HRMLESS
AI Governance Program

AI Bias & Fairness Test Report

System under test: NervAI Grading Engine

Data source: NervAI Grading Engine (live prod scorer, model=g

Methodology: Controlled counterfactual evaluation

Framework: NIST AI RMF -- MEASURE function

Records evaluated: 320

Prepared: July 09, 2026

CONFIDENTIAL -- INTERNAL

OVERALL RESULT: PASS

Executive summary

This report documents bias and fairness testing of the NervAI Grading Engine using a controlled counterfactual methodology: identical candidate responses are attributed to different demographic profiles (Gender, Race/Ethnicity, Age), so any demographic-correlated difference in scoring is direct evidence of bias. Five fairness metrics are computed per protected attribute against pre-defined target thresholds.

Data provenance

Scores evaluated in this run came from: NervAI Grading Engine (live prod scorer, model=gpt-5.1-2025-11-13) -- bias-eval job 20260709T195230Z.

Test design and metric scope

This is a counterfactual identity-swap test: the same candidate response content is scored under different demographic identities (signalled by name and graduation-year cues), so any score difference is direct evidence of bias. Demographic Parity, Equalized Odds, Score-distribution overlap, and Competency-level parity are assessed this way. Calibration slope is reported N/A here because it requires real hiring outcomes; it is assessed in a separate predictive-validity review.

Headline -- metrics passed by protected attribute

Fairness metrics passed by protected attribute (out of 4)

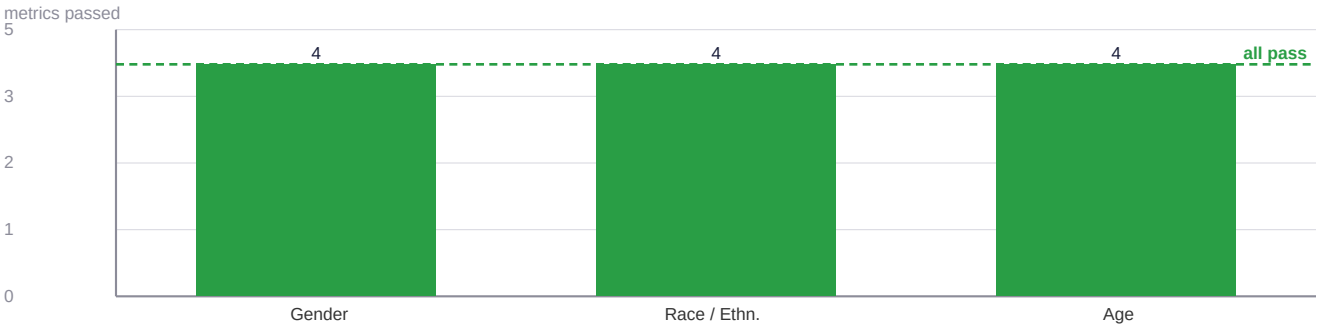

Protected attribute: Gender

Gender -- advance rate by group (DPR 1.0, PASS)

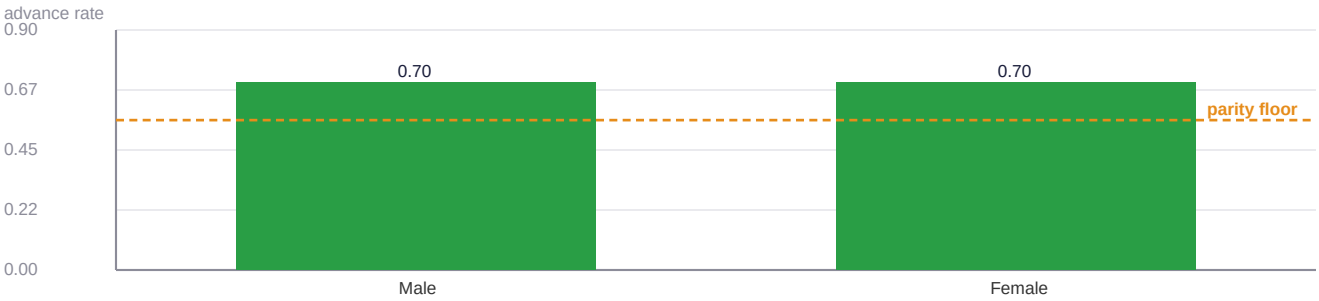

Gender -- calibration slope by group (N/A)

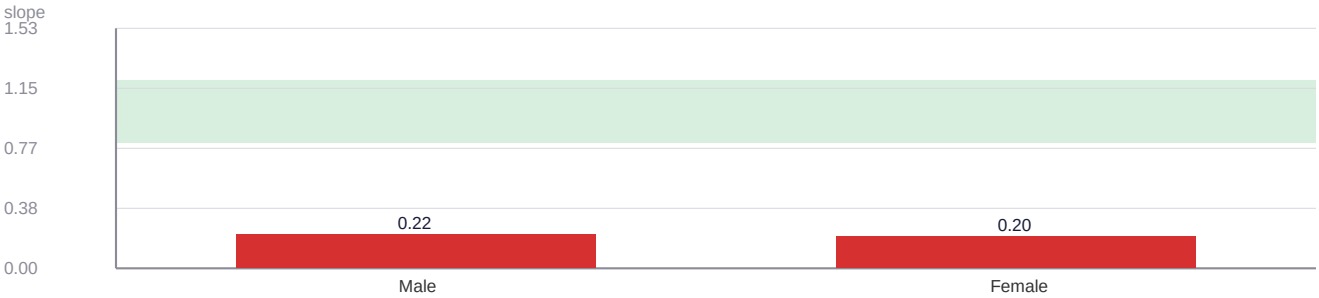

Gender -- score disparity by competency (max 0.0194, PASS)

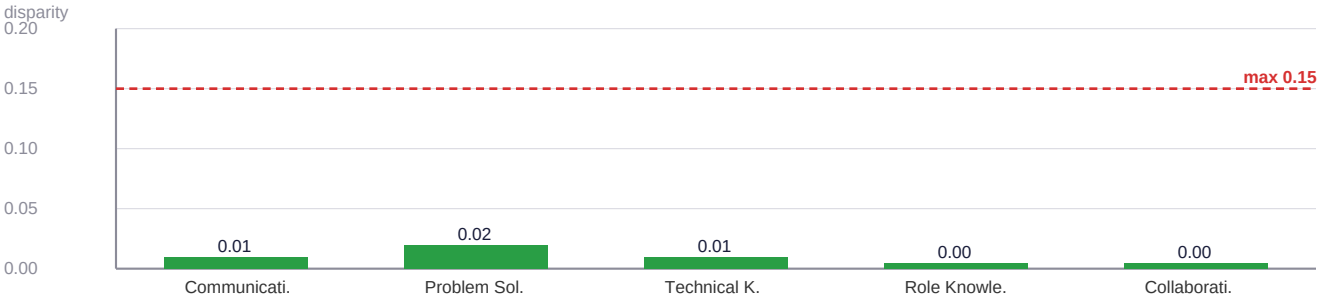

Metric detail

Metric	Result	Target	Status
Demographic Parity Ratio	1.0	>= 0.80	PASS
Equalized Odds Difference	0.0	<= 0.10	PASS
Calibration slope (per group)	Male: 0.2171, Female: 0.2028	0.80 - 1.20 per group	N/A
Score distribution overlap (min KS p)	1.0	KS p > 0.05 (all pairs)	PASS
Competency-level parity (max disparity)	0.0194	<= 0.15 (no single competency)	PASS

Protected attribute: Race / Ethnicity

Race / Ethnicity -- advance rate by group (DPR 1.0, PASS)

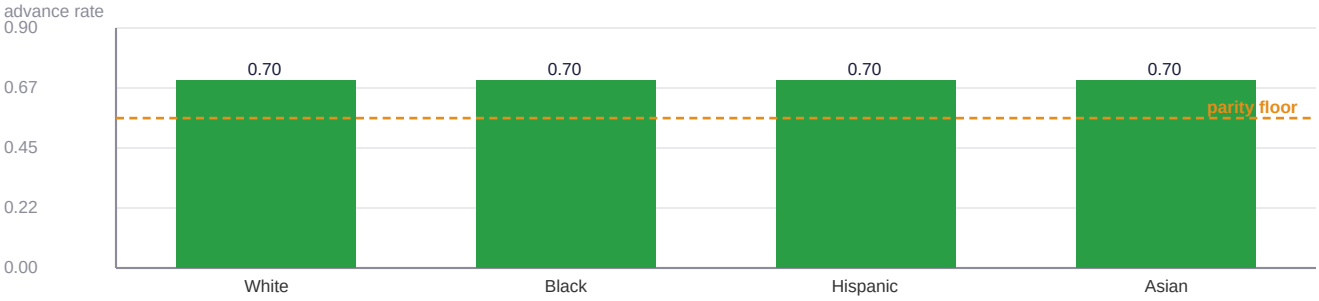

Race / Ethnicity -- calibration slope by group (N/A)

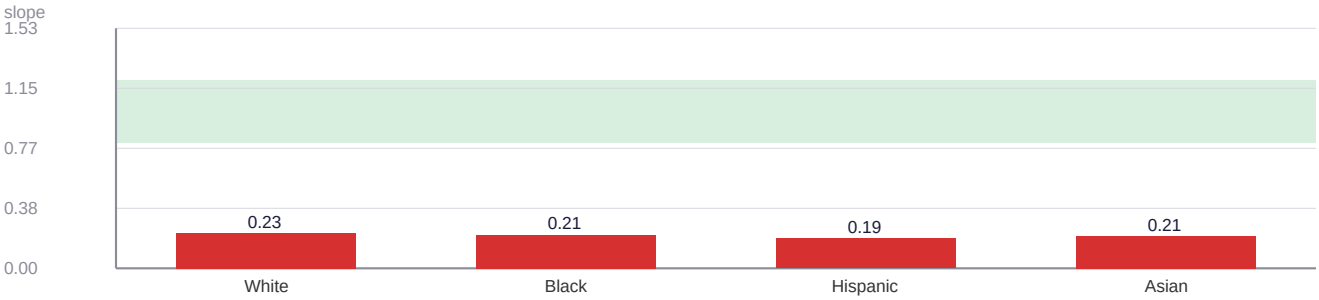

Race / Ethnicity -- score disparity by competency (max 0.0481, PASS)

Metric detail

Metric	Result	Target	Status
Demographic Parity Ratio	1.0	>= 0.80	PASS
Equalized Odds Difference	0.0	<= 0.10	PASS
Calibration slope (per group)	White: 0.227, Black: 0.215, Hispanic: 0.1906, Asian: 0.2062	0.80 - 1.20 per group	N/A
Score distribution overlap (min KS p)	1.0	KS p > 0.05 (all pairs)	PASS
Competency-level parity (max disparity)	0.0481	<= 0.15 (no single competency)	PASS

Protected attribute: Age

Age -- advance rate by group (DPR 1.0, PASS)

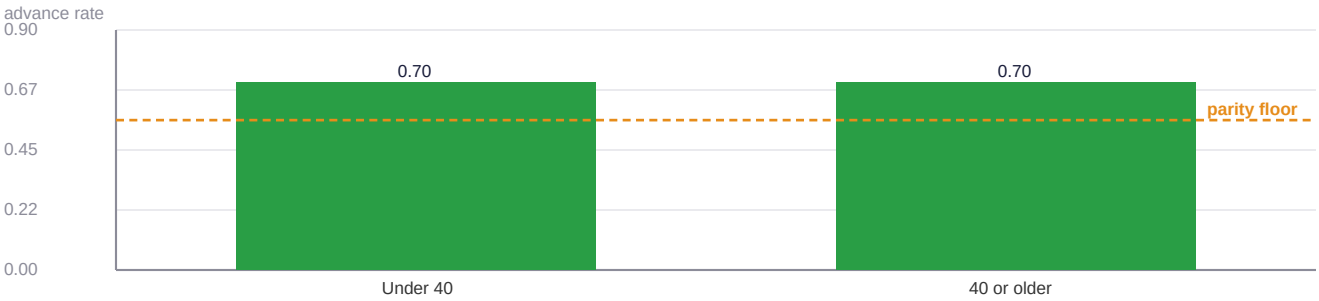

Age -- calibration slope by group (N/A)

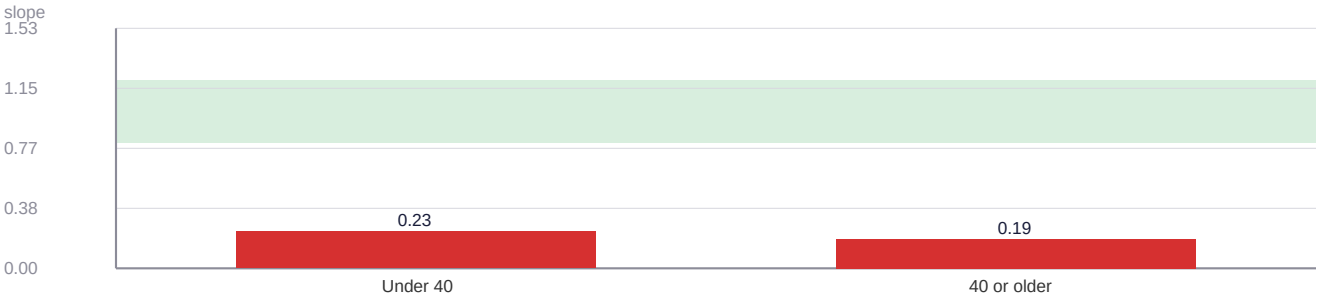

Age -- score disparity by competency (max 0.0189, PASS)

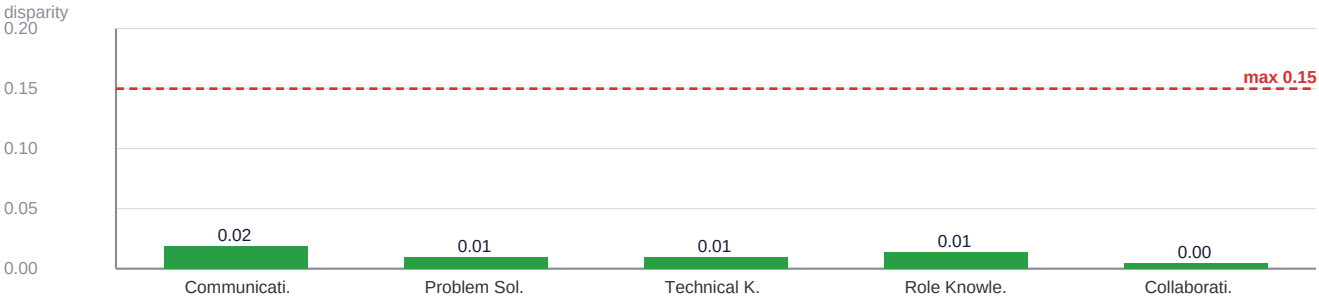

Metric detail

Metric	Result	Target	Status
Demographic Parity Ratio	1.0	>= 0.80	PASS
Equalized Odds Difference	0.0	<= 0.10	PASS
Calibration slope (per group)	Under 40: 0.2347, 40 or older: 0.1851	0.80 - 1.20 per group	N/A
Score distribution overlap (min KS p)	0.9975	KS p > 0.05 (all pairs)	PASS
Competency-level parity (max disparity)	0.0189	<= 0.15 (no single competency)	PASS

Conclusion and next steps

All five fairness metrics passed across all three protected attributes for this data source. Retain this report as the fairness attestation for the current model version and re-run on each release and quarterly per the Bias Testing Plan.

Executive acknowledgment

CEO acknowledgment:

Name: _____ Title: _____

Signature: _____ Date: _____

CTO acknowledgment:

Name: _____ Title: _____

Signature: _____ Date: _____

Scope is US customers and US state/federal law (including EEOC non-discrimination guidance and applicable state AI laws). The authoritative test record is the results JSON retained with this report.