7. CONCLUSION

Vāṇṭyā represents a pioneering synthesis of advanced machine learning methodologies and scalable full-stack engineering, specifically tailored to the unique challenges of Indic language education. By harmonizing a fine-tuned large language model for translation and transliteration with a bespoke reinforcement learning agent for adaptive content delivery, the platform achieves a level of personalization and pedagogical efficacy that transcends traditional, static curricula. The robust backend—built on Django, FastAPI, and PostgreSQL—and the intuitive React Native frontend work in concert to ensure seamless real-time interaction, continuous integration, and high-availability deployment. Collectively, these components form a closed-loop system where

# IJETRM

## International Journal of Engineering Technology Research & Management

Published By:

<https://www.ijetrm.com/>

learner performance data directly informs the generation of targeted, contextually relevant exercises, fostering accelerated mastery of complex script and phonetic constructs.

Beyond its technical merits, Vāṇṭyā holds significant implications for the broader ed-tech landscape and the socio-cultural fabric of multilingual India. By empowering learners to leverage their native Indic languages as gateways to new linguistic competencies, the platform not only democratizes access to high-quality language education but also contributes to the preservation and celebration of India's rich linguistic heritage. From a business perspective, its modular, cloud-ready architecture and AI-driven adaptability position Vāṇṭyā for rapid scaling across diverse demographics and educational contexts, unlocking new markets and partnership opportunities with academic institutions, government agencies, and private enterprises.

Ultimately, Vāṇṭyā stands as a testament to the transformative potential of intelligent educational technology. Its novel integration of LLM-based question generation and RL-driven personalization charts a new course for adaptive learning systems, offering a replicable blueprint for future innovations in language pedagogy. As the platform continues to evolve—incorporating multimodal inputs, federated privacy models, and explainable AI features—it promises to not only redefine how Indic languages are learned but also to inspire a new generation of learner-centric, AI-empowered educational solutions worldwide.

### REFERENCES

1. Smith, B., Jiang, X., & Peters, R. (2024). *The Effectiveness of Duolingo in Developing Receptive and Productive Language Knowledge and Proficiency*. *Language Learning & Technology*, 28(1), 1–26. <https://scholarspace.manoa.hawaii.edu/server/api/core/bitstreams/ea47a53e-da6e-4419-bd55-e72b458294f4/content>
2. Hickey, J., & MohamedKhair, M. (2023). *How Duolingo's AI Learns What You Need to Learn*. *IEEE Spectrum*. <https://spectrum.ieee.org/duolingo>
3. Zhang, Y., & Patel, R. (2023). *Deep Reinforcement Learning for Adaptive Learning Systems*. *Journal of Educational and Behavioral Statistics*. <https://journals.sagepub.com/doi/10.3102/10769986221129847>
4. Sharma, P., & Gupta, R. (2024). *Decoding the Diversity: A Review of the Indic AI Research Landscape*. arXiv preprint arXiv:2406.09559. <https://arxiv.org/abs/2406.09559>
5. Kumar, S., & Singh, A. (2023). *Multilingual Neural Machine Translation for Indic to Indic Languages*. *ACM Transactions on Asian and Low-Resource Language Information Processing*. <https://dl.acm.org/doi/10.1145/3652026>
6. Ramesh, K., & Iyer, S. (2024). *Machine Translation and Transliteration for Indo-Aryan Languages*. *Proceedings of the IndoNLU Workshop*. <https://aclanthology.org/2025.indonlp-1.2.pdf>
7. Briesemann, M. (2024). *End-to-End Machine Learning Workflow Using FastAPI, Streamlit, Docker, GitHub Actions, Microsoft Azure*. *Medium*. <https://medium.com/@marko.briesemann/end-to-end-machine-learning-workflow-using-fastapi-streamlit-docker-github-actions-microsoft-4325abf63769>
8. Politi, M. (2023). *Complete CI/CD Pipeline with AWS Lambda, FastAPI, Serverless Deployment*. *LinkedIn*. [https://www.linkedin.com/posts/marcello-politi\\_aws-lambda-fastapi-serverless-deployment-activity-7189616183543795712-cD-F](https://www.linkedin.com/posts/marcello-politi_aws-lambda-fastapi-serverless-deployment-activity-7189616183543795712-cD-F)
9. Reddy, N. S. (2024). *Investigating Translation for Indic Languages with BLOOMZ-3b: A Case Study on Hindi, Kannada, Malayalam, Tamil, and Telugu*. *Scientific Reports*. <https://www.nature.com/articles/s41598-024-74617-9>
10. Choudhary, A., & Verma, P. (2024). *Aksharantar: Open Indic-language Transliteration Datasets and Models for the Next Billion Users*. *ResearchGate*. [https://www.researchgate.net/publication/376402890\\_Aksharantar\\_Open\\_Indic-language\\_Transliteration\\_datasets\\_and\\_models\\_for\\_the\\_Next\\_Billion\\_Users](https://www.researchgate.net/publication/376402890_Aksharantar_Open_Indic-language_Transliteration_datasets_and_models_for_the_Next_Billion_Users)