discovr.

TALENT INTELLIGENCE REPORT

Anthropic London Research Engineering Talent Map

A comprehensive analysis of 94 candidates, 42 target companies, and 61 current team members — strategic insights for building a world-class AI safety research team.

94	42	5	50%
CANDIDATES MAPPED	TARGET COMPANIES	OPEN ROLES	PHD RATE

Executive Summary

This report presents a comprehensive analysis of the research engineering talent landscape for Anthropic's London office expansion across five key roles: Alignment Science, Production Post-Training, Pretraining Scaling, Pre-training Core, and Multimodal.

33%

Of candidates from a single source

Google DeepMind dominates the candidate pool with 31 of 94 candidates (33%), reflecting its position as the primary talent source for frontier AI research in London.

61%

Have relocated internationally

Compared to only 30% of the current Anthropic team. This candidate pool is highly mobile and international, reducing relocation friction concerns.

46%

Currently UK-based

43 candidates are already in the UK, critical for the Pretraining Scaling role which requires 5 days in-office and on-call availability.

8.4/10

Average candidate score

Exceptionally strong pipeline with 15 candidates scoring 10/10 and 80% scoring 7 or above.

KEY STRATEGIC INSIGHT

The candidate pool shows a striking contrast to the current Anthropic team: **50% hold PhDs** (vs 46% at Anthropic), but more notably, **84% of candidates have Masters degrees or higher**. The pool is also significantly more internationally mobile — a key advantage for London hiring where visa sponsorship is table stakes.

SECTION 01

Source Company Analysis

Understanding where top research engineering talent currently works is critical for effective sourcing strategy. The candidate pool reveals a highly concentrated talent market.

Top Source Companies

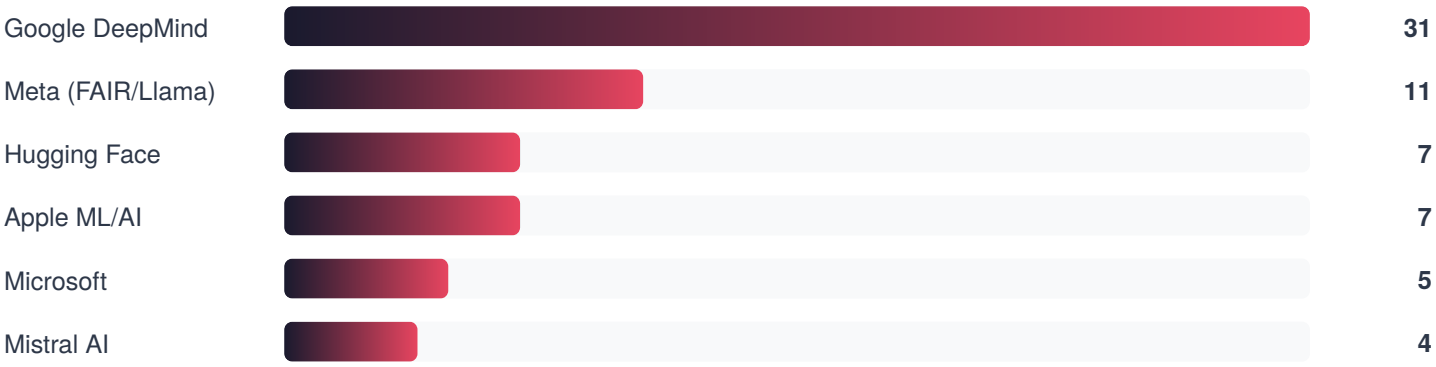

DEEPMIND DOMINANCE

Google DeepMind accounts for **33% of the entire candidate pool**. This creates both opportunity and risk: deep bench for sourcing, but candidates may require significant compensation uplift to move. DeepMind's London HQ makes these candidates particularly valuable for the Pretraining Scaling role.

Gender Diversity by Source Company

COMPANY	TOTAL CANDIDATES	FEMALE	FEMALE %	ASSESSMENT
Google DeepMind	31	9	29.0%	Above Average
Meta	11	2	18.2%	Average
Microsoft	5	3	60.0%	Leader
Apple	7	1	14.3%	Below Average
Hugging Face	7	1	14.3%	Below Average
Mistral AI	4	0	0.0%	Critical

SECTION 02

Target Company Ecosystem

42 companies have been mapped as sourcing targets, categorized by region, type, and relevance to each of the five open positions.

By Region

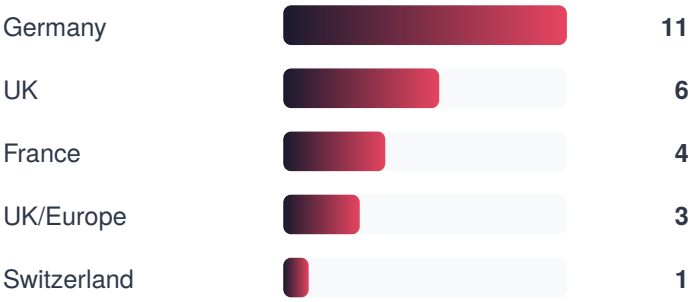

By Primary City

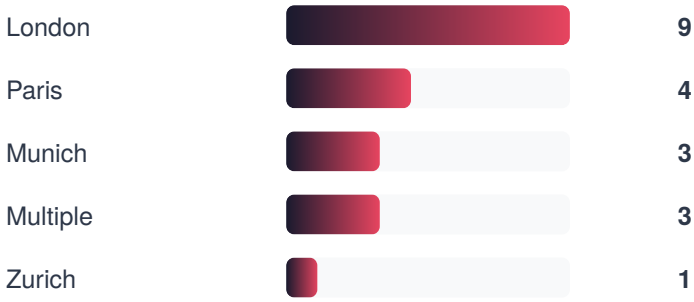

Companies by Role Relevance

ROLE	TARGET COMPANIES	KEY SOURCES
Pretraining Scaling	27	DeepMind, Meta FAIR, OpenAI, Mistral, xAI
Pre-Training Core	22	DeepMind, Meta FAIR, OpenAI, Mistral, Microsoft
Multimodal	20	DeepMind, Meta FAIR, Apple, Black Forest Labs
Post-Training	18	DeepMind, Meta, OpenAI, Cohere, Mistral
Alignment Science	17	DeepMind Safety, OpenAI, UK AI Safety Institute

MULTI-ROLE COMPANIES

Nine companies are relevant across 4-5 roles: **Google DeepMind, Meta AI/FAIR, Microsoft, OpenAI, Anthropic, Google Gemini, xAI, Mistral AI, and Aleph Alpha**. These represent the highest-value sourcing targets as candidates from these organizations are likely qualified for multiple positions.

SECTION 03

Candidate Demographics & Mobility

The candidate pool shows distinctive patterns in geography, mobility, and diversity that inform sourcing and hiring strategy.

Geographic Distribution

Current Country

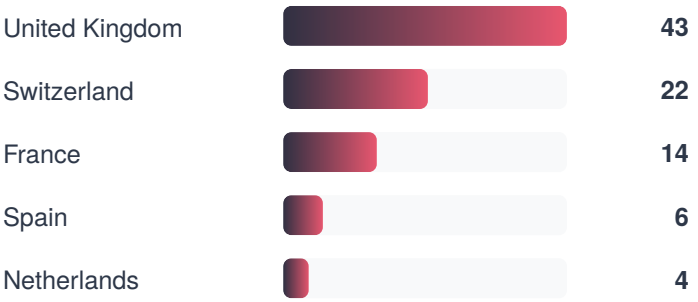

Nationality (Origin)

International Relocation Patterns

57 of 94 candidates (61%) have previously relocated internationally for work, demonstrating high mobility in this talent pool.

FROM	TO	COUNT	IMPLICATION
India	England	5	Well-established pathway
USA	England	4	Reverse migration opportunity
Italy	England	3	Strong EU→UK pipeline
Germany	Switzerland	3	May need extra incentive for UK
India	Switzerland	3	Zurich candidates for London

MOBILITY INSIGHT

The most common destination is **England (47% of relocations)**, followed by Switzerland (37%). Switzerland-based candidates represent a key opportunity — they've already demonstrated willingness to relocate and are in closer geographic/cultural proximity to London than US-based candidates.

SECTION 04

Education & Credentials

Education backgrounds reveal both the depth of expertise in the candidate pool and key feeder institutions for future sourcing.

50%

Hold PhD degrees

47 of 94 candidates have completed doctoral programs, slightly above the Anthropic team rate of 46%.

84%

Masters or higher

Nearly all candidates have advanced degrees, reflecting the research-intensive nature of these roles.

Top Feeder Universities — Candidate Pool

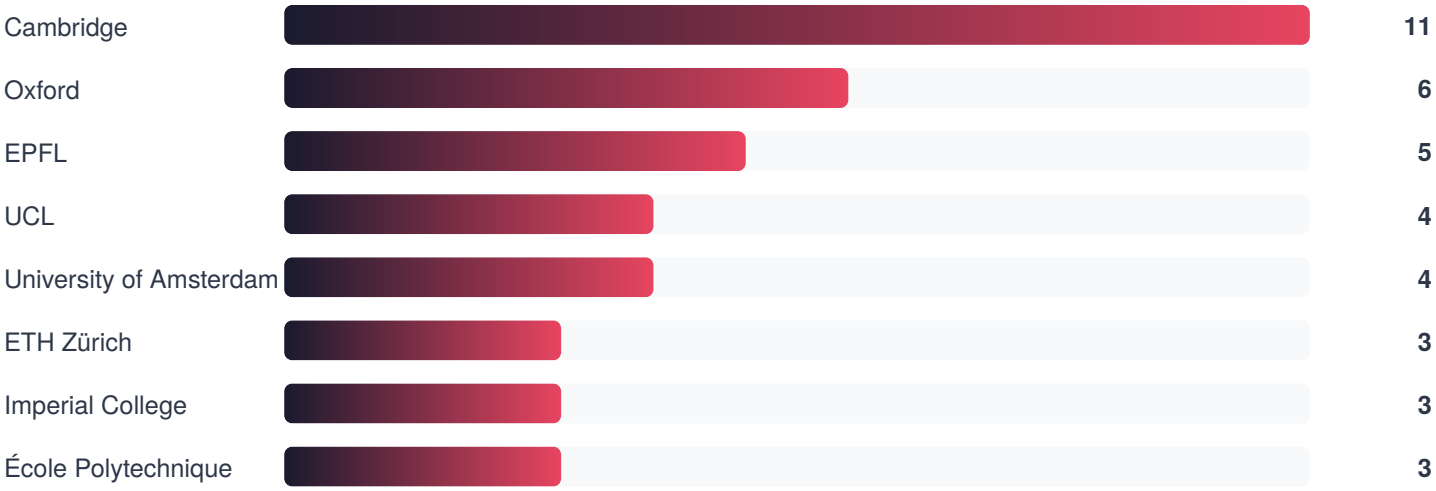

UNIVERSITY OVERLAP ANALYSIS

Only **11 universities appear in both the candidate pool and current Anthropic team**. Notable overlaps: Cambridge, Oxford, Stanford, Berkeley, ETH Zürich, and Imperial. The candidate pool skews heavily European (Cambridge, Oxford, EPFL, ETH) while the current team skews US (Berkeley, Stanford, MIT). This diversification could strengthen research perspectives.

SECTION 05

Candidate Pool vs Current Team

Comparing the candidate pipeline to Anthropic's current research engineering team reveals important differences in background and composition.

METRIC	CANDIDATE POOL	ANTHROPIC CURRENT	DELTA
Total Count	94	61	+33
PhD Rate	50.0%	45.9%	+4.1%
Female Representation	23.4%	28.6%	-5.2%
Prior Relocation	60.6%	29.5%	+31.1%

Previous Employer Comparison

Current Team - Came From

- → Google (including DeepMind) — 9
- → Meta — 5
- → MATS Research — 5
- → OpenAI — 4
- → Amazon — 4
- → Stripe — 4

Candidate Pool - Currently At

- → Google DeepMind — 31
- → Meta — 11
- → Hugging Face — 7
- → Apple — 7
- → Microsoft — 5
- → Mistral AI — 4

DIVERSITY GAP ALERT

The candidate pool has **lower female representation (23.4%)** than the current Anthropic team (28.6%). To maintain or improve diversity, targeted sourcing from companies like Microsoft (60% female in pool) and DeepMind (29% female) should be prioritized. The Mistral AI pipeline (0% female) should be supplemented with alternative sources.

SECTION 06

Seniority & Experience Analysis

Understanding the experience distribution helps calibrate expectations and identify optimal matches for each role's requirements.

Seniority Distribution

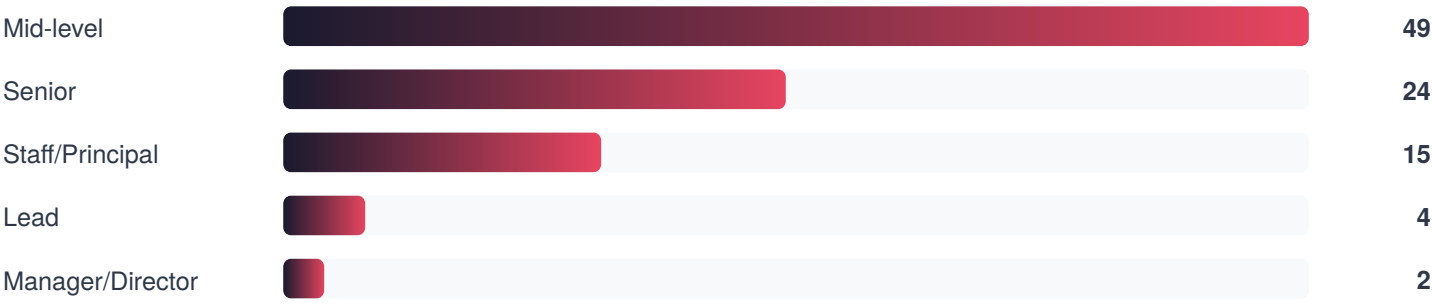

Role Type Distribution

ROLE TYPE	COUNT	PERCENTAGE
Research Engineer	48	51%
Research Scientist	28	30%
Other (ML Engineer, etc.)	11	12%
Researcher	6	6%

SENIORITY & ROLE ANALYSIS

The candidate pool is predominantly **mid-level (52%)**, which aligns well with Anthropic's growth stage and salary bands. However, the **15 Staff/Principal candidates** represent a valuable senior cohort for technical leadership. DeepMind contributes disproportionately to senior talent, with 9 Staff/Principal engineers in the pool. The 51% Research Engineer composition matches the target roles well, while the 30% Research Scientists offer flexibility for research-heavy positions like Alignment Science.

Tenure in Current Role

28.2

Mean tenure (months)

Average time in current role is just over 2 years, indicating a mobile talent pool.

15

Median tenure (months)

Half of candidates have been in role for 15 months or less — high-mobility market.

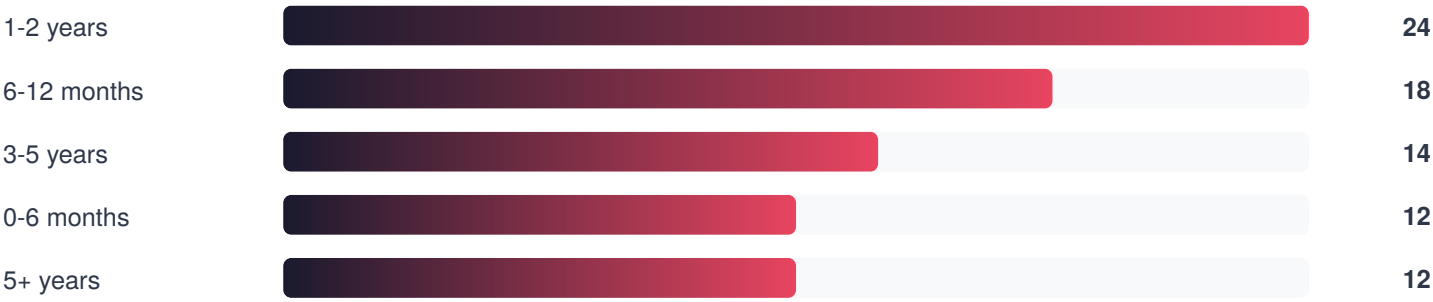

TIMING CONSIDERATION

32% of candidates (30 people) have been in their current role for **less than 12 months**. These candidates may have unvested equity or recent commitment that makes them harder to move. Conversely, candidates in the **1-3 year range (28 candidates)** represent the optimal timing window — established enough to demonstrate value, but potentially ready for new challenges.

SECTION 07

Role Coverage Analysis

The candidate pool provides varying depth of coverage for each of the five target positions.

ROLE	SALARY RANGE	LOCATION	COVERAGE	NOTES
Alignment Science	£250-270k	Hybrid (25% office)	Strong	Deep bench from DeepMind Safety, former Anthropic employees, academic safety researchers
Post-Training	£270-340k	Hybrid (25% office)	Strong	Multiple candidates with direct Gemini/Llama RLHF experience
Pretraining Scaling	£250-435k	5 days in-office	Strong	43 UK-based candidates critical for on-call requirements
Pre-Training Core	£250-270k	Hybrid (25% office)	Moderate	Good architecture/algorithm expertise, fewer at Staff+ level
Multimodal	£250-270k	Hybrid (25% office)	Limited	Only 2-3 candidates with direct multimodal LLM experience; audio/video generation backgrounds available

COVERAGE GAP: MULTIMODAL

The Multimodal role has the thinnest coverage. While there are candidates from audio generation (ElevenLabs alumni), video generation (Synthesia), and image generation (Black Forest Labs), **true multimodal LLM expertise is rare**. Consider expanding sourcing to: DeepMind Lyria/Gemini Vision, Meta's multimodal FAIR team, Apple's Foundation Models group, and academic vision-language researchers.

SECTION 08

Strategic Recommendations

Based on the data analysis, here are actionable recommendations for the Anthropic London hiring initiative.

1. Prioritize UK-Based Candidates for Pretraining Scaling

With the Pretraining Scaling role requiring 5 days in-office and on-call availability, the 43 UK-based candidates should be prioritized. Relocation from Zurich (22 candidates) is the next-best option given geographic proximity and demonstrated mobility patterns.

2. Address the Diversity Gap Proactively

The candidate pool's 23.4% female representation trails the current team's 28.6%. To maintain diversity standards:

→ Prioritize Microsoft and DeepMind sources (highest female representation)

→ Supplement Mistral sourcing (0% female) with alternative targets

→ Consider academic channels where gender balance may be higher at junior levels

3. Expand Multimodal Sourcing

With only 2-3 strong multimodal fits, expand targeting to:

→ DeepMind's Lyria and Gemini Vision teams

→ Meta FAIR's vision-language research group

→ Apple's multimodal Foundation Models team

→ Academic labs (Cambridge, Oxford ML groups doing vision-language)

4. Optimize for the 1-3 Year Tenure Window

28 candidates are in the optimal 1-3 year tenure window — enough time to prove value, but potentially ready for new challenges. Candidates in the 0-12 month window may face unvested equity constraints.

5. Leverage Staff+ Candidates Strategically

15 Staff/Principal-level candidates represent senior leadership potential. DeepMind's Staff Research Engineers and Meta's Staff positions should be highest priority for the £270-340k Post-Training role.

> **PROCESS SPEED MATTERS**
>
> With 45% of ML engineers changing roles within 12 months (per market intelligence), **hiring speed is a competitive differentiator**. Top candidates receive multiple offers within 2-3 weeks. Target total process time of under 3 weeks from first contact to offer.

Methodology

Data Sources: LinkedIn profile data, job descriptions for 5 target roles, company mapping across 42 organizations

Sample Size: 94 candidates (target), 61 current Anthropic engineers

Analysis Date: December 2025

Note: All analysis is anonymized. No individual names are included to protect candidate privacy.

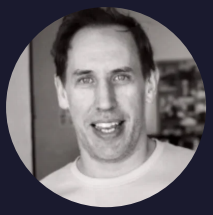

Douglas Gambling

Founder, Discovr

I help scaling AI companies build world-class ML teams through data-driven talent strategy and executive search. Previously Head of Talent at multiple high-growth tech companies, with experience across Europe, the US, and Asia.

LinkedIn | Book a Call | Email