

Annex IV Technical Documentation Skeleton — Anonymised Sample

Annex IV Technical Documentation Skeleton — Anonymised Sample

Anonymised — Do not redistribute. Excerpt from a Documentation Sprint deliverable. The customer "Acme Decisions Ltd" and the AI system "AcmeMatch" are illustrative.

Customer: Acme Decisions Ltd · **AI System:** AcmeMatch (AI-assisted candidate-screening for SaaS recruiters) · **Risk classification:** High-risk per EU AI Act Annex III §4(a) (AI systems used in employment for recruitment, in particular for screening or filtering applications)

Purpose of this document

This skeleton is the structured outline a Provider must populate to satisfy **EU AI Act Article 11** (Technical Documentation) for high-risk AI systems. It maps directly to **Annex IV** of the AI Act. The Provider keeps the populated version and provides extracts to notified bodies, market surveillance authorities, downstream Deployers, and procurement reviewers as required.

This sample shows **sections 1-3 of 9**. The full Documentation Sprint deliverable populates all 9 sections to first-draft level for one AI system.

1. General description of the AI system

1.1 Intended purpose

- AcmeMatch ranks candidate applications for SaaS recruiter customers against a customer-defined role profile, returning a ranked shortlist with structured rationale tags (e.g. "7+ years SaaS sales", "managed quota >£500k").
- Used by recruiter end users in the Acme Decisions Ltd customer base; final hiring decisions are made by the recruiter and the hiring manager, not by AcmeMatch.

1.2 Provider, authorised representative, and Deployer relationships

- **Provider:** Acme Decisions Ltd (UK).
- **Authorised representative in the EU:** [TBD — required before placing AcmeMatch on the EU market — see §9 *Outstanding actions*].
- **Deployers:** SaaS recruiting agencies and in-house talent teams, primarily UK + EU + US.

1.3 Versions covered by this documentation

- AcmeMatch v3.4.x — current production line. Documentation is updated on every minor version that materially changes scope, training data, or human oversight controls.

1.4 Software ecosystem

- AcmeMatch ingests CV PDFs and structured ATS records, processes via the AcmeMatch API, returns ranked output to the Acme Decisions web UI and via API to downstream ATS integrations.

1.5 Hardware on which AcmeMatch is intended to run

- Hosted SaaS only. Provider-controlled cloud infrastructure (UK + EU regions). No on-prem or edge deployment offered in v3.4.

1.6 Form of placing on the market

- SaaS subscription. No standalone software is shipped to the Deployer.

1.7 Description of user interface

- Web UI for recruiters: shortlist view with a per-candidate rationale panel, an "override and re-rank" control, an export-to-ATS button, and an "explain this rank" link to a per-candidate model-card view.

1.8 Instructions for use (Deployer-facing)

- Linked separately as the *AcmeMatch Deployer Use Manual*. Covers human oversight obligations under AI Act Article 26, retention of human-in-the-loop, and prohibited use cases (e.g. "*Do not use AcmeMatch as the sole basis for adverse hiring outcomes*").

2. Detailed description of the elements of the AI system and the development process

2.1 Methods and steps performed for the development

- Foundation: pretrained large language model accessed via licensed API; **no continued pre-training** by Acme Decisions Ltd.
- Fine-tuning: instruction-tuned on a curated dataset of (anonymised) prior recruiter shortlists with role-fit labels, plus a small synthetic-augmentation set covering rare role types.
- Post-processing: rule-based filters for **AI Act Article 5(1)(c) social scoring** patterns (rejected) and **Article 5(1)(f) emotion-recognition** patterns (rejected at input).

2.2 Design specifications

- Goal: rank order candidates by predicted role-fit, *not* to make hire/no-hire decisions.
- **Out of scope by design:** automated rejection emails, automated interview invitations without recruiter confirmation, demographic feature use as a ranking signal, voice/video assessment.

2.3 Description of the system architecture

- See architecture diagram referenced as `acmematch-architecture-2026-Q1.png` in the populated full Annex IV. Inputs → ingestion service → vectorisation → ranking model → rationale generator → human-oversight UI → ATS integration. Each component has a separate log stream with tamper-evident retention per Article 12.

2.4 Where relevant, the data requirements in terms of datasheets

- Training and evaluation datasets, their provenance, scope, main characteristics, the methods to obtain and select data, labelling procedures (including for sensitive attributes), and data cleaning methodologies — populated in §2.4.1 to §2.4.6 in the full deliverable.
- Sensitive-attribute handling: explicit demographic features (e.g. name, age, gender) are stripped before ranking. Implicit-bias auditing is run quarterly on the evaluation set.

2.5 Assessment of human oversight measures (cross-reference)

- Detailed in **§4 Human oversight measures** of this Annex IV (and in `04-human-oversight-transparency-checklist-sample.md`).

2.6 Pre-determined changes to the AI system and its performance

- The Provider may update model weights, retrain on newer recruiter feedback, or change rationale prompts. Material changes are notified to Deployers 30 days before deployment and re-tested against the original conformity assessment scope.

3. Detailed information about the monitoring, functioning, and control of the AI system

3.1 Capabilities and limitations in performance

- **In scope:** ranking against a recruiter-defined role profile in English, French, German, Spanish.
- **Out of scope:** recommending hire/no-hire; predicting protected-characteristic outcomes; assessing candidates from incomplete CVs.

3.2 Expected accuracy, including its metrics

- Top-10 recall vs. recruiter ground-truth shortlists: **target ≥ 0.78** , measured monthly on the held-out evaluation set.
- Calibration: Brier score on rationale-tag confidence: **target ≤ 0.18** .
- Drift detection: Population Stability Index (PSI) on input embedding distributions per role family, alert at $PSI > 0.15$.

3.3 Foreseeable unintended outcomes and sources of risk

- Indirect proxy bias from CV-language correlations with demographic groups. **Mitigated by** explicit demographic-feature stripping + quarterly bias audit (see [03-risk-management-register-sample.md](#), R-04).
- Over-reliance by the recruiter ("automation bias"). **Mitigated by** mandatory human-in-the-loop interaction before any candidate is moved to an interview stage.
- Data poisoning via crafted CV uploads. **Mitigated by** input sanitisation, rate limiting per recruiter, and anomaly detection on output distributions.

3.4 Human oversight measures, including the technical measures put in place

- See [04-human-oversight-transparency-checklist-sample.md](#). Summary: every shortlist returned to a recruiter has explicit "override and re-rank" UI, every candidate-level rank carries a "explain" link, and the recruiter must affirmatively select candidates before they are moved to interview.

3.5 Specifications of input data, as appropriate

- CV PDFs ≤ 5 MB, structured ATS records (JSON), role profile (structured).
 - Personal data minimisation: only fields strictly needed for role-fit ranking are processed; sensitive demographic fields are explicitly excluded at ingestion.
-

4-9. (Headings only in this sample; populated in the full deliverable)

- §4 Human oversight measures (cross-references the human oversight checklist and the Deployer Use Manual)
 - §5 Description of the changes to the system through its lifecycle
 - §6 List of harmonised standards applied (or where not, what alternative the Provider used)
 - §7 EU declaration of conformity (signed copy or reference)
 - §8 Description of the system in place to evaluate AI system performance in the post-market phase
 - §9 Outstanding actions and known gaps as of the documentation date — *e.g. EU authorised representative appointment outstanding; bias audit Q3 due*
-

Document control

Item	Detail
Sample produced	2026-04-25
Anonymised customer	Acme Decisions Ltd (fictional UK SME)
AI system in scope	AcmeMatch v3.4.x
Risk classification	High-risk (Annex III §4(a))
Reviewer	Customer's solicitor or DPO before any external reliance

Prepared with support from GeraCompliance. This document is not legal advice and does not guarantee compliance. The customer remains responsible for legal review, implementation, and regulatory obligations.