

System Card

Firstview • Version 1.0 • Updated: 21 October 2026

This card describes the AI systems in Firstview: what they are for, how they work, how we test them, and their known limits. It is written for customers, their legal and procurement teams, and candidates who want more detail than our AI notice gives.

1. Intended use

Firstview helps employers and recruitment agencies run hiring. The AI gives the hiring team information about candidates. **People make every hiring decision.** No part of the system rejects, ranks out or moves a candidate on its own.

The system is not intended for: deciding on promotion, termination or pay; assessing people who have not applied or been sourced for a role; any use that infers emotions or personality; inferring protected characteristics such as age or gender. The system does not infer them by itself; the one case where it may, when a customer writes an age, gender or military-service criterion for CV screening, is described in section 6.

Under the EU AI Act, AI used to evaluate job candidates is a high-risk system (Annex III, point 4). Firstview is the provider; the customer is the deployer. The high-risk obligations apply from 2 December 2027.

2. Components

Component	Input	Output	How it works
CV screening	CV text, the job's criteria written by the customer	For each criterion: a result (for example met, not met, unclear) and an explanation	A decision model answers each criterion from the CV, and a language model writes the explanations. If the decision model is unavailable, a language model agent checks the criteria instead, using tools for structured steps such as calculating years of experience.
AI interviewer	The customer's interview plan and questions, the candidate's speech	Spoken questions and follow-ups	Speech-to-text turns the candidate's speech into text; a language model decides the next question or follow-up; text-to-speech speaks it.
Interview report	Interview transcript, the competencies the customer chose	A summary, an assessment per competency with transcript quotes, and highlights for the team	A language model reads the full transcript after the interview.
Notetaker	Recording of a meeting the customer records	Transcript, summary and notes	A meeting bot records the meeting; speech-to-text and a language model produce the notes.

Component	Input	Output	How it works
Candidate search	A job description or search query from the customer	A list of matching profiles	Text embeddings and a ranking model match the query with profiles in the customer's candidates and our sourcing index.
Assistant	A question from a customer user	An answer	A language model answers using only data the user can already see. Turned on per customer.

3. Models and providers

The language models come from Anthropic, OpenAI and Google, and the CV screening decision model (Jev) from TypeSafe, through the OpenRouter gateway and Google Cloud Vertex AI. Speech-to-text and text-to-speech come from Soniox, with Cartesia and Deepgram as backups. Embeddings and ranking for search come from Cohere. We pick models per component by running our evaluations (section 5) and change them as better models appear. The sub-processor list at <https://firstview.ai/trust#subprocessors> names every provider and its region.

None of these providers may train models on our data. Firstview does not train or fine-tune models on customer or candidate data.

4. Design choices that limit risk

- **People decide.** Results are information for the hiring team. Every result comes with reasons and evidence (CV text or transcript quotes), and the recording and transcript stay available, so a reviewer can check the AI.
- **The customer sets the criteria.** CV screening and interview reports assess only against criteria and competencies the customer writes or chooses for the job.
- **No face, emotion or personality analysis.** The system does not analyze faces, voice tone, emotions or personality. The camera check during interviews measures only image brightness.
- **Protected topics.** The AI interviewer is instructed not to ask follow-up questions about protected characteristics (age, family, pregnancy, health, disability, religion, ethnicity, politics, union membership, sexual orientation, gender identity, military-service reasons), and not to probe when a candidate mentions one. Reports describe what the candidate said and did, not their emotional state.
- **Transparency.** Candidates are told that an AI conducts the interview. Proctoring events (tab changes, full-screen exits and similar) are shown to reviewers next to the interview and do not change the AI's assessment.
- **Human review.** Candidates can ask for a person to review any AI result (security@firstview.ai or the employer).

5. How we test

- **Interview behaviour.** Changes to the interviewer (prompts, models, flow) run against a set of scripted interview scenarios. A separate model judges each conversation against written criteria, such as staying on the plan, asking relevant follow-ups and not asking about protected topics. We compare results with the previous version before release.
- **Interview reports.** When we change the report model, we run candidate models on the same transcripts and compare their output before switching.
- **CV screening.** We test against labelled CVs and criteria, including production cases that went wrong, and check that explanations do not state facts the CV does not support.

Not yet done: we have not run a formal bias audit across demographic groups, and we do not yet publish accuracy figures. Both are part of our preparation for the EU AI Act.

6. Known limits and risks

- **Speech recognition errors.** Background noise, a weak connection or a strong accent can cause mistranscription, which can affect the report. Reviewers can check the recording.
- **Model errors.** Language models can misread an answer or a CV, miss relevant evidence or overstate it. Results are information, not decisions, for this reason.
- **Language coverage.** Interviews run in Turkish, English, German, Spanish and French. Quality can differ between languages.
- **Customer-written content.** The questions and criteria the customer writes are used as written. If a customer writes an unlawful or unfair question or criterion, the system will apply it.
- **Age, gender and military-service criteria in CV screening.** If a customer writes a criterion about age, gender or military service, CV screening evaluates it like any other criterion, and may estimate age from dates in the CV and gender from explicit statements or the first name. A military-service criterion can also lead it to estimate gender. These estimates can be wrong. Using age or gender in hiring is unlawful in most countries unless the job genuinely requires it. The customer chooses the criteria and is responsible for their lawfulness.
- **Proctoring events are signals, not proof.** A tab change can have an innocent reason. Reviewers should ask the candidate before drawing conclusions.

7. Human oversight

The customer's hiring team reviews every result before acting on it. Customers should make sure reviewers understand that results can be wrong, check the evidence, and do not rely only on the AI. Firstview staff do not review individual candidates except when the customer asks for support. When a candidate asks us for a human review, we pass the request to the customer, whose team makes the decision.

8. Contact

Questions or concerns about the AI: security@firstview.ai