

NOTICE

AI in Hiring — Transparency Notice

What Gradii's AI does in the hiring process, what it does not do, what data it uses, and how a person stays in control of the decision.

VERSION	EFFECTIVE	OWNER	REVIEW	CLASSIFICATION
1.0	9 October 2026	AI Lead	Every six months, and on any material change to scoring	Public

Who this is for

This notice is for **candidates** who are interviewed, called or verified through Gradii, and for the **employers** who use Gradii and must explain their hiring process to those candidates. Employers may share this notice with candidates or link to it.

The short version

- Gradii uses AI to read resumes, ask interview questions, hold screening calls, score answers and summarise interviews.
- Scores are **recommendations for the hiring team**. Gradii can also move a candidate to the next stage, or out of the process, when a score crosses a threshold that the employer has set. The hiring team can see every one of those changes and reverse it.
- Gradii does **not** ask for protected characteristics such as gender, religion, caste, age, disability or ethnicity, and does not score them. It does not infer emotions, personality or honesty from your face or voice.
- Gradii does **not** use your data to train AI models.

What the AI does, step by step

Stage	What the AI does	Data it uses	Who makes the decision
Resume parsing	Extracts skills, experience and education from your resume	Your resume	Not a decision; produces structured data
Talent-fit score	Compares the extracted profile with the job's requirements and produces a percentage	Your structured profile; the job description	The employer sets the threshold; see "Automatic stage changes"

Stage	What the AI does	Data it uses	Who makes the decision
Interview questions	Generates questions for the role, or uses questions the employer wrote	Job description; employer's question bank	—
AI interview	Presents questions by voice, video or text; records your answers	Your answers: audio, video, text and code	—
Iris screening call	An AI voice agent calls you, asks screening questions in your chosen language and records the call	Call audio and transcript	See "Automatic stage changes"
Scoring and summary	Scores each answer against the question's criteria and writes a summary for the hiring team	Transcript of your answers; the question's evaluation criteria	Hiring team reviews
Interview integrity	Checks that the same person is present throughout, and records events such as leaving the interview tab or pasting text	Reference photo and live camera frames, as face-match scores; browser events	Flags only; a person reviews
Background verification	Matches names and details across documents you share, and screens against public sanctions and adverse-media sources	Documents and identifiers you consent to share	The employer reviews the report

Automatic stage changes

Some employers let Gradii move candidates automatically, to keep a process moving:

- **Pass:** at or above the pass mark, you move to the next round.
- **Under review:** between the two thresholds, a person at the employer decides.
- **Automatic rejection:** at or below the employer's automatic-reject threshold, you are moved out of the process. For Iris screening calls, the threshold is currently 70%.

Every automatic change is recorded with its score and reason, and the hiring team can reverse it. If you think an automated outcome was wrong, you can ask the employer for a human review. If the employer is unreachable, write to team@gradii.ai and we will pass your request on.

How we try to keep scoring fair

- **Same questions, same criteria.** Candidates for the same role are scored against the same criteria and rubric.
- **Content, not presentation.** Scoring is designed to assess what you said, wrote or coded against the question's criteria. It is not designed to assess your appearance, accent or background.
- **No protected inputs.** Gradii does not ask candidates for protected characteristics, and the scoring rubric contains none.
- **Languages.** Iris calls are available in 11 languages. Speech is transcribed in the language you speak.
- **Explainability.** Every score is stored with the question, your answer and the model's reasoning, so that a person can check it.

- **Monitoring.** Regular bias testing of scores across candidate groups is **planned**. It is not yet in place. Until it is, we recommend that employers review automatic outcomes, especially rejections.

Known limitations

- Speech-to-text can make mistakes, especially with background noise, strong accents or mixed languages. A transcription error can affect a score.
- Language models can misjudge an unusual but correct answer. That is why scores come with reasoning, and why a person should review borderline and negative outcomes.
- Face matching can fail in poor lighting or with low-quality cameras. A failed match produces a flag for a person to review. It never leads to rejection on its own.

Your rights

Depending on where you live, you may have the right to know what data was used, to correct it, to have it erased, to withdraw consent, to ask for a human to review an automated decision, and to contest it. Contact the employer that invited you first. You can also contact Gradii's Grievance Officer at team@gradii.ai.

Models and providers

Gradii uses large language models from Google (Gemini) and OpenAI, and speech models from Deepgram, ElevenLabs and Sarvam AI. All of them are accessed through their commercial APIs. The full list, with locations, is in our [Sub-processor List](#).

Manjari Technologies Private Limited ("Gradii") · 316, Neelkanth Corporate Park, Vidyavihar West, Mumbai, Maharashtra, India · gradii.ai/trust-vault

The current version of this document is published at gradii.ai/trust-vault. Printed or downloaded copies are uncontrolled.