CORSEN AI Academy

Voxtral Audio: Transcription and Speech-to-Text

Course syllabus

Version 2026-11

Overview

Offline and real-time audio transcription with Voxtral: diarization, context biasing, microphone streaming, and latency optimization.

Learning objectives

- Transcribe an audio file with the offline API retrieving timestamps and temporal segmentation
- Enable diarization to attribute each transcribed segment to the correct speaker
- Improve accuracy on business vocabulary thanks to context biasing
- Implement real-time transcription in Python fed by the microphone via PyAudio
- Choose between offline and real-time transcription based on latency and volume to process

Audience and prerequisites

Developers, engineers, and integrators looking to leverage Mistral AI via API in their applications.

Ability to read Python and JSON, command line basics. A Mistral AI Studio API key to try the examples.

Duration and format

Estimated duration : 1 h 45 min. 100% online self-paced course, accessible 24/7 from a browser. Progress saved, resume anytime. Estimated duration: 1 h 45 min.

Content

- Audio Overview (3 lessons)
- Offline Transcription (5 lessons)
- Real-Time Transcription (5 lessons)
- Conclusion (3 lessons)

Assessment

Application exercises provided throughout the lessons. Each lesson ends with its Key takeaways to support self-assessment. A final quiz of 5 corrected questions at the end of the last lesson validates skills.

Accessibility

Self-paced course accessible to people with disabilities: text content compatible with screen readers, keyboard navigation, and compliant contrasts. For any adaptation needs, email hello@corsen.ai.

Price and recognition

Free

Certificate of completion issued by Corsen AI.

Access time

Immediate after account creation and email confirmation.

Corsen AI — Paris, France — SIRET 10297040700015
academy@corsen.ai — https://academy.corsen.ai