

AI BIAS TEST

NinjaHire

AI Bias Audit Report

Independent bias evaluation of NinjaHire's system, using a purpose-built matched-pair test dataset.



CLIENT
NinjaHire



AUDIT DATE
September 22, 2026



PREPARED BY
AramGRC
www.AramGRC.com



TABLE OF CONTENTS

Table of Contents

1. Executive Summary

1.1 Audit information

1.2 Results overview

2. Auditor & Engagement

2.1 About AramGRC

2.2 Independence statement

2.3 Company information

3. System & Test Scope

3.1 System overview

3.2 Test dataset

3.3 Protected attributes & scope

3.4 Coverage of regulatory requirements

4. Audit Overview

5. Findings

5.1 Sex bias

5.2 Race/Ethnicity bias

5.3 Age bias

5.4 Intersectional (Sex × Race/Ethnicity) bias

6. Audit Approach

6.1 Methodology overview

6.2 Disparate impact analysis

6.3 Counterfactual analysis

6.4 Grading criteria

7. Interpretation & Limitations

EXECUTIVE SUMMARY

Executive Summary

AramGRC was engaged by NinjaHire to perform an independent bias audit of its AI system. This report summarizes that audit, conducted using a purpose-built matched-pair test dataset spanning multiple job roles across NinjaHire's hiring pipeline, and reflects the system's behavior under controlled test conditions at the time of evaluation.

This report is not a formal certification or regulatory compliance determination, and its findings do not predict outcomes for any specific job opportunity or real-world deployment.

Audit information

SYSTEM AUDITED

NinjaHire AI system

AUDIT DATE

September 22, 2026

SAMPLE SIZE

1630

PROTECTED CATEGORIES TESTED

Sex, Race/Ethnicity, Age, Sex × Race/Ethnicity

BIAS DETECTION METHODS

Disparate impact analysis,
Counterfactual analysis

Results overview

CATEGORY	GROUP BIAS	INDIVIDUAL BIAS
Sex bias	Pass	Clear
Race/Ethnicity bias	Pass	Clear
Age bias	Pass	Clear
Intersectional (Sex × Race/Ethnicity) bias	Pass	Clear

Auditor & Engagement

About AramGRC

AramGRC is an independent, third-party AI assurance authority. We audit, assess, and test AI systems against recognized standards, giving organizations independent verification of AI safety, fairness, and compliance for the customers, regulators, and boards who rely on them.

AramGRC's assessors bring audit competency in independent audits and certification support, including ISO/IEC 42001 and EU AI Act conformity, alongside assessment competency spanning AI risk, impact, Responsible AI maturity, and agent governance assessments. Our testing competency extends to AI red-teaming, covering safety, security, and privacy risks.

Independence statement

AramGRC's evaluation and the findings in this report are independent of NinjaHire's engagement of this service. Our compensation for this audit is not tied to the outcome of the results. Our role is limited to testing and reporting on the trustworthiness of the system under test; we do not form part of the solution or otherwise affect how the audited system works.

Company information

WEBSITE

www.AramGRC.com

CONTACT

Contact@AramGRC.com

SYSTEM & TEST SCOPE

System & Test Scope

System overview

NinjaHire's AI system scores inbound and sourced candidates against a job posting, producing a numeric score per candidate that NinjaHire's screening product uses to help recruiters shortlist candidates.

Test dataset

- A purpose-built, synthetic matched-pair test dataset was used, as historical candidate data with sufficient demographic coverage was not available.
- The dataset spans multiple job roles across NinjaHire's hiring pipeline.
- Each matched set consists of candidate resumes that are identical in skills, experience, and qualifications, and differ only in demographic indicators (name, pronouns, graduation year, and related proxies) tied to Sex, Race/Ethnicity, and Age - so any disparity in outcomes can be attributed to the system's treatment of those attributes rather than candidate quality.
- Candidate demographic identity was blinded from NinjaHire at scoring time: every candidate was submitted under an opaque, randomly-generated source ID, with no demographic labels in any field sent to the system under test.

Protected attributes & scope

The audit evaluates potential bias across 4 protected categories: Sex, Race/Ethnicity, Age, and the Sex × Race/Ethnicity intersection, using two bias detection techniques:

- **Disparate impact analysis** - evaluates whether the system adversely impacts a protected group compared to the highest-scoring group. An impact ratio of 0.80 or higher is considered acceptable (the four-fifths rule).
- **Counterfactual analysis** - evaluates whether the system produces consistent outcomes for matched candidates who differ only in a protected attribute. A consistency score of 95% or higher is considered acceptable.

Note: for the disparate impact analysis, each job role's own median score was used as its cutoff for classifying a candidate as "selected," rather than a single cutoff pooled across all roles, since different roles have materially different scoring distributions.

SYSTEM & TEST SCOPE

Coverage of Regulatory Requirements

This audit considers the provisions of applicable AI laws and frameworks that concern evaluating system behaviour for material bias that may lead to unfair outcomes:

- **New York City Local Law 144:** The testing scope and methodology were designed with reference to the bias audit this law requires for automated employment decision tools. The audit reports scoring rates and impact ratios for sex, race/ethnicity, and the intersection of sex and race/ethnicity. It uses a synthetic test dataset because sufficient historical data was not available (see Section 3.2).
- **Additional testing:** The audit also evaluates age (40 and over compared with under 40) and adds counterfactual (matched-pair) analysis of individual-level consistency.
- **California Fair Employment and Housing Act (FEHA):** The testing scope and methodology take into account the anti-bias testing that the FEHA regulations on automated decision systems consider.
- **Texas Responsible Artificial Intelligence Governance Act (TRAIGA, HB 149):** The testing scope and methodology take into account the testing-based defense this law makes available.
- **Illinois HB 3773 and the EEOC Uniform Guidelines:** The testing scope and methodology draw on Illinois's anti-discrimination requirements and on the four-fifths rule.
- **EU AI Act:** The AI Act classifies AI systems used in recruitment and candidate evaluation as high-risk. The testing scope and methodology take into account the Act's expectations for examining these systems for possible biases.
- **NIST AI Risk Management Framework (AI RMF 1.0):** The testing scope and methodology draw on the Measure function, particularly MEASURE 2.11, which calls for evaluating fairness and bias and documenting the results.

Regulatory scope note: This document focuses on provisions related to evaluating system behaviour for material bias that may lead to unfair outcomes. Accordingly, this report should not be interpreted as evidence of full compliance with regulatory obligations.

AUDIT OVERVIEW

Audit Overview

SYSTEM

NinjaHire AI system

AUDIT DATE

September 22, 2026

SAMPLE SIZE

1630

CATEGORIES TESTED

Sex, Race/Ethnicity, Age, Sex ×
Race/Ethnicity

ANALYSIS METHODS

Disparate impact, Counterfactual

Findings at a glance

Sex

GROUP BIAS

Pass

INDIVIDUAL BIAS

Clear

Race/Ethnicity

GROUP BIAS

Pass

INDIVIDUAL BIAS

Clear

Age

GROUP BIAS

Pass

INDIVIDUAL BIAS

Clear

Sex × Race/Ethnicity

GROUP BIAS

Pass

INDIVIDUAL BIAS

Clear

See Findings (Section 5) for the full detailed tables and evidence behind each result.

FINDINGS

Sex bias

GROUP BIAS

Disparate impact analysis

Result: **PASS**

Sample size: 1,630

SEX	SAMPLES	SELECTED	SCORING RATE	IMPACT RATIO	RESULT
Male · reference	543	407	74.95%	1.00	PASS
Non-binary	544	404	74.26%	0.99	PASS
Female	543	407	74.95%	1.00	PASS

AUDIT NOTE

The results indicate equitable outcomes across all groups.

INDIVIDUAL BIAS

Counterfactual analysis

Result: **CLEAR**

Sample size: 1,630

SEX	SAMPLES	CONSISTENCY SCORE	RESULT
Male	543	99.54%	CLEAR
Non-binary	544	99.26%	CLEAR
Female	543	99.35%	CLEAR

AUDIT NOTE

The results indicate highly consistent outcomes across all groups.

FINDINGS

Race/Ethnicity bias

GROUP BIAS

Disparate impact analysis

Result: **PASS**

Sample size: 1,630

RACE/ETHNICITY	SAMPLES	SELECTED	SCORING RATE	IMPACT RATIO	RESULT
White	408	304	74.51%	1.00	PASS
Black or African American	408	305	74.75%	1.00	PASS
Asian · reference	406	304	74.88%	1.00	PASS
Hispanic or Latino	408	305	74.75%	1.00	PASS

AUDIT NOTE

The results indicate equitable outcomes across all groups.

INDIVIDUAL BIAS

Counterfactual analysis

Result: **CLEAR**

Sample size: 1,630

RACE/ETHNICITY	SAMPLES	CONSISTENCY SCORE	RESULT
White	408	99.18%	CLEAR
Asian	406	99.51%	CLEAR
Hispanic or Latino	408	99.67%	CLEAR
Black or African American	408	99.67%	CLEAR

AUDIT NOTE

The results indicate highly consistent outcomes across all groups.

FINDINGS

Age bias

GROUP BIAS

Disparate impact analysis

Result: **PASS**

Sample size: 1,630

AGE	SAMPLES	SELECTED	SCORING RATE	IMPACT RATIO	RESULT
40 and over · reference	815	611	74.97%	1.00	PASS
Under 40	815	607	74.48%	0.99	PASS

AUDIT NOTE

The results indicate equitable outcomes across all groups.

INDIVIDUAL BIAS

Counterfactual analysis

Result: **CLEAR**

Sample size: 1,630

AGE	SAMPLES	CONSISTENCY SCORE	RESULT
40 and over	815	99.26%	CLEAR
Under 40	815	99.26%	CLEAR

AUDIT NOTE

The results indicate highly consistent outcomes across all groups.

FINDINGS

Intersectional (Sex × Race/Ethnicity) bias

GROUP BIAS

Disparate impact analysis

Result: **PASS**

Sample size: 1,630

SEX × RACE/ETHNICITY	SAMPLES	SELECTED	SCORING RATE	IMPACT RATIO	RESULT
Male × White	136	102	75.00%	0.99	PASS
Non-binary × Black or African American	136	101	74.26%	0.98	PASS
Non-binary × White	136	101	74.26%	0.98	PASS
Male × Black or African American	136	102	75.00%	0.99	PASS
Male × Asian	135	101	74.81%	0.99	PASS
Non-binary × Asian	136	101	74.26%	0.98	PASS
Male × Hispanic or Latino	136	102	75.00%	0.99	PASS
Non-binary × Hispanic or Latino	136	101	74.26%	0.98	PASS
Female × Asian · reference	135	102	75.56%	1.00	PASS
Female × Hispanic or Latino	136	102	75.00%	0.99	PASS
Female × Black or African American	136	102	75.00%	0.99	PASS
Female × White	136	101	74.26%	0.98	PASS

AUDIT NOTE

The results indicate equitable outcomes across all groups.

FINDINGS

Intersectional (Sex × Race/Ethnicity) bias

INDIVIDUAL BIAS

Counterfactual analysis

Result: **CLEAR**

Sample size: 1,630

SEX × RACE/ETHNICITY	SAMPLES	CONSISTENCY SCORE	RESULT
Male × White	136	99.60%	CLEAR
Non-binary × Black or African American	136	99.46%	CLEAR
Male × Asian	135	99.46%	CLEAR
Non-binary × Hispanic or Latino	136	99.46%	CLEAR
Male × Hispanic or Latino	136	99.60%	CLEAR
Non-binary × White	136	98.93%	CLEAR
Non-binary × Asian	136	99.46%	CLEAR
Female × Black or African American	136	99.60%	CLEAR
Male × Black or African American	136	99.60%	CLEAR
Female × Hispanic or Latino	136	99.60%	CLEAR
Female × White	136	99.46%	CLEAR
Female × Asian	135	99.60%	CLEAR

AUDIT NOTE

The results indicate highly consistent outcomes across all groups.

AUDIT APPROACH

Methodology Overview

Our methodology for evaluating AI systems is designed to ensure fairness and transparency. Our approach combines controlled, purpose-built test data with two complementary bias-detection techniques, and reports results using a standard, threshold-based grading system.

Black-box testing

We use black-box testing: the audited system is evaluated purely on its inputs and outputs (candidate resume and job description in, match score out), without needing access to its internal model or weights. This lets us evaluate systems built on different underlying models using the same methodology.

Matched-pair (counterfactual) test data

The dataset is anchored in counterfactual profiles: matched sets of candidate resumes that are identical in skills, experience, and qualifications, and differ only in demographic indicators and related proxies. This isolates the effect of the tested attribute from genuine differences in candidate quality, so any disparity in outcomes can be attributed to the system's treatment of that attribute.

Multiple bias detection techniques

Disparate impact analysis and counterfactual analysis provide complementary views: the former measures whether outcomes are equitable at the group level, while the latter measures whether the system treats individuals consistently when only a protected attribute changes. Using both gives a more complete picture than either alone - a system can pass one and still fail the other.

Hybrid auditing

Automated testing is paired with human review at every stage. Running the matched-pair dataset through the audited system programmatically is what makes an audit of this size practical, scoring every candidate in every matched set without the cost or delay of a manual process. But the dataset itself, and the results that come out of it, are both checked by hand: an auditor reviews the generated candidate profiles for realism and consistency before they are used, and reviews the computed results afterward to confirm the grading reflects genuine model behavior.

Grounded in established standards

Our thresholds and definitions draw on established fairness research and regulatory guidance, including the EEOC's Uniform Guidelines (the four-fifths rule) and the academic literature on counterfactual fairness and matched-pair audit testing.

AUDIT APPROACH

Disparate Impact Analysis

Disparate impact analysis evaluates whether a protected demographic group is adversely affected compared to the highest-scoring group. This is achieved by comparing each group's scoring rate to that of the highest-scoring group. The goal is to ensure the system does not disproportionately disadvantage any group based on a protected characteristic.

Scoring rate

Scoring rate measures the proportion of individuals in a group who receive a favorable outcome (a score at or above the applicable cutoff).

$$\text{Scoring rate} = \frac{\text{\# candidates in group scoring} \geq \text{the cutoff}}{\text{Total candidates in group}}$$

Impact ratio

Impact ratio measures potential adverse impact on a group by comparing its scoring rate to that of the highest-scoring group.

$$\text{Impact ratio} = \frac{\text{Scoring rate for the group}}{\text{Scoring rate of the highest-scoring group}}$$

An impact ratio of 1.00 indicates no adverse impact; a lower ratio indicates a higher likelihood of adverse impact. Per the four-fifths rule (EEOC Uniform Guidelines, 1978), an impact ratio of 0.80 or higher is considered acceptable.

AUDIT APPROACH

Counterfactual Analysis

Counterfactual analysis assesses fairness by examining how the system's decisions would change if a candidate's demographic attributes (or proxies for them) were altered while everything else about their resume stayed the same. This determines whether outcomes are influenced by protected attributes, so that individuals receive similar treatment regardless of their demographic characteristics.

Matched-set process

- **Generate** - candidates are grouped into matched sets sharing identical resume content (template, role, and every other tested attribute held fixed) and differing only in the one attribute under test.
- **Execute** - every candidate in every matched set is run through the AI system being tested.
- **Measure** - within each matched set, every pairwise comparison is checked for whether both candidates land on the same side of the cutoff (both selected, or both not selected). A pair is "inconsistent" if one is selected and the other is not.

$$\text{Consistency score} = \frac{\text{Consistent pairs}}{\text{Total pairs compared}}$$

A higher consistency score means the system's decisions are less influenced by the attribute under test.

Grading criteria

Results are graded using a threshold-based system indicating the severity of any bias detected:

GRADE	DESCRIPTION	DISPARATE IMPACT ANALYSIS	COUNTERFACTUAL ANALYSIS
Pass / Clear	No issues detected	Impact ratio $\geq 80\%$	Consistency score $\geq 95\%$
Consider	Minor issue(s) detected	n/a	$90\% \leq \text{Consistency score} < 95\%$
Flag for review / Concern	Issue(s) requiring attention	Impact ratio $< 80\%$	Consistency score $< 90\%$

INTERPRETATION & LIMITATIONS

Interpretation & Limitations

This report has been prepared by AramGRC to provide an independent audit of the AI system developed by NinjaHire, based on the methodology described above and a purpose-built synthetic test dataset. The limitations below define the boundaries within which its findings should be read.

Scope limitations

The results and conclusions presented in this report reflect our professional judgment based on the information available at the time of evaluation. While reasonable steps were taken to ensure accuracy and completeness, the evaluation may not be exhaustive and cannot be guaranteed to be entirely free of error.

Dataset limitations

The system was evaluated on a purpose-built, synthetic matched-pair test dataset, as sufficient historical production data with controlled demographic coverage was not available. A synthetic dataset, however carefully constructed, cannot capture every real-world candidate profile or scenario the system may encounter in production.

Interpretation limitations

Our methodology is designed to identify potential issues of bias in the system under examination. However, our approach, like any evaluation methodology, has limitations. Our findings do not guarantee the absence of any bias, flaws, or limitations within the audited system; they indicate that, based on our specific testing framework and within the scope of our analysis, no significant issues were identified.

Deployment limitations

This report is intended for informational purposes only and should not be interpreted as a guarantee of the system's performance, fairness, or suitability for any specific purpose or use case. Its findings do not predict outcomes for any specific job opportunity or real-world deployment.

Regulatory limitations

This report is not a regulatory compliance determination, and does not determine whether the audited system is legally classified as an AEDT or other regulated technology.

AI BIAS TEST

NinjaHire

AI Bias Audit Report



CLIENT
NinjaHire



AUDIT DATE
September 22, 2026



PREPARED BY
AramGRC
www.AramGRC.com

