

Dataset & Feature Risk Assessment

Examines data provenance, labels, representativeness, missingness, proxy risks, accessibility, and feature controls for EMP-AI-01 ResumeRank.

SIMULATED OPERATING DOCUMENT

Prepared for a fictional employer using synthetic systems and synthetic applicant data. This document is a governance demonstration, not legal advice.

1. Data governance questions

- What population produced the training, validation, and operational data, and does it resemble the deployment population?
- Are labels based on job performance, manager preference, historical hiring, or another construct?
- Which fields may act as proxies for protected characteristics, disability, socioeconomic access, or prior structural disadvantage?
- Can the feature be tied to an essential or documented job-relevant qualification?
- Does missingness differ by group or candidate pathway?
- Can a candidate correct inaccurate parsed information?
- Are protected-class data used only in a segregated audit environment where lawful and appropriate?
- What event makes the dataset or feature approval expire?

2. Dataset register and provenance

Dataset	Source / owner	Purpose	Key risk	Control / validity period
DS-01 Job profile & competency map	HR job analysis	Defines approved scoring criteria	Outdated or inflated requirements	Annual job-family review; owner attestation
DS-02 Candidate resumes	Applicant / ATS	Operational ranking	Formatting, multilingual, missingness, name/address proxies	Parser QA; excluded fields; correction route; 12-month model-use approval
DS-03 Application answers	Applicant	Minimum qualifications	Inaccessible questions; unnecessary sensitive data	Accessible forms; legal/privacy review; version control
DS-04 Historical recruiter advances	ATS	Vendor proposed tuning	Prior human bias and inconsistent standards	Not approved as primary label
DS-05 Post-hire performance	HRIS	Potential criterion validation	Manager-rating bias; survivor bias; role mismatch	Validation study only after bias/quality review
DS-06 Demographic audit file	Voluntary / legally collected HR data	Adverse-impact monitoring	Sensitive data governance	Separate access; aggregate reporting; no operational scoring; retention per legal schedule

3. Feature decision matrix

Feature	Decision	Primary risk / rationale	Required control
Required certification	Permit	Direct job requirement where documented	Validate requirement; do not overgeneralize across jobs
Job-relevant skill evidence	Permit	Direct evidence	Use skill taxonomy and evidence tags
Relevant experience duration	Permit with cap	Potential age/career-path proxy	Use broad bands; no graduation-year derivation
Current job title	Restrict	Title prestige and industry-path proxy	Normalize to skill/level; low weight; ablation test

Feature	Decision	Primary risk / rationale	Required control
Employer / school prestige	Prohibit for scoring	Socioeconomic and demographic proxy risk	Strip entity prestige signals
Employment gap length	Prohibit for scoring	Caregiving/disability/nontraditional-path risk	Do not penalize gaps
Number of employers / job changes	Restrict	Age/career-pattern proxy	Only if job analysis establishes relevance; default exclude
Graduation year	Prohibit	Strong age proxy	Remove before model input
Name / honorifics / pronouns	Prohibit	Protected-trait proxy	Mask before scoring
Street address / ZIP	Prohibit	Race/socioeconomic proxy; usually unnecessary	Remove before scoring
Photo / video	Prohibit for résumé ranking	Not necessary; biometric/protected-trait exposure	Never ingest
Disability or accommodation data	Prohibit for scoring	Sensitive and legally protected	Separate accommodation workflow
Citizenship / nationality inference	Prohibit	National-origin risk	Use only lawful work-authorization question outside model
Language fluency	Restrict	National-origin impact	Only when job-relevant; directly assess if necessary
Writing quality from résumé	Restrict	Language/disability formatting effects	Use only where writing is essential and validate direct assessment
Volunteer / unpaid experience	Permit	Can capture transferable skill	Count equivalent evidence; avoid prestige weighting
Inferred personality / emotion	Prohibit	Weak job linkage / disability and protected-class concerns	No deployment
Keyword density	Restrict	Resume-gaming and formatting bias	Use semantic evidence + human review; monitor parser errors

4. Label validity assessment

The vendor proposed tuning the model to historical recruiter advances and hires. That label is rejected as the primary ground truth because it measures past selection behavior, not job performance. A future criterion-related validation study must identify a defensible outcome tied to the actual role, examine the quality and bias of that outcome measure, and document whether the model adds useful prediction beyond simpler methods.

DATA DECISION

Do not use raw historical hire/advance decisions as a success label for production ranking. Use job-analysis-grounded features and, if criterion validation is pursued, evaluate post-hire outcome quality before treating it as ground truth.

5. Data quality and representativeness controls

Control	Method	Evidence
Parser error testing	Sample résumés across PDF/DOCX/plain text, multi-column layouts, assistive-tech outputs, nontraditional histories	Error-rate report; failure examples; remediation
Missingness review	Compare missing fields and parser failures across process paths and audit groups where lawful	Missingness table + significance/practical review
Population fit	Compare validation dataset with job family, geography, level, and applicant source	Representativeness memo
Data lineage	Track source, transformations, derived features, model version, and retention	Data lineage diagram / data dictionary
Sensitive audit data segregation	Separate demographic audit file from operational model features	Access-control evidence + query logs
Expiry / refresh	Invalidate approval after material role or model change or 12 months	Dataset approval record

6. Assessment conclusion and release evidence

Dataset approval is conditional. Resume parsing and job-profile data may support a limited pilot only after the risky features above are technically excluded, parser accessibility is tested, historical recruiter decisions are barred as the primary success label, and the audit dataset remains segregated from operational scoring.

Release evidence	Owner	Gate
Approved job analysis and feature dictionary	HR + Validation	Required before validation
Parser QA across accessible and nontraditional résumé formats	Product + Accessibility	Required before pilot
Feature-exclusion and lineage evidence for prohibited fields	Vendor + Data / Analytics	Required before pilot
Protected-class audit-data access and retention controls	Privacy + Legal	Required before subgroup testing
Dataset/model expiry date and material-change trigger	AI Governance	Required before approval