

Skills, Not Keywords: Inference, Auditability, and the Claims Employers Should Not Make

Heather Grizzle · May 18, 2025

Replacing keyword filters with AI skills inference does not remove proxies. It moves them somewhere the employer cannot inspect, and raises the burden of proof.

Contents

1. The proposition
2. What the evidence shows about the practice
3. Inference does not remove the proxy
4. The taxonomy is a construct, and someone built it
5. An auditable instrument traded for an opaque one
6. Skills-based framing raises the validation burden, not lowers it
7. The claim not to make
8. What to establish before deployment
9. Conclusion
10. References

The proposition

The most substantive claim currently made for artificial intelligence in recruitment is that it permits hiring for skills rather than keywords, and that doing so produces fairer outcomes, wider access, and more diverse teams. It is the strongest of the available arguments

because, unlike claims about efficiency or engagement, it identifies a real defect in the thing it proposes to replace. Keyword matching is crude. It rewards résumé vocabulary rather than capability, filters out people who describe the same work in different language, and encodes credential proxies that exclude candidates who acquired their skills outside conventional paths.

All of that is true, and it does not establish that the replacement is an improvement in the respects claimed. Novara Consulting Group's position is that the substitution of inferred skills for matched keywords changes the character of the instrument in ways that matter for governance, that most of those changes make the employer's position harder rather than easier to defend, and that one of the claims routinely attached to it should not be made in public at all.

What the evidence shows about the practice

Before examining the technology it is worth noting what has actually happened where organizations adopted skills-based hiring as policy.

In February 2024 the Harvard Business School Project on Managing the Future of Work and the Burning Glass Institute published a study of 11,300 roles at large firms, examining hiring at least one year before and after degree requirements were removed. Dropping the requirement produced an average increase of roughly 3.5 percentage points in the hiring of workers without a bachelor's degree, and because the change applied only to the small share of roles that dropped a requirement, the net effect across the labor market amounted to fewer than one in seven hundred hires. Close to half the firms studied fell into a category the authors termed in name only, where a stated commitment produced no measurable change in hiring behavior.

BY THE NUMBERS

11,300

roles at large firms examined before and after degree requirements were removed

3.5 pp

average increase in hiring of workers without a bachelor's degree

"~50% "

"of

Source: Fuller et al., Skills-Based Hiring: The Long Road from Pronouncements to Practice, February 2024

The finding is not that skills-based hiring cannot work. Firms that changed their processes rather than only their postings did see movement. The finding is that announcement and practice diverged almost completely, and that the divergence went unnoticed because nobody was measuring the second one. An organization intending to adopt this approach should assume its own results will resemble the average unless it measures, and should treat the gap between what it says and what its hiring data shows as the thing most likely to be produced in a future dispute.

Inference does not remove the proxy

The technical claim is that a skills model reads past surface vocabulary to the underlying capability. What it actually does is derive skill attributions from the same document the keyword filter read, using a longer and less inspectable chain of inference.

The model does not observe that a candidate can manage a project. It observes a job title, at an employer, for a duration, described in particular language, and assigns a skill on the basis of associations learned from prior data about what people holding similar titles at similar employers were taken to be able to do. Every proxy available to the keyword filter remains available to the inference model, and several more become available, because the model can draw on institutional prestige, career continuity, phrasing conventions, and patterns of progression that a keyword list could not represent.

This matters most for the candidates the approach is meant to reach. Skills acquired through caregiving, community organizing, military service, informal or family enterprise, disability self-management, or an interrupted career are precisely the skills least well represented in training data drawn from

conventional employment histories. The candidate whose capability is not legible in the ordinary corpus was disadvantaged by the keyword filter and remains disadvantaged by the model, with the difference that she can no longer see why.

The candidate whose capability is not legible in the ordinary corpus was disadvantaged by the keyword filter and remains disadvantaged by the model, with the difference that she can no longer see why.

The taxonomy is a construct, and someone built it

Every skills-based system rests on a taxonomy: an enumerated set of skills, their definitions, their relationships, and the evidence taken to demonstrate each. That

taxonomy is not a natural object. It was constructed by people, encoding judgments about what counts as a skill, which skills are distinguishable from one another, and what demonstrates possession.

Organizations adopting these systems almost never examine the taxonomy, and in many products cannot. The result is that the definition of qualification for a role has been outsourced to a vendor's ontology, and the employer using it will nonetheless be the party asked to justify it. When the question arrives asking why a particular attribute was treated as required, the answer that it was in the vendor's skills library is not responsive.

An auditable instrument traded for an opaque one

This is the governance point and it is the one most consistently missed.

A keyword filter is a bad instrument with one substantial virtue: it is legible. The list can be read. A person can look at it and observe that the filter requires a term that has nothing to do with the job, or that excludes a synonym common in a particular sector or region. Where it produces adverse impact, the component responsible can be identified, because the components are enumerated.

A skills inference model has neither property. Its criteria cannot be read off, its components cannot be individually inspected by the employer, and the reason a given candidate scored as she did is not recoverable from anything the employer holds. The organization has replaced a crude transparent instrument with a sophisticated opaque one and described the exchange as an improvement in fairness. What it improved is sophistication. What it gave up is the only thing that made the prior instrument defensible.

This has direct consequences under existing federal law. The Uniform Guidelines on Employee Selection Procedures require that where a total selection process shows adverse impact, the user identify which components carry it and maintain evidence of validity for

those components. A process built from an enumerated filter permits that analysis. A process built from an inference model, on vendor infrastructure, with scores discarded after ranking, generally does not.

Skills-based framing raises the validation burden, not lowers it

There is a further consequence that organizations adopting this language have not typically considered.

A keyword filter makes no claim about what it measures. It matches strings. A skills-based system makes an explicit assertion: that it measures identified, job-related competencies. That assertion is a claim of construct validity, made by the employer, in public, about its own selection process.

Under the Uniform Guidelines, an employer whose selection procedure produces adverse impact is expected to produce evidence that the procedure is valid for the job. An employer that has publicly characterized its process as measuring specific job-related skills has volunteered the framework against which it will be judged, and has done so before establishing that it can meet it. The organization should therefore be able to answer, for each scored construct, which documented duty of the role it corresponds to and on what evidence the correspondence rests. Where the answer is that the vendor named the skill, the employer has adopted a validity claim it did not make and cannot support.

The claim not to make

One assertion attached to skills-based hiring should be removed from all external communication: that the approach produces more diverse teams.

The difficulty is not that diversity is an improper goal. It is that the sentence, read by a lawyer, states that the organization's selection process affects the protected-class composition of its workforce. Title VII, as amended in 1991, makes it unlawful in connection with the selection or referral of applicants to adjust scores, use different cutoff scores, or otherwise alter the results of employment-related tests on the basis of race, color, religion, sex, or national origin. Nothing here suggests that skills-based tools do this. What the public claim does is invite the question, in a document the employer wrote and cannot retract, in an environment where such questions are being asked with increasing frequency and where Executive Order 14281 has directed federal agencies to deprioritize disparate impact enforcement while doing nothing to reduce the exposure of employers who make affirmative representations about demographic outcomes.

There is also the ordinary problem of substantiation. An organization asserting that its process yields more diverse teams should be able to produce the before-and-after data supporting it, controlled for applicant pool composition and hiring volume. Most cannot, which means the claim is unverified as well as unwise.

THE DEFENSIBLE FORMULATION

The organization is broadening the range of qualifications it will consider, has documented the job-related basis for each, and monitors outcomes. That statement is accurate, is consistent with the objective, and asserts nothing the organization would rather not be asked to prove.

The defensible formulation is available and costs nothing. The organization is broadening the range of qualifications it will consider, has documented the job-related basis for each, and monitors outcomes. That statement is accurate, is consistent with the objective, and asserts nothing the organization would rather not be asked to prove.

What to establish before deployment

Five things, answerable in writing and from evidence rather than from vendor description, before a skills inference system touches a selection decision.

1. **What is the taxonomy, who built it, and can the organization read it.** A system whose skill definitions cannot be inspected has placed the definition of qualification outside the organization's control while leaving the justification inside it.
2. **For each scored construct, which documented duty of the role does it correspond to.** This should be written by the organization, not supplied by the vendor, and it should exist before the system runs rather than after a charge is filed.
3. **Are the score, the inputs, and the model version retained for each decision, and can a specific decision be reproduced later.** If not, the organization cannot perform the component analysis the Uniform Guidelines contemplate, and should know that before it needs to.
4. **What does the organization's own hiring data show, measured before and after.** Given the published gap between pronouncement and practice, this is the only way to know whether anything changed, and it is also the record that establishes good faith.
5. **What is the organization saying publicly, and can it substantiate every sentence.** Marketing copy about hiring outcomes is discoverable, and it is written by people who are not consulted when it is later produced.

Conclusion

Hiring for skills rather than keywords is a sound objective and the criticism of keyword matching that motivates it is correct. What the current technology delivers is not the removal of proxies but their migration into a layer the employer cannot inspect, accompanied by a public claim of construct validity the employer has not established and a claim about demographic outcomes it would be better off not making.

The organizations that will benefit from this shift are the ones that treat it as a measurement problem rather than a purchasing decision: that define the skills themselves, document the job-related basis, retain the record, measure what actually changed, and describe the result in language they could defend under oath. That is more work than adopting the vocabulary. It is also the only version of skills-based hiring that has been shown to move anything.

References

- Executive Order 14281, Restoring Equality of Opportunity and Meritocracy, April 23, 2025.
- Fuller, Joseph, et al. *Skills-Based Hiring: The Long Road from Pronouncements to Practice*. Harvard Business School Project on Managing the Future of Work and The Burning Glass Institute, February 2024.
- Title VII of the Civil Rights Act of 1964, as amended, 42 U.S.C. § 2000e-2, including § 2000e-2(l).
- Uniform Guidelines on Employee Selection Procedures, 29 C.F.R. Part 1607, including §§ 1607.4 and 1607.15; EEOC recordkeeping requirements, 29 C.F.R. Part 1602.