

LinguaLeap Product Requirements Document



Table of Contents

Table of Contents.....	1
Product Overview.....	1
Scope and Objectives.....	2
User Personas.....	3
Jobs To Be Done.....	4
Functional Requirements.....	6
Non-Functional Requirements.....	7
Prototypes.....	7
Constraints and Assumptions.....	9
Milestones and Timeline.....	10
Resource Requirements.....	10
Testing and Quality Assurance.....	12
Metrics and Success Criteria.....	12
Go To Market.....	13
References.....	15

Product Overview

Language-learning platforms today overwhelmingly emphasize passive study: vocabulary drills, grammar flashcards, and isolated listening exercises. Popular apps like Duolingo, Memrise, and HelloChinese prioritize gamified repetition over spontaneous conversation, while tools like Rosetta Stone and Babbel offer structured lessons but minimal real-time interaction. Even platforms designed for speaking, such as HelloTalk or Tandem, rely on unstructured, unpredictable partner availability and lack guided, pedagogically aligned feedback. As a result, most learners cannot reliably build the conversational fluency required for academic or professional contexts.

Overall, these existing language learning apps are not rigorous enough, focus too much on memorizing vocabulary & grammar, and/or don't provide social connection. Students often have sparse opportunities to practice real conversation in a psychologically safe environment, making it difficult to receive specific and real-time feedback, which is necessary for students to improve and maintain fluency. In classroom settings, limited instructor time and uneven participation further restrict opportunities for consistent, personalized speaking practice.

LinguaLeap is an AI-driven language companion that supports real-time, in-person conversations between learners. The system provides immediate feedback, generates prompts, and evaluates performance so students can track measurable improvement. This companion transforms language practice from isolated memorization into authentic, guided interaction that feels both natural and supportive. By combining Elo-based skill matching with structured conversational guidance,

LinguaLeap ensures that learners practice at an appropriate difficulty level, remain engaged, and build real fluency through repeated, meaningful interaction.

Scope and Objectives

Scope

LinguaLeap is building an AI-powered conversation platform that aims to address the core challenges language learners face: speaking anxiety and lack of consistent conversation opportunities. Unlike existing tools that prioritize vocabulary drills or unstructured language exchange, LinguaLeap integrates directly into university courses, providing both students and instructors with measurable progress tracking and psychological safety during practice.

The platform will be accessible through both mobile and desktop interfaces, designed for seamless integration into existing academic workflows. By focusing on the university segment, LinguaLeap positions itself as infrastructure rather than just another app and becomes the default tool that language departments depend on for speaking assessment and practice coordination.

Objectives

Primary Objectives

Strengthening Speaking Habits: Establish a consistent practice routine for students by providing easy-to-access sessions, structured conversation flows, and supportive real-time guidance that encourages repeated, meaningful use.

Expanding University Partnerships and Adoption: Grow LinguaLeap's institutional footprint by supporting university-level deployment, ensuring smooth course integration, and building long-term departmental relationships.

Demonstrating Speaking Improvement: Provide students with actionable feedback and instructors with clear visibility into progress so that measurable gains in oral proficiency can be tracked over time.

Improving Platform Reliability and Cross-Platform Performance: Enhance overall stability and ensure that mobile and desktop experiences meet the performance demands of large-scale academic use.

Secondary Objectives

Building Instructor Alignment and Trust: Deliver tools, dashboards, and reporting that allow instructors to incorporate Elo insights and AI feedback into assessments, participation grading, and course planning.

Enhancing the Learner Community Experience: Create spaces for students to reflect on conversations, share progress, and stay motivated through social accountability and a supportive peer environment.

Include and Not Include

Include	Not Include
AI-powered Elo-based matching system for in-person peer conversations	Comprehensive grammar lessons or vocabulary drills
Real-time conversation guidance with AI-generated prompts and structured flows	Corporate LMS integrations (initial phase)
Post-conversation feedback on pronunciation, grammar, vocabulary, and fluency	Certain gamification features like streaks or points (MVP)
Solo practice modes and chatbot conversations	International market expansion
Instructor dashboard with student progress, Elo tracking, and class-wide patterns	Automated language proficiency certification
Mobile and desktop platforms with feature parity	
Asynchronous text-based exchange	
Scheduling and notifications for regular practice sessions	
University integration tools for course-level and department-wide deployment	

User Personas

Our target users are college students who are learning a language with the goal of becoming conversationally fluent. Many feel a deep connection to their family heritage through learning their chosen language and are concerned they are losing that connection by not practicing or speaking for long periods of time. They are willing to go out of their way to practice and surround themselves with others who speak the language and want to use it in a conversational setting. While some might be taking classes or learned the language through growing up in a multi-language household, our target user wants to continue honing their language capability as well as strengthen their connection to their language's heritage.

Through conducting interviews with university students at Stanford and other schools, we learned that many students feel like current language learning apps like Duolingo lack in-person functionality and conversation practice. While they appreciate the ease and convenience of such apps, they feel like it is surface-level and doesn't have a deeper emotional connection to the culture of the language. Since most students live in proximity, there is a need for a way to easily practice their language, meet fellow learners, and feel a connection to their heritage through language.

Personas

Ari is a strong auditory learner thanks to his theater training and easily grasps pronunciation and grammar, but struggles with vocabulary retention and the leap between beginner and advanced content. He enjoys learning through meaning and story. Understanding the etymology and cultural

context of words makes the language feel alive. Ari values a balance between structured learning (classes) and self-directed exploration (films, articles, conversations).

Jacob is a Stanford student who has been learning Spanish from a young age. He has used Duolingo since childhood, as well as studied Spanish in high school. His main drive to pursue Spanish is to be able to have conversations with his grandparents, and he feels like he doesn't have enough opportunities to use Spanish in a conversational setting.

Kevin is a university student in LA who is studying Chinese to speak with family members and connect with his cultural heritage. He has participated in study abroad & language tables, takes formal classes at his university, and is a regular user of Anki as a study tool. He has previously used Duolingo to study Japanese.

Jobs To Be Done

<i>When...</i>	<i>... I want to ...</i>	<i>... so that I/we can ...</i>
When I'm anxious about speaking in my target language,	I want to practice with a peer who won't judge my mistakes	So that I can build confidence without the fear of embarrassment or correction in the moment.
When I have an upcoming oral exam or presentation,	I want structured conversation practice with someone at my level	So that I can prepare effectively and perform well without freezing up.
When I'm learning a heritage language,	I want to practice conversations about my culture and family	So that I can connect with my relatives and feel closer to my roots.
When I finish a conversation practice session,	I want to receive specific, private feedback on my pronunciation and grammar	So that I can improve without feeling judged during the actual conversation.
When I'm trying to stay consistent with language learning,	I want accountability that doesn't punish me for missing days	So that I can build a sustainable practice habit without guilt or stress.
When I'm frustrated with surface-level app exercises,	I want real conversations with cultural context and meaning	So that the language feels alive and relevant to my actual life.
When I don't know anyone else learning my target language,	I want to be matched with peers at similar proficiency levels	So that I can find practice partners easily without awkward mismatches.
When my professor assigns speaking practice,	I want an easy way to coordinate with classmates and track my progress	So that I can complete assignments efficiently and see measurable improvement.

When I'm between beginner and advanced levels,

I want conversation prompts that match my current ability

So that I'm challenged but not overwhelmed, and can stay engaged.

When I'm preparing for study abroad or work opportunities,

I want to track my speaking progress with concrete metrics

So that I can prove my fluency and feel confident in professional or academic settings.

User Story Maps



Grace just started her third quarter of Russian. She's excited... but still freezes every time she tries speaking to her grandmother



She practices on Duolingo, but 'The cat drinks milk' won't help her talk about family memories




Midterms include an oral exam. She knows vocab, but speaking aloud—especially with corrections—terrifies her




She talks to her friend, who feels the same way. They wish for a low-pressure way to practice together.




Their professor mentions a new tool that matches students for real conversations—and gives feedback privately afterward.




Grace creates her profile. LinguaLeap pairs her with another learner at her level and suggests a time to meet



With structured prompts, the conversation feels natural—no awkward silence, no pressure



She gets pronunciation and vocabulary tips privately—no judgment, just support. For the first time, she feels like she's actually improving



Functional Requirements

Priority Level	Requirement	Why
High/Now	AI feedback accuracy	If feedback isn't accurate, users won't trust the AI or see improvement. A previously reported problem was users not being able to tell the difference between random recommendations and AI features so we want to make this more clear.
High	Conversation matching + scheduling	Without easy matching and scheduling, students can't practice consistently. This feature directly drives engagement and retention.
Medium	Instructor dashboard	The dashboard proves value to universities. It visualizes student progress, builds trust in the AI's results, and supports institutional adoption. It's less urgent short-term but will be important for more university partnerships.
Medium	"Voice calibration" feature	Calibration personalizes the AI to each user's voice and device, improving accuracy and user trust. Especially for when we expand to offering less common languages and want to improve our existing AI, we want our tech to be able to work across accents, unideal recording environments,

		etc.
Low	Corporate learning integrations	Corporate use cases came up during our brainstorming and would be a valuable vertical, but for right now, they only distract from the university focus and add complexity early on. We plan on working on these integrations once the student model is validated.
Low	Gamification features	Gamification may help long-term retention and engagement, but it can't fix weak feedback or poor matching which is why those are higher priorities.
Low	International market	Expanding internationally is another idea that came up during our brainstorming, but focusing on this expansion too early risks stretching our resources.

Non-Functional Requirements


Other than the core functionality described above, LinguaLeap must comply with certain non-functional requirements in order to maintain the quality of our product.

LinguaLeap's matching algorithm should be able to execute in a reasonable amount of time depending on the context (geographic region, number of viable matches, etc.). That can range from under <1 to 15 seconds depending on the complexity of matching. Our app's live transcription should also have very low latency (ideally between <1 to 3 seconds depending on utterance length) and post-conversation feedback should be completed within 10 to 30 seconds depending on conversation length. Furthermore, our performance should scale such that we are able to support tens of thousands of simultaneous conversations and other user interactions given that our current user base already comprises 25,000 active monthly users.

Our data security infrastructure should comply with industry standards (AES-256 encryption, robust authentication, etc.). Since we are using AI to provide feedback, we will need training data to improve our models, which should be anonymized and scrubbed of sensitive data. Furthermore, since we are working with student data, we should adhere to FERPA requirements. When the product expands internationally, we must have GDPR compliance (particularly in the EU) such that users can voluntarily consent to data collection and be able to port and delete data upon request.


Prototypes

We have two separate demos for the [landing page](#) and [the app](#) itself. In addition to the images below, we have prepared a separate [slide deck](#) that includes more screenshots of the app interface and functionality. For the sake of brevity, we have not included those images here in this document. Students can access key information and features in the app such as past conversations, feedback logs, session requests, active streaks, Elo score, and other data.




Landing Page

Student Profile

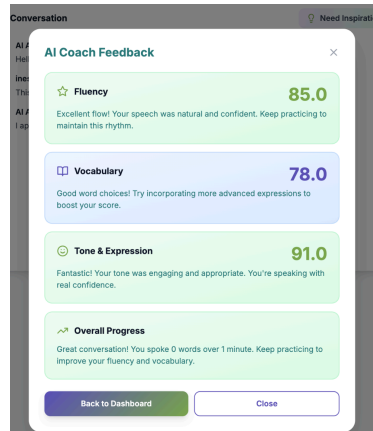


Partner and Individual Practice


Conversation Topics



AI Practice Interface



Post-Conversation Feedback



Schedule Next Conversation

Constraints and Assumptions

Key Constraints and Limitations

We have identified the following high-level constraints on LinguaLeap's product development and deployment:

- **Runway and growth pressure:** With \$350k monthly burn and about 5 months of runway, new features must show retention and revenue impact quickly.
- **AI capabilities:** Personalization is currently mostly rule-based, and beginner speech feedback has only been tested in narrow scenarios, constraining how complex and autonomous AI coaching can be.
- **Student willingness and engagement:** Current collaborative and peer features are rarely used and six-month individual retention is around 50% so we must validate desirability.
- **University budgets and competition:** Paid university contracts must compete with free options like Duolingo for Schools, making decision-makers highly sensitive to cost and perceived differentiation.
- **Student price sensitivity:** Student interviews suggest a realistic willingness to pay around \$10/month, which limits ARPU and increases reliance on retention and institutional deals.
- **Privacy and institutional requirements:** Handling student speech and progress data in university settings requires FERPA-aligned privacy and data controls, constraining how recordings, transcripts, and performance data can be stored and reused.

Core Assumptions and Test Results

Given these constraints, we selected three assumptions across desirability, viability and feasibility:

Desirability: Students value in person peer conversations and will meet regularly

Across roughly ten user tests and interviews, students repeatedly said that structured, low pressure in-person conversations with peers helped their speaking more than apps or flashcards, and sessions with light prompts and gentle feedback felt more productive and less awkward. This supports keeping guided in person peer conversations at the center of LinguaLeap; if this stops being true, the product loses its demand advantage over solo practice tools.

Viability: Students or universities will pay for rigorous AI conversational feedback

After trying an AI guided conversation prototype, students linked speaking confidence directly to internships and jobs, noted that they already pay for classes or tutoring, and said they would be comfortable paying around ten dollars per month for something that clearly improves speaking and tracks progress. This confirms willingness to pay for our value proposition, while indicating that sustainable pricing needs to stay near the low to mid end of our original range.

Feasibility: AI can accurately understand beginner speech and provide constructive feedback

In French speaking trials with beginner learners, the model almost always understood accented, hesitant speech and correctly inferred intent. It often distinguished unclear pronunciation, although it

sometimes over-corrected grammar. This shows that we can rely on AI for recognition and intent understanding and should design a more constrained grammar feedback layer; if the model doesn't reach this level, LinguaLeap becomes a matching tool without real learning value.

Milestones and Timeline

Timeframe	Milestone	Key Deliverables
Weeks 1–3	Staff a dedicated LinguaLeap squad	4 engineers, 2 designers, and 1 education specialist hired and onboarded; roles and ownership for the speaking and retention workstreams clearly defined.
Weeks 2–4	Lock AI integration plan and infrastructure upgrade scope	Selected speech and feedback models; signed-off integration design with the existing platform; prioritized infra backlog to support real-time sessions and logging at current scale.
Weeks 4–8	Ship first wave of retention features with mobile parity V1	New retention and guided speaking flows live to production; core journey usable on both desktop and mobile; tracking for activation and weekly usage.
Month 3	Run structured pilots with ~8 university partners	Pilots launched in up to eight courses or campuses; instructors onboarded with simple playbooks; each pilot has clear success metrics and control vs. test groups.
Month 4	Quantify impact on retention and ARR and iterate once	Report comparing pilot vs. baseline cohorts on engagement and retention proxies; early ARR lift (conversion or upsell) documented; at least one iteration to prompts/feedback shipped based on findings.
Month 5	Demonstrate traction and Series B readiness	Investor-ready story summarizing pilot outcomes, retention uplift, student and instructor satisfaction, and projected ARR if rolled out to all university partners; next-phase roadmap agreed with leadership.

Resource Requirements

This section outlines the staffing, infrastructure, and tools required to build, maintain, and scale LinguaLeap's real-time conversational learning platform.

Human Resources

Resource	Cost	Description
----------	------	-------------

Engineering	4 FTE \$720k–\$880k/yr	These engineers will develop and maintain the real-time speech pipeline, build and optimize the matching engine, expand the instructor dashboard, ensure mobile and desktop feature parity, and handle reliability improvements, load testing, and infrastructure scaling
Design	2 FTE \$300k–\$360k/yr	The design team will create the conversation flow UX, shape the instructor dashboard UI and visualizations, ensure accessibility across platforms, and iterate on prompt, transcription, and feedback surfaces.
AI/ML + Speech Specialist	0.5–1 FTE \$160k–\$220k/yr	This specialist will select and evaluate models, design the voice calibration pipeline, and refine grammar, pronunciation, and vocabulary feedback layers.
Education Specialist	1 FTE \$110k–\$150k/yr	This role will ensure that feedback, prompts, and learning flows are pedagogically grounded and aligned with university curricula.
Go-to-Market and Partnerships	1–1.5 FTE \$150k–\$220k/yr	This team member will manage university procurement cycles, structure and oversee pilots, support departmental adoption, and onboard instructors.
Customer Success and Support	0.5–1 FTE \$60k–\$120k/yr	This role will provide onboarding for instructors and students, support early pilots, troubleshoot issues, and coordinate with engineering on reliability concerns.

Total annual personnel cost: \$1.5M–\$1.95M.

Technical Resources

Resource	Cost	Description
AI/ASR Infrastructure	\$15k–\$40k/mo	This infrastructure will support real-time speech-to-text, model hosting and fine-tuning, and GPU-backed feedback processing.
Backend Infrastructure	\$8k–\$15k/mo	These systems will run conversation sessions, store sessions/Elo/progress data, and provide monitoring, analytics, and logging.
Frontend + Mobile Tooling	\$1k–\$3k/mo	This environment will support CI/CD pipelines, device testing, and release management.
Security + Compliance	\$20k–\$40k/yr	These tools will offer FERPA-aligned access controls, data anonymization capabilities, and regular security and vulnerability scanning.

Total infrastructure cost: \$24k–\$58k per month.

Additional Tools: Product analytics (Mixpanel/Amplitude), bug tracking (Linear/Jira), load-testing frameworks, and internal tools for instructor demos will support development and deployment.

Testing and Quality Assurance

LinguaLeap’s QA process ensures reliable real-time conversations, accurate AI feedback, and university-level security.

- **Functional Testing:** We will validate matching, scheduling, session flow, feedback delivery, and mobile–desktop parity, supported by regression tests and automated UI checks for onboarding, profile setup, and instructor dashboard use.
- **AI Feedback Validation:** Speech recognition will be tested on accented and beginner audio, and expert review will confirm accuracy of pronunciation, grammar, and vocabulary feedback while monitoring over-corrections. Calibration consistency and Elo stability, convergence, and fair partner matching will also be evaluated.
- **Performance & Load:** We will simulate 5k–50k concurrent users, verify latency targets (transcription <1–3s, feedback 10–30s, matching 1–15s), stress-test peak usage, and track crash and dropout rates to ensure reliability.
- **Security & Compliance:** FERPA-aligned data handling, encryption, role-based access control, vulnerability scans, and penetration tests will ensure secure operation.
- **Pilot Validation:** Pilots will measure speaking frequency, conversation completion, instructor satisfaction with dashboard accuracy, and student perceptions of clarity, fairness, and usefulness of AI feedback, supported by canary releases, rollback thresholds, weekly QA checks, and end-of-term audits.
- **Success Indicators:** Session failure <5%, mobile crash rate <1%, instructor acceptance of AI feedback ≥80%, and no unresolved high-severity security issues at release.

Metrics and Success Criteria

We focused on the parts of the product that matter for universities and for students such as consistent speaking practice, measurable progress, and reliable course-level adoption. The table below outlines the specific objectives we’re holding ourselves to and the indicators we’ll use to evaluate whether the product is moving in the right direction.

Objective	Key Results	Primary KPIs
Strengthen Retention by Making Speaking a Habit	• 6-month retention increases from 50% to 70% • Avg. speaking frequency increase from <1/mo → 2x month • 80% of active users complete 1+ matched conversation biweekly	• Weekly Active Speakers (WAS) • Time-to-first-conversation • Conversation completion rate • Elo match stability
Grow ARR Through University Partnerships	• University partners jump from 12 to 20 in 6 months • University ARR \$1.17M → \$3.5M • 90% of partnered universities use LinguaLeap in weekly / bi-weekly routines	• Signed/renewed contracts • Instructor adoption rate • % of courses integrating conversation sessions • Renewal likelihood (NPS + qualitative signals)

**Demonstrate
Measurable
Speaking
Improvement**

- Average Elo increase $\geq 12\%$ in 4 weeks
- 75% of students use AI feedback
- 60% of instructors reference Elo in assessments

- Elo progression curves
- AI feedback completion rate
- Instructor dashboard activity
- Speaking confidence (pre/post)

**Improve Mobile +
Platform Reliability**

- Mobile parity 60% to 90%
- Session failure rate $< 5\%$
- Concurrent capacity 5k to 50k users

- Feature parity index
- Crash/dropout rate
- Load-testing thresholds

These indicators give us a straightforward way to assess whether the pivot is working. If students are completing conversations regularly, their Elo scores show steady progression, and instructors incorporate the tool into their routines, that's a strong signal that the product is delivering real value. If those patterns don't show up, or they show up inconsistently, that's a clear sign that we need to adjust the product before scaling further.

Go To Market

Distribution Channels

We analyzed interviews, competitive analysis, and professor / student feedback and found that our strategy should be to become the infrastructure that their classes already depend on. Even though our primary users are students and all of our product features are tailored towards their needs, we prioritize universities for distribution and revenue rather than trying to "market" to students.

Primary Distribution Channel: University Partnerships (B2B)

Our main path is through universities, not only because of their "whale" efficiency in taking them on as clients, but because it matches how language learning actually happens. Students are already doing weekly speaking activities in class where TAs already run conversation sections, and professors already need a structured way to measure speaking progress. LinguaLeap fits naturally into all of those routines which we will distribute in three ways:

- **Department-wide or university-wide software contracts:** Lets entire programs adopt LinguaLeap as their default speaking practice tool = reducing fragmentation.
- **Course-level integration:** Professors embed sessions directly into weekly schedules.
- **TA-led and department-run conversation hours:** These events already exist on many campuses. By powering them with LinguaLeap, we make the experience more structured and less awkward, and we give departments a simple way to guarantee consistent practice .

How and why this works

Since universities already budget for academic technology and care about measurable outcomes (think of university contracts with LeetCode, CaseCoach etc), we know that the service benefits them as it improves not only the students' language ability, but the professor's access to student improvement and a continuation of the learning environment. Not only that, but the classroom provides built-in trust and oversight, which is especially important because our whole value proposition depends on real-life interaction and as we have seen with examples like Tandem.

Secondary Distribution Channel: Student Word-of-Mouth

We're intentionally avoiding ads and broad marketing because students don't adopt academic tools with "influencers" they don't trust academically. Instead, they adopt trusted tools that their professors and TAs are using. Thus, our word-of-mouth strategy is quiet but extremely effective with:

- **Peer-driven adoption inside each course:** Students talk about the tools that help them perform better (especially if they're anxious about speaking). Once a few students see that matched partners make conversations easier, the tool spreads quickly.
- **Micro-ambassadors inside classes:** These aren't official roles but rather they're the students professors already rely on to explain assignments or walk others through technical tools.
- **Department-backed events:** When a department sponsors a conversation hour using LinguaLeap, the tool gains instant credibility.

Our "marketing" is the social atmosphere of the classroom where if LinguaLeap becomes part of how students prepare for their oral work, we know that usage will spread naturally and sustainably. This sales approach also fits our runway constraints because universities already budget for academic software, and the pilot-to-contract cycle is shorter than corporate / K-12 procurement.

Sales Strategy

Our sales strategy focuses on programs that already run structured speaking components, since they feel the pain of coordinating conversation practice most directly. The sales motion starts with outreach to language department chairs, program directors, and centers for teaching and learning. From there, we offer a short pilot inside one course sequence so instructors can test Elo-matched conversations and the feedback tools with real students.

Contracts are priced at a standardized annual rate and aligned to academic calendars so departments can plan around existing instructional budgets. Our goal in each pilot is to show clear improvements in speaking frequency, Elo progression, and professor visibility into student progress.

Customer Support & Service

To support universities after adoption, we provide onboarding for professors, TAs, and students at the beginning of each term. For instructors, this includes a walkthrough of the dashboard, Elo interpretation, and how to integrate conversation sessions into their weekly routine. Students receive a short onboarding module that helps them set up their profiles and schedule their first conversation.

Proceed, Pivot, or Perish Criteria

Proceed

- Students complete conversations regularly (ideally biweekly to begin with)
- Elo scores show consistent progression across courses
- Instructors integrate LinguaLeap into their weekly routines
- University partners expand usage or renew without hesitation

Pivot

- Departments adopt the tool but students don't consistently participate
- Elo data shows weak or uneven progression across levels
- TA-led or department-run conversation hours have low turnout

Perish or fully
shift strategy

- Reliability or mobile parity issues continue to create friction
- Speaking frequency remains close to baseline despite classroom integration
- Retention does not improve beyond current levels
- Universities decline renewals or pilots do not convert
- It's technically unreliable (chat and AI feedback don't meet our standards)

References

"Cabrillo College Success Story." *Mheducation.com*, 15 Aug. 2024, www.mheducation.com/highered/blog/2024/08/cabrillo-college-success-story.html. Accessed 9 Oct. 2025.

Catalano, Frank. "Surveys Find Districts Are Using More Edtech Tools — and Teachers Are Bearing the Costs." *EdSurge*, 11 Jan. 2021, www.edsurge.com/news/2021-01-11-surveys-find-districts-are-using-more-edtech-tools-and-teachers-are-bearing-the-costs.

"Cengage Unlimited for Institutions Expands to Offer Support for Online Instruction | Cengage Group." *Cengagegroup.com*, 2020, www.cengagegroup.com/news/press-releases/2020/cengage-expands-institutional-offerings-to-support-colleges-and-universities-moving-to-online-learning/. Accessed 9 Oct. 2025.

"Corporate Online Language Learning Market Size, Share, 2033." *Businessresearchinsights.com*, 2025, www.businessresearchinsights.com/market-reports/corporate-online-language-learning-market-121007.

Duolingo. "Duolingo Works." *Duolingo*, www.duolingo.com/efficacy.

Global Market Insights. "Language Learning Market Size, Share & Forecast, 2023." Global Market Insights, 2023, www.gminsights.com/industry-analysis/language-learning-market

Goodwin, Emily. "The Word on Language Learning Apps – What's behind Their Growing Popularity?" *CivicScience*, 30 Jan. 2023, civicscience.com/the-word-on-language-learning-apps-whats-behind-their-growing-popularity/.

OpenAI. "Introducing Whisper." *Openai.com*, 2023, openai.com/research/whisper.

Rocheft, Simone De. "Duolingo Users Are in Turmoil over the App's AI Lessons." *Polygon*, 4 June 2025, www.polygon.com/ai-artificial-intelligence/603216/duolingo-ai-language-lessons/.

Statista "Market Size of the Global Language Services Industry from 2009 to 2022 with a Forecast until 2026." Statista, 2025, www.statista.com/statistics/257656/size-of-the-global-language-services-market

Statista "Leading Language Learning Apps Worldwide in July 2024, by Downloads." *Statista*, 2024, www.statista.com/statistics/1239522/most-popular-language-learning-apps-worldwide-by-downloads

Statista Consumer Insights. Target Audience Report: Online Language Course Users in the United States. Statista, Mar. 2025, www.statista.com/study/137188/target-audience-language-course-users