

Algorithmic Impact Assessment / Employment AI Assessment

Pre-deployment assessment for EMP-AI-01 ResumeRank, a high-impact résumé-ranking tool used to recommend applicants for recruiter review.

SIMULATED OPERATING DOCUMENT

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

SHOWCASE ARTIFACT

Decision: baseline configuration NOT APPROVED. A mitigated pilot is conditionally approvable only after required validation and controls are evidenced.

1. System and decision context

Field	Assessment
System	EMP-AI-01 ResumeRank (fictional third-party tool)
Purpose	Prioritize applicants for recruiter review for Customer Operations Specialist roles.
Business justification	Reduce manual sorting time while maintaining a consistent job-relevant first review. Efficiency alone does not justify criteria unrelated to job performance.
Decision point	After minimum eligibility checks and before recruiter phone screen.
Population	External applicants in the United States; NYC applicants may trigger additional AEDT controls.
Output	0-100 score, ranked list, top-signal explanation tags.
Human decision role	Recruiter must review screen-outs and may override. No auto-reject from AI score alone.
Materiality	High: ranking determines which candidates receive scarce recruiter attention.

Authority notes: R1, R2, R3, R6.

2. Process map and decision rights

Step	System action	Human action	Control
1. Apply	ATS collects résumé and structured application answers	Candidate submits; accommodation path available	Notice + data minimization
2. Eligibility	Deterministic rules check required work authorization / certification where lawful	Recruiter handles exceptions	AI does not infer eligibility
3. Ranking	ResumeRank generates score and evidence tags	No decision yet	Protected/sensitive attributes excluded from model input
4. Review queue	System orders candidates	Recruiter reviews top and screen-out queue using structured rubric	No auto-reject; controlled reason codes

Step	System action	Human action	Control
5. Advance / reconsider	System records workflow event	Recruiter advances, overrides, or requests second review	Override logged; low-score screen-outs sampled / reviewed
6. Candidate recourse	System stores audit trail	Candidate may request accommodation, alternative process, or human reconsideration	SLA + documented resolution
7. Monitoring	Analytics calculates subgroup outcomes	Governance team reviews thresholds and complaints	Pause / rollback triggers

3. Stakeholders and affected groups

Stakeholder	Potential benefit	Potential harm / concern
Applicants	Faster review; more consistent criteria	Unfair screen-out, inaccessible process, opaque reason, inability to challenge
Applicants with disabilities	Alternative formats can improve access	Résumé gaps, assistive-technology formatting, timed steps, or proxy traits may depress scores
Protected demographic groups	Standardized review can reduce some subjective variation	Historical data or proxy features may reproduce inequity
Recruiters	Prioritized queue and reduced administrative load	Automation bias; excessive trust; inadequate time for real review
Hiring managers	More consistent candidate pipeline	False confidence in fit scores; narrowed candidate diversity
HR / Legal / Governance	Traceable decision process	Insufficient vendor evidence; difficult incident investigation
Vendor	Clear requirements and feedback	Contract burden; limits on model changes / data reuse

4. Data sources and feature posture

Data category	Source	Use	Risk posture
Résumé text	Candidate	Experience, skills, certifications	Allowed with feature restrictions and parser testing
Structured application answers	Candidate	Job-specific eligibility / experience	Allowed if validated and accessible
Job description / competency profile	Employer	Reference criteria	Allowed; must be current and job-relevant
Historical hiring decisions	Employer	Vendor proposed tuning label	RESTRICTED - may encode past discrimination and inconsistent manager preferences
Post-hire performance data	Employer	Potential validation criterion	Restricted; must assess measurement bias and role comparability
Names, photos, addresses / ZIP	Candidate	Not needed for ranking	Excluded from scoring
Disability / accommodation data	Candidate / HR	Accommodation administration only	Segregated; never used for ranking
Inferred personality / emotion	Derived	Not necessary	Prohibited

5. Risk assessment

Risk ID	Risk scenario	How harm could occur	Inherent	Required controls	Residual
AIA-01	Historical label bias	Model tuning to past advances/hiring decisions may reproduce prior manager preferences.	High	Do not train/tune to raw historical hire decisions; prefer job-analysis-grounded criteria; test alternatives.	Medium
AIA-02	Career continuity proxy	Employment-gap and tenure-continuity signals may disadvantage caregivers, people with disabilities, or candidates with nontraditional paths.	High	Remove gap-length and continuity-derived features; validate remaining experience measures.	Low-Med
AIA-03	Title prestige proxy	Current title / employer prestige may proxy socioeconomic access rather than job skill.	High	Map to skill evidence; cap weight; test feature ablation.	Medium
AIA-04	Résumé parser accessibility	Nonstandard formatting or assistive-tech-generated files may parse poorly.	High	Parser QA; plain-text alternative; candidate correction route; no penalty for format artifacts.	Low-Med
AIA-05	Disability screen-out	Tool may score a candidate poorly for a disability-related reason unrelated to essential job functions.	High	Accommodation notice; alternative process; direct-measurement preference; human reassessment.	Medium
AIA-06	Automation bias	Recruiters may treat score as authoritative.	High	Hide rank percentile until rubric review where feasible; reviewer training; override metrics; second review for screen-outs.	Medium
AIA-07	Adverse impact	Baseline synthetic test shows Black selection-rate ratio of 0.75 vs highest-rate group.	High	Pause baseline; remove risky features; re-test; evaluate job relevance and less-discriminatory alternative.	Low-Med if mitigated
AIA-08	Vendor update drift	Vendor may update model weights without client validation.	High	Contractual advance notice; version lock; change-triggered reapproval.	Low
AIA-09	NYC compliance	Covered AEDT use may lack timely bias audit/public summary/notice.	High	Jurisdiction gate; independent audit workflow; 10-business-day notice process.	Low
AIA-10	Privacy / secondary use	Candidate data may be retained or reused beyond hiring purpose.	High	Purpose limitation, no vendor model training without approval, deletion SLA, subprocessor controls.	Low
AIA-11	Recourse failure	Candidate has no meaningful way to challenge an AI-influenced outcome.	High	Human reconsideration, accommodation and alternative-process route; documented SLA.	Low
AIA-12	Monitoring blind spot	One-time prelaunch audit misses drift and operational behavior.	High	Monthly early-life monitoring; subgroup, override, complaint, accessibility, and version metrics.	Low

Authority notes: R2, R3, R5, R6, R7.

6. Alternatives assessed

Alternative	Discrimination / rights impact	Operational impact	Assessment
Manual unstructured résumé review	May preserve or increase subjective bias; weak auditability	High recruiter effort	Not preferred without structured rubric

Alternative	Discrimination / rights impact	Operational impact	Assessment
Structured rules based only on minimum qualifications	More explainable; lower proxy risk	May miss transferable skills	Viable for eligibility, not full prioritization
Mitigated ranking model excluding continuity, graduation-year, ZIP/address, name/entity prestige, and inferred traits	Lower observed disparity in synthetic test; more defensible feature set	Similar prioritization benefit	Preferred pilot candidate
Skills work sample before ranking	Directly measures some job requirements; accessibility must be designed	More candidate effort	Use for selected job families after validation
No AI ranking; structured human scoring rubric	High transparency; human variability remains	Higher recruiter effort	Fallback if model validation/control evidence remains inadequate

7. Decision and conditions

BASELINE: NO-GO

The original configuration may not proceed because the disparity signal is unresolved and job-relatedness evidence for continuity/title-prestige features is insufficient.

1. Remove or disable identified proxy-risk features.
2. Complete client-job-family validation using a defensible criterion and documented job analysis.
3. Re-run subgroup and intersectional outcome testing; assess statistical uncertainty and practical significance.
4. Complete accessibility review and accommodation / alternative-process design.
5. Configure no-auto-reject control and recruiter review workflow.
6. Complete NYC applicability gate, independent bias audit process, public-summary template, and notice timing where covered.
7. Execute contract controls for model versioning, data use, audit rights, and material-change notice.
8. Launch only as a 90-day monitored pilot with monthly governance review.

8. Sign-off

Role	Name / Decision	Date
HR / Talent Acquisition	HOLD baseline; supports mitigated pilot after controls	
Employment Legal	Required before pilot	
Privacy	Required before pilot	
Security	Required before pilot	
Validation / People Analytics	Required after re-test	
AI Governance Committee	Final go/no-go	