

JAD-IT-UZ: AN AI-POWERED PLATFORM FOR TEACHING ITALIAN AS A FOREIGN LANGUAGE WITH DOMAIN-SPECIFIC VOCABULARY TRACKS

Ganisher Rakhimov

Doctor Sciences, Professor

Uzbekistan State University of World Languages

rakhimovdtpi@gmail.com

Abstract. This paper presents the architecture and design rationale of JAD-IT-UZ, an AI-powered platform for teaching Italian as a foreign language, developed jointly by Uzbek and Italian university partners. The platform combines a four-layer client–content–AI–data architecture with a trilingual Uzbek–Italian–English terminology database and domain-specific vocabulary tracks in medicine, tourism, and business. We describe the passive and active AI components (text-to-speech, contextual translation, speech scoring, adaptive tutoring), the CEFR-aligned assessment design, and outline four open technical decisions — cloud versus on-premises inference, native versus cross-platform mobile development, the assessment method, and build-versus-buy for AI components — that will shape the project's development roadmap.

Keywords: artificial intelligence, Italian as a foreign language, CEFR assessment, trilingual terminology database, adaptive learning, speech recognition, e-learning platform, domain-specific vocabulary.

Abstrakt. Ushbu maqolada O'zbekiston va Italiya universitet hamkorlari tomonidan birgalikda ishlab chiqilayotgan, italyan tilini xorijiy til sifatida o'rgatuvchi sun'iy intellektga asoslangan JAD-IT-UZ platformasining arxitekturasi va loyihalash mantig'i taqdim etiladi. Platforma to'rt qatlamli (mijoz–kontent–sun'iy intellekt–ma'lumotlar) arxitekturani o'zbek–italyan–ingliz uch tilli terminologik ma'lumotlar bazasi hamda tibbiyot, turizm va biznes sohalariga oid maxsus lug'at yo'nalishlari bilan uyg'unlashtiradi. Maqolada passiv va faol sun'iy intellekt komponentlari, CEFR asosidagi baholash tizimi tavsiflanadi, shuningdek loyiha rivojlanish yo'l xaritasini belgilaydigan to'rtta ochiq texnik qaror yoritiladi.

Kalit so'zlar: sun'iy intellekt, xorijiy til sifatida italyan tili, CEFR baholash, uch tilli terminologik ma'lumotlar bazasi, moslashuvchan ta'lim, nutqni tanish, elektron ta'lim platformasi, sohaviy lug'at.

Аннотация. В статье представлены архитектура и обоснование проектных решений платформы JAD-IT-UZ — основанной на искусственном интеллекте платформы для обучения итальянскому языку как иностранному, совместно разрабатываемой узбекскими и итальянскими университетами-партнёрами. Платформа объединяет четырёхуровневую архитектуру (клиент–контент–ИИ–данные) с трёхязычной узбекско-итальянско-английской терминологической базой данных и специализированными лексическими модулями по медицине, туризму и бизнесу. Рассматриваются пассивные и активные компоненты ИИ, система оценивания на основе CEFR, а также четыре открытых технических решения, определяющих дорожную карту проекта.

Ключевые слова: искусственный интеллект, итальянский язык как иностранный, оценивание по CEFR, трёхязычная терминологическая база данных, адаптивное обучение, распознавание речи, платформа электронного обучения, отраслевая лексика.

1. Introduction

Italian is taught across a number of Uzbek universities, yet the tools available to learners rarely combine sound language pedagogy with vocabulary relevant to their professional field. Commercial applications such as Duolingo or Babbel offer broad, general-purpose instruction but lack Uzbek–Italian trilingual grounding, domain-specific vocabulary, and assessment aligned to the Common European Framework of Reference for Languages (CEFR). At the same time, recent advances in speech recognition, text-to-speech, and large language models have made it practical to build adaptive tutoring and pronunciation-scoring features on a modest infrastructure budget.

JAD-IT-UZ (Joint AI-based Digital Italian Teaching — Uzbekistan) is a joint Italy–Uzbekistan initiative that responds to this gap. It proposes an AI-powered platform for teaching Italian as a foreign language, with vocabulary tracks tailored to three in-demand domains: medicine, tourism, and business. This paper describes the platform's architecture, its AI/NLP design, its assessment strategy, and the technical decisions that will shape its development between 2026 and 2027.

2. The JAD-IT-UZ Project: Partners and Objectives

The project is led on the Uzbek side by Uzbekistan State University of World Languages (UzSWLU / OʻzDJTU), and on the Italian side by the Università per Stranieri di Perugia and the Università degli Studi di Bergamo. It

is funded through two parallel tracks: an Italian application, which supports curriculum design and content curation by the Italian partners, and an Uzbek Innovation grant, which funds platform development and the server infrastructure hosted at the Uzbek university. The core development team includes a lead programmer alongside academic staff who define content and assessment requirements.

The project's stated objective is to combine linguistic expertise — curriculum design, terminology curation, and pedagogical validation — with technical delivery, producing a platform that is usable across web and mobile clients in Uzbek, Italian, and English.

3. Platform Architecture

The platform is organised into four layers, each communicating only with the layer directly above or below it through a defined API. This separation is a deliberate design rule: no client application queries the database directly, and no lesson module calls a speech-recognition or translation model directly — all such calls pass through the AI engine's own API. This makes it possible to change a vendor, for instance a speech-to-text provider, without touching the client applications.

Layer 1 (client apps) covers the web, iOS, and Android interfaces in Uzbek, Italian, and English. Layer 2 (learning and content modules) covers CEFR assessment, A1–C1 lessons, the lexicon, and the three domain tracks. Layer 3 (the AI/NLP engine) provides passive AI — text-to-speech and dictionary lookups — and active AI — an adaptive tutor, speech scoring, and grammar checking. Layer 4 (data and content) holds the trilingual terminology database, user-progress records, and the media store. Two supporting processes run alongside the four layers: a human curation loop, in which professors and translators validate new terminology before it enters the database, and an admin analytics dashboard that aggregates usage and progress data for the funders' reporting requirements.

4. The Platform from the Learner's Perspective: Three Core Areas

While the technical architecture is organised in layers, the platform presents itself to learners and institutional users as three core, user-facing areas, which map onto the underlying layers described above.

Language courses (Italian and Uzbek). Structured, interactive lessons spanning CEFR levels A1 to C2, covering speaking and listening, grammar and vocabulary, and progress certification. This area corresponds to Layer 2 of the architecture.

Information hub. A comparative resource on facilities, infrastructure, transport, and services in Italy and Uzbekistan — covering education, healthcare, business, tourism, and culture — intended to support learners who are studying, working, or travelling between the two countries. This area extends Layer 2's content model beyond language instruction into practical, up-to-date country information.

Automatic translation tool. An Italian–Uzbek translation tool grounded in lexicographic research, offering context-aware, up-to-date translations in the medical, economic, and touristic domains. This is the learner-facing expression of Layer 3's contextual-translation component and Layer 4's trilingual terminology database.

5. The AI/NLP Engine: Passive and Active Interaction

The platform distinguishes between passive and active AI interaction. Passive AI includes text-to-speech generation for Italian pronunciation of any term in the database, on-demand dictionary lookups with pronunciation and example sentences, and contextual translation lookup, which returns the sense appropriate to the surrounding sentence rather than a raw dictionary hit. Active AI includes an adaptive-sequencing model that recommends a learner's next lesson or exercise based on their history, speech recognition with phoneme-level pronunciation scoring, and a grammar-checking module for free-text exercises that flags errors and suggests corrections.

Not all of these components require custom-trained models from the outset. The recommended approach, consistent with the project's own phased plan — prototype in 2026, refine in 2027 — is to start with an existing large-language-model API for translation and grammar checking and a commercial speech-to-text/text-to-speech API for pronunciation features, then evaluate in the second year whether a custom-trained model is justified by usage data.

6. Domain-Specific Vocabulary and the Trilingual Terminology Database

The platform's vocabulary tracks target three domains chosen for their practical relevance to Uzbek learners of Italian: medicine (clinical vocabulary, patient-interaction phrases, and medical documentation terms), tourism (hospitality, travel logistics, and cultural-heritage vocabulary), and business (commerce, negotiation, correspondence, and administrative Italian).

All domain vocabulary is drawn from a trilingual Uzbek–Italian–English terminology database. Its schema records, for each entry, the term, its translation, pronunciation, an example sentence, a domain tag, a source, and a review status.

Every new term passes through the human curation loop — professors and translators — before it is added to the database, ensuring that the automatic translation tool described in Section 4 remains accurate and context-aware rather than a raw machine-translation output.

7. CEFR-Aligned Assessment Design

Placement and progress within the platform are measured against the Common European Framework of Reference for Languages. Two assessment methods were considered. A fixed-form test is simpler to build, validate, and explain to learners, but presents the same items to every learner and is therefore less precise near level boundaries, requiring a longer test for comparable accuracy. An adaptive, Item-Response-Theory-style test is shorter for the same precision and pinpoints level boundaries more accurately, but requires an item bank calibrated with difficulty parameters and is more complex to build, test, and maintain.

This decision has direct downstream consequences: it determines the scope of both the CEFR-scoring service and the adaptive-sequencing model described in Section 5. Before committing to either method, the project's AI/ML engineer has been directed to trial DILIT's placement test and Duolingo's adaptive-testing flow, since DILIT is the closest existing analogue to JAD-IT-UZ in its CEFR structuring and Italian-language focus.

8. Benchmarking Against Existing Platforms

Design decisions throughout the project are informed by a set of reference platforms. DILIT is the closest analogue, being Italian-specific and CEFR-structured, and serves as the primary reference for placement testing. Duolingo offers a well-tested adaptive placement-test flow, though it is broad rather than domain-specific. Babbel is notable for its structured, conversation-first lesson design. Rosetta Stone offers an immersive, visual approach to vocabulary acquisition. Nuovo Espresso, published by Alma Edizioni, provides an established textbook progression for Italian as a foreign language. Vocabolario Treccani serves as an authoritative Italian lexical reference for validating terminology before it enters the database.

9. Open Technical Decisions

Four technical decisions remain open and require early resolution, as each has significant downstream effects on scope and timeline.

Cloud versus on-premises AI inference. The equipment budget assumes a single on-site server, while speech recognition and text-to-speech are the most compute-intensive workloads in the AI engine. On-premises inference offers full

data control, which matters for the data-privacy policy required before speech features go live with real users, but represents a hard capacity ceiling. A cloud AI API dependency offers elastic capacity and faster prototyping at the cost of recurring fees and data-residency questions for student recordings; a hybrid approach — on-premises for the database and application, cloud for inference — is also compatible with the existing budget.

Native versus cross-platform mobile development. Because mobile development depends on the completion of the web client, this decision mainly affects team effort and timeline rather than the earliest milestones. Native development offers the best audio-recording latency and offline performance for speech scoring, at the cost of maintaining two codebases with a single-programmer team; a cross-platform framework such as Flutter or React Native offers a single codebase and faster delivery, provided the audio pipeline is tested carefully against the speech and text-to-speech components.

The CEFR assessment method. As discussed in Section 7, the choice between a fixed-form and an adaptive test changes the scope of both the scoring service and the personalization model.

Build versus buy for speech and translation AI. Commercial APIs offer the fastest path to a working 2026 prototype and are recommended for the initial phase; custom-trained models offer long-term control over Uzbek–Italian speech accuracy but carry higher upfront cost and timeline risk, and are better reserved for a Year 2 evaluation once usage data exists.

10. Roadmap and Conclusion

Development is sequenced quarterly against the project's own timeline. The second half of 2026 establishes foundations: the API contract, authentication, the content service, the terminology database schema, the server, and the continuous-integration pipeline. The first half of 2027 targets a first working prototype, covering translation, text-to-speech, and speech scoring, along with the dashboard and lexicon interfaces, tested internally. The fourth quarter of 2027 adds progress tracking, the personalization hookup, localization, and the data-privacy policy. The first half of 2027 covers grammar checking, the admin analytics dashboard, the mobile applications, and the pilot study, comparing a control group against an experimental group using the platform. The second half of 2027 focuses on hardening, scaling for wider rollout, and integrating results into final project reporting.

JAD-IT-UZ illustrates how a joint Italy–Uzbekistan partnership can combine linguistic curation with a modest but carefully layered technical

architecture to deliver a trilingual, AI-assisted platform for teaching Italian as a foreign language. Resolving the four open technical decisions early, and grounding design choices in established reference platforms such as DILIT and Duolingo, will be central to keeping the 2026–2027 development roadmap on track.

References

1. Council of Europe. (2020). Common European Framework of Reference for Languages: Learning, Teaching, Assessment – Companion Volume. Council of Europe Publishing.
2. Wainer, H. (Ed.). (2000). Computerized Adaptive Testing: A Primer (2nd ed.). Routledge.
3. Chapelle, C.A., & Sauro, S. (Eds.). (2017). The Handbook of Technology and Second Language Teaching and Learning. Wiley-Blackwell.
4. Nation, I.S.P. (2013). Learning Vocabulary in Another Language (2nd ed.). Cambridge University Press.
5. JAD-IT-UZ Project Consortium. (2025). Technical Task Breakdown and Architecture Specification [Internal project document]. Uzbekistan State University of World Languages & Università per Stranieri di Perugia.