VIET NAM NATIONAL UNIVERSITY HO CHI MINH CITY

VNUHCM Innovation Hub

IESF PARTNERS & MANAGING DIRECTORS

Match. Evaluate. Rank.

What AI can and cannot do in leadership search

Dr. Trinh Quang Viet

Director, Institute of Applied Artificial Intelligence · VNUHCM Innovation Hub

September 2026

THE ARGUMENT

Three verbs. Three very different answers.

VNUHCM Innovation Hub

MATCH

Who is in the universe?

AI is transformative here.

Volume, language, structure.
Machines are unreasonably good at all three.

EVALUATE

What are they capable of?

AI helps — indirectly.

It cannot judge. It can force structure and capture evidence, which is what actually predicts.

RANK

Who should get the job?

AI is a liability here.

Ordering people is the act regulators, courts and plaintiffs are now looking at.

The mistake is treating them as one problem and buying one tool for all three.

WHY NOW

Four numbers that changed the job

7.1 years

average tenure of an outgoing CEO in 2025 — down from 8.3 in 2021

Russell Reynolds

446

CEO exits at US public companies in 2025 — a record since tracking began in 2002

Challenger, Gray & Christmas

11,000/min

job applications on LinkedIn, up 45% in one year, generative AI credited

LinkedIn via NYT

11%

of talent leaders say their executives are ready to lead the AI transition

Korn Ferry, 1,600+ leaders

More decisions. Less time. Worse signal. And clients who do not trust their own bench. That is not a profession being disintermediated — it is a higher price on being right.

Sources: Russell Reynolds Associates 2025 Global CEO Turnover Index; Challenger, Gray & Christmas 2025 CEO Turnover Report; LinkedIn data via The New York Times (2025); Korn Ferry Talent Trends 2026.

THE PROBLEM UNDERNEATH

Every written artifact just lost its signal

ARTIFACT	WHAT AI CAN NOW FABRICATE	SIGNAL REMAINING
CV / résumé	Fully generatable	None
Cover letter, application answers	Fully generatable	None
Public profile & thought leadership	Largely generatable	Consistency over time
Video interview presence	Increasingly synthesisable	Live, unscripted probing
Attributable reference testimony	Not generatable	Substantially intact

Gartner projects that by 2028 roughly one in four candidate profiles worldwide could be fake.

KnowBe4 — a security firm — hired a North Korean operative on a stolen identity. He cleared four video interviews and a background check.

Sources: Gartner projection reported via HR Dive; KnowBe4 published incident report (2024). Column three is the author's assessment, not measured data.

VIET NAM NATIONAL UNIVERSITY HO CHI MINH CITY

VNUHCM Innovation Hub

01

MATCHING

Who is in the universe — and how do we know?

Three generations

Building a map

Where it fails

MATCHING · 01

Three generations of matching

GEN 1 1998–2017

Keyword & Boolean

Matches on: literal strings

Fails on: synonyms, translation, any career that does not use your vocabulary

GEN 2 2018–2023

Semantic embeddings

Matches on: meaning in vector space

Fails on: it inherits every association in the training corpus, including names

GEN 3 2024–

Capability graphs & inference

Matches on: inferred skills, trajectory, adjacency

Fails on: it asserts things about people that nobody told it

Each generation widened the net and weakened the audit trail. Gen 3 makes claims no human ever stated.

MATCHING · 02

How a market map is actually built now

VNUHCM Innovation Hub

1

Define the universe

Sectors, geographies, company archetypes, scale bands

2

Harvest public signal

Filings, registries, press, conference and publication records

3

Resolve identity

De-duplicate, disambiguate, reconstruct org charts

4

Infer capability

Trajectory, scope, adjacency — machine-asserted, not stated

5

Verify with humans

Confirm, correct, and reject what the machine asserted

Steps 1 - 4 are now hours of compute, not weeks of research

Step 5 is the product. It is also the only step a client cannot buy as software.

Six signals in an executive match

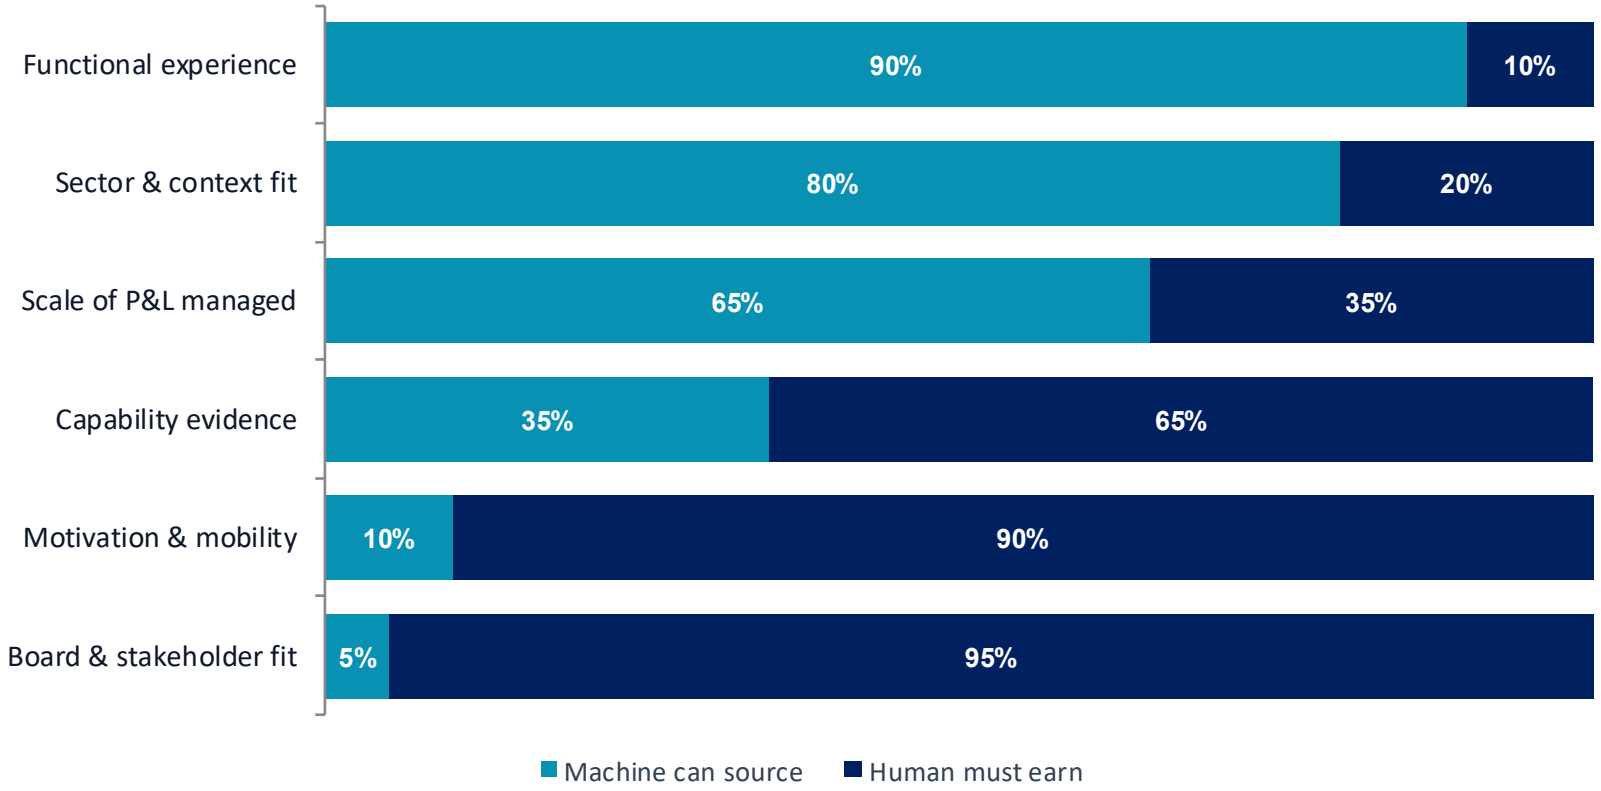

Read it top-down

The signals a machine can supply are the ones a client can also buy directly. The signals it cannot are the ones that decide whether the appointment survives year two.

A framework, not a finding

These proportions are my working model for structuring a search, offered for argument. No study measures them. Treat any vendor who presents similar splits as measured fact with suspicion.

VNUHCM Innovation Hub

How matching goes wrong: proxies

INNOCUOUS VARIABLE	PROTECTED CHARACTERISTIC
ZIP / postal code	Race, national origin
University attended	Class, nationality, age
Continuous tenure	Parental leave, disability
Name on the CV	Race and gender

Measured, not alleged

Wilson & Caliskan (2025) tested the embedding models underneath modern CV screening:

- White-associated names favoured in 85.1% of cases
- Female-associated names favoured in 11.1%
- Black male names disadvantaged in up to 100% of comparisons

Already illegal in one US state

Illinois HB 3773, in force since 1 January 2026, specifically prohibits using ZIP codes as a proxy for a protected class.

Nobody builds a discriminatory model. They build a model that learns what past hiring rewarded.

Sources: Wilson & Caliskan, Brookings AI Equity Lab / AAAI-ACM AIES (2025); Illinois Public Act 103-0804 (HB 3773).

VIET NAM NATIONAL UNIVERSITY HO CHI MINH CITY

VNUHCM Innovation Hub

02

EVALUATING

What is this person actually capable of?

- The ceiling
- Structure over scores
- The red list

The ceiling on prediction

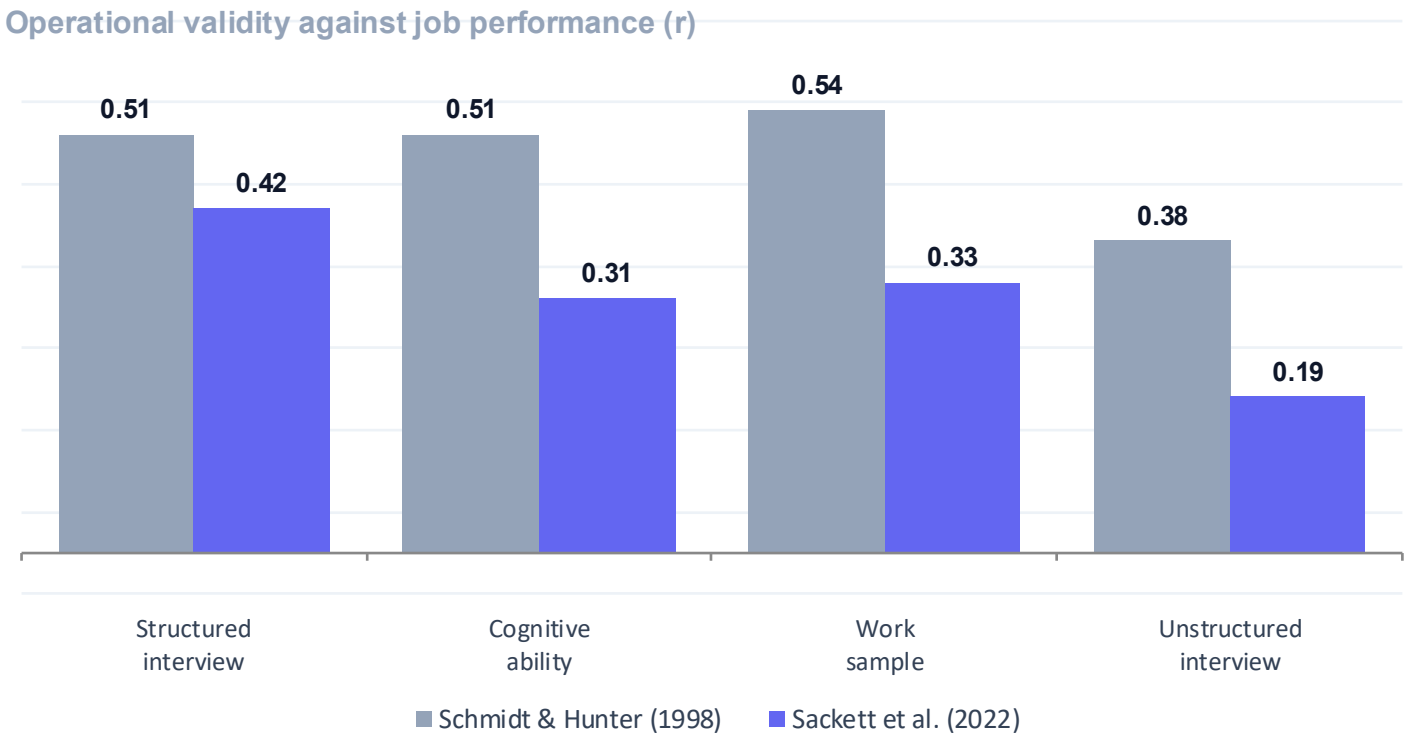

$0.42^2 \approx 0.18$

The best single method explains under a fifth of the variance in performance.

Three questions that end a sales pitch

When a platform claims 90% accuracy, ask:

1. What was the criterion measure?
2. What was the sample?
3. What was the time lag?

Most cannot answer all three.

Structured interviews now outrank cognitive ability — reversing what the profession has quoted since 1998.

Source: Sackett, Zhang, Berry & Lievens (2022), Journal of Applied Psychology 107(11), 2040–2068, compared with Schmidt & Hunter (1998).

VNUHCM Innovation Hub

Score-first versus evidence-first

SCORE-FIRST

- Candidate
- **Opaque model**
- 4.2 / 5
- Rank order

What survives a challenge?

Nothing. You cannot show the client — or a court — why this person is above that person.

EVIDENCE-FIRST

- Candidate
- Same structured probes
- Evidence captured verbatim
- Human judgment, recorded

What survives a challenge?

All of it. The reasoning is on the record, attributable to a named person.

AI belongs in the middle two boxes on the right — never in the black box on the left.

EVALUATING · 03

The highest-return application available today

1

Consent & record

Explicit, documented, per jurisdiction

2

Transcribe

Multilingual, speaker-separated

3

Map to competencies

Against your model, not the vendor's

4

Extract evidence

Verbatim quotes, not summaries

5

Comparable scorecard

Same fields for every candidate

What it removes

Recall bias, note-taking gaps, and the consultant who probed differently on a Friday.

What it creates

Two consultants in two countries producing genuinely comparable assessments of one role.

What it must not do

Score. Rank. Or infer emotion, personality or motivation from voice or face.

EVALUATING · 04

The red list

Emotion or affect inference	Prohibited in EU workplaces under AI Act Article 5 — in force since February 2025, and untouched by the 2026 deferral.
Personality inferred from voice or face	HireVue removed visual analysis from new models in 2021 after its own research found it added little.
Rejection with no human in the loop	The core allegation in Mobley v. Workday: candidates screened out before any person saw them.
Scoring the candidate cannot see	Kistler v. Eightfold AI, filed January 2026 — a hidden 0–5 'likelihood of success' score.
Third-party data gathered without notice	Scraped profiles, tracking data and inferred attributes now carry consumer-reporting exposure.

If a demo shows you any of these five, the meeting is over. That is a cheap rule to enforce.

Sources: Regulation (EU) 2024/1689 Art. 5; HireVue / ORCAA disclosure (Jan 2021); Mobley v. Workday, N.D. Cal.; Kistler v. Eightfold AI (Cal., Jan 2026).

VIET NAM NATIONAL UNIVERSITY HO CHI MINH CITY

VNUHCM Innovation Hub

03

RANKING

Ordering people is the act that creates liability.

Two landmark cases

The defensible slate

The clock

RANKING · 01

Two cases, two different theories

Mobley v. Workday — the outcome theory

The claim is that the tool produced discriminatory outcomes.

- Preliminary ADEA collective certified 16 May 2025
- Extended July 2025 to applicants screened via HiredScore, with an order to identify which customers had it enabled
- Workday's filings referenced roughly 1.1 billion applications rejected in the relevant period
- The court has treated a screening tool as capable of acting as the employer's agent

Kistler v. Eightfold — the secrecy theory

The claim is not that the score was biased. It is that it was hidden.

- Class action filed 20 January 2026
- Alleges a concealed 0–5 'likelihood of success' score used to rank and reject
- Argues the vendor is an unregistered consumer reporting agency under FCRA and California's ICRAA
- Brought by former EEOC chair Jenny Yang with Towards Justice and Outten & Golden

Between them the two theories close the gap: proving your tool is fair no longer protects you if the candidate was never told it existed.

Sources: Mobley v. Workday, No. 3:23-cv-00770 (N.D. Cal.); Kistler v. Eightfold AI, filed 20 Jan 2026 (California); law-firm analyses (Holland & Knight, Epstein Becker Green, Jones Walker).

RANKING · 02

What a defensible slate looks like

RANK	CAPABILITY EVIDENCE	MOTIVATION / MOBILITY	RISK FLAGS	WHO DECIDED
1	4 of 4 competencies evidenced; verbatim record	Confirmed in two conversations	None	*R. Ngo, partner*
2	4 of 4 evidenced; one thinner on scale	Confirmed; notice period 9 months	Notice period	*R. Ngo, partner*
3	3 of 4 evidenced; turnaround untested	Willing, conditional on equity	Comp gap	*A. Silva, partner*
—	Machine-surfaced, not yet verified	Unknown	Unverified identity	*Not yet assessed*

No score anywhere on the slate. The order is an argument a named human being can defend — to the client, and to a court.

RANKING · 03

Testing a ranking — and why most do not

NYC Local Law 144 has required published bias audits since July 2023

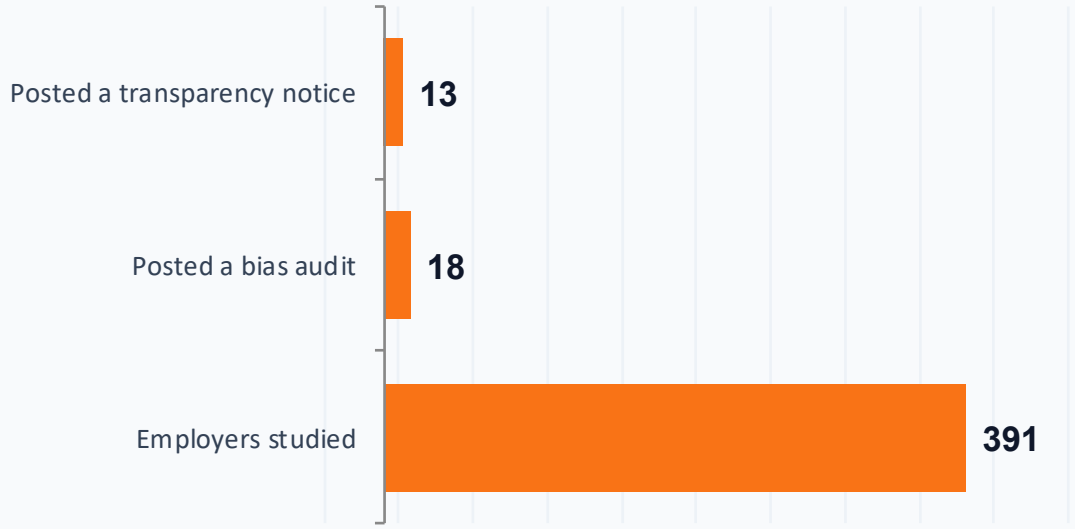

”Null compliance”

4.6% published an audit. 3.3% posted a notice. The law lets the employer decide whether its tool "substantially assists" the decision — so employers decided it did not.

The four-fifths rule, in one line

Selection rate of the lowest group ÷ selection rate of the highest. Below 0.8 signals adverse impact.

The same study found nearly every published audit reported a ratio above 0.8 — which tells you more about who chooses to publish than about the tools.

Run the ratio on your own placements. Quietly, annually, and before someone else does.

Source: Wright, Muenster et al., "Null Compliance: NYC Local Law 144 and the challenges of algorithm accountability", ACM FAccT 2024 (Cornell / Data & Society / Consumer Reports).

RANKING · 04

The clock

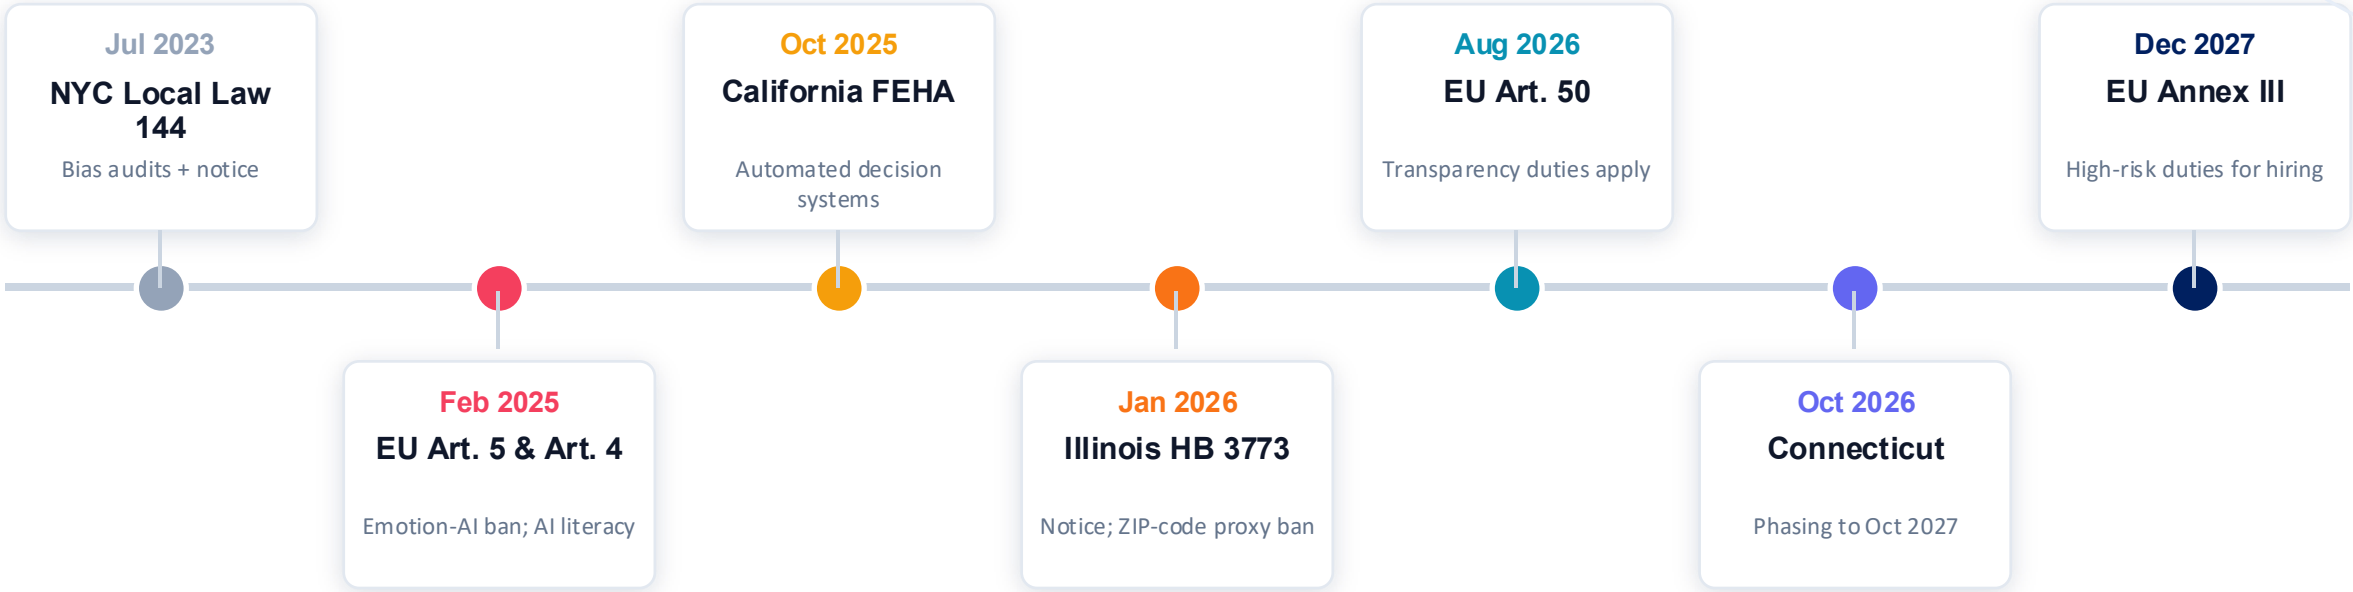

Europe deferred the high-risk hiring rules from August 2026 to December 2027 under Regulation (EU) 2026/1744. The Article 5 prohibitions, the AI-literacy duty and the Article 50 transparency rules did not move. Deferral is preparation time, not permission.

Sources: Regulation (EU) 2024/1689 and Regulation (EU) 2026/1744; NYC Local Law 144; California FEHA ADS regulations; Illinois PA 103-0804; Connecticut AI Responsibility and Transparency Act.

PUTTING IT TOGETHER

The division of labour

AI LEVERAGE

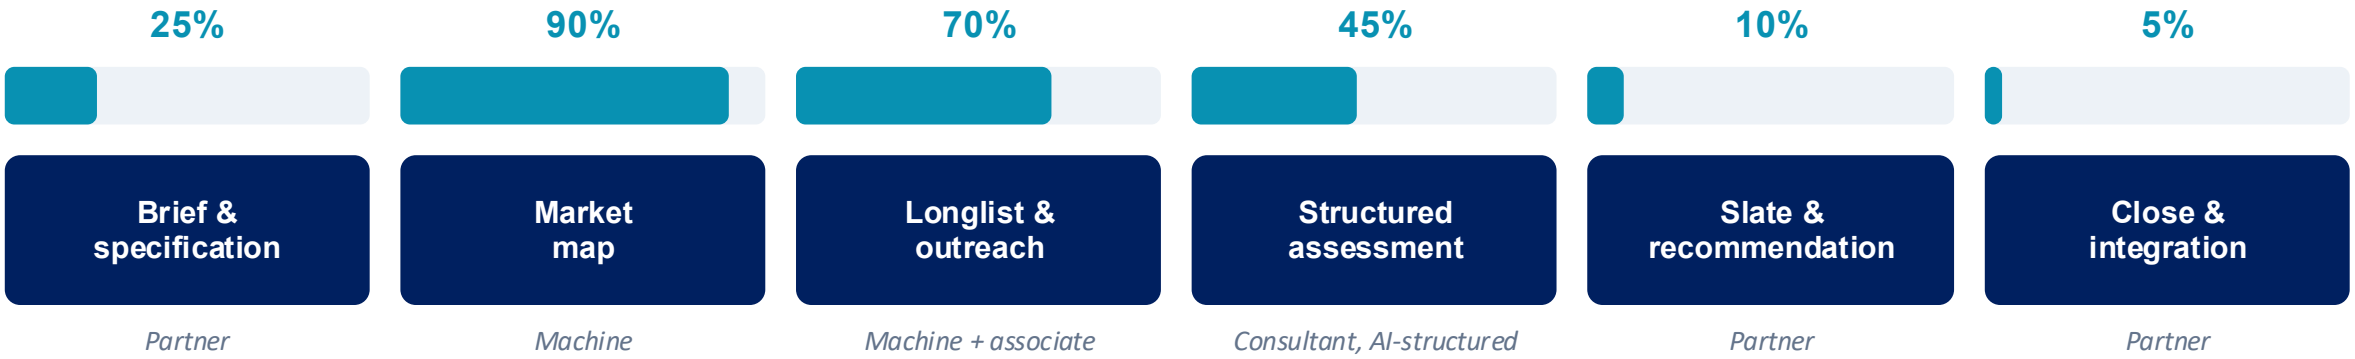

Where speed comes from

Stages two and three. If your competitor takes two weeks out of longlist delivery and you do not, the client notices that long before they notice your judgment.

Where the fee comes from

Stages one, five and six — specification, recommendation, and making the appointment work. These barely move, and they are what the client is actually buying.

ACTION

Ninety days, six moves

VNUHCM Innovation Hub

01 Inventory

Every tool that scores, ranks or filters a person — including features inside platforms you already licence.

02 Baseline five metrics

Time-to-longlist, time-to-shortlist, time-to-fill, offer-acceptance, 18-month retention.

03 Structure the interview

Your competency model, same probes, verbatim evidence. Highest return available today.

04 Kill the score

Replace ranked scores with evidence columns and a named decision-maker.

05 Contract for evidence

Bias testing, documentation and incident notice as vendor obligations, not assurances.

06 Publish your method

Clients will ask how you use AI in their search. A written answer is a differentiator this year.

If you do only one: number two. You cannot prove a return on a process you have never measured.

Thank you

Let the machine find them.
Let the machine structure you.
Do not let it rank them.

Dr. Trinh Quang Viet

Director, Institute of Applied Artificial Intelligence
VNUHCM Innovation Hub · Vietnam National University Ho Chi Minh City