pave + nua group

2025

AI & ML Compensation Trends & Practices Report

© 2025 Pave

Table of Contents

01 Research Methodology PG. 3

02 Introduction to AI & Machine Learning PG. 8

03 Pay Levels & Premiums PG. 16

04 Digging Deeper: Attrition, Sign-on Bonuses & Vesting PG. 23

05 Offer & Hiring Trends PG. 30

06 Bringing AI & ML into Your Job Architecture with Nua Group PG. 36

01

Research Methodology



METHODOLOGY

To create this report, we utilized Pave's real-time compensation dataset to analyze pay for 4,503 employees mapped to Artificial Intelligence (AI) and Machine Learning (ML) jobs and 169,000 employees mapped to software engineering jobs. Overall, 73% of these employees works at private companies and 27% work at public companies.

All analyses presented in this report include a minimum sample size of 100 companies, and all charts reflect data for US-based employees, unless otherwise noted (e.g., when examining locations for AI and ML engineers).

Pave's real-time dataset is maintained through automated and persistent connections to HRIS, ATS and equity management systems (EMS). Connecting to EMS platforms allows us to capture information on individual equity awards, including the intended, gross, and net values of awards, vesting schedules, equity vehicles and other equity practices.

Company firmographic data (e.g., capital raised, company size, private vs public, etc.) comes directly from customers or is enriched from Pitchbook.

For charts that show prevalence across companies, we have filtered the company set to companies within a precursor category (e.g., companies containing at least one software engineer) to capture true prevalence by excluding companies with missing job family classifications. For details on the job levels referenced in this report, visit our **Job Levels support page.**

INTERPRETING OUR DATA

This report is intended to provide directional information on the state of compensation for Artificial Intelligence (AI) and Machine Learning (ML) engineers. Given the volume of media coverage discussing the actual or estimated size of compensation packages for a select number of AI/ML professionals, you may be surprised by some of the numbers in this report. Frankly, the numbers may be lower than you expect.

As experienced compensation professionals know, this mismatch between perception and real market data, especially when dealing with "hot skills" is commonplace. Median pay levels reflect the midpoint of market practices and will not match widely reported outliers at the extreme ends of the market. For this reason, we encourage readers of this document to focus not only on the actual pay levels we report, but the premiums we observe between AI/ML engineers and general software engineers. These premiums, especially for equity compensation, are material and reflect an important point of differentiation in how AI/ML engineers are paid.

Furthermore, as we note in the job architecture section of this report, companies are still in the early days of formalizing job titles, job descriptions, job families, and career paths for AI/ML engineers. This means the structure of the data we receive from clients on their AI/ML talent will continue to improve. For example, we expect the specificity of job titles (e.g., AI/ML researchers vs AI/ML application developers) to evolve over time. As companies mature their handling of AI/ML jobs and career paths, a clearer view of the market will emerge.

This said, Pave's real-time approach to collecting and analyzing compensation data, coupled with the capabilities of our data science team, allows us to present an early view of the market for AI/ML talent that traditional compensation survey providers, with slower timelines to collect and publish data on new jobs, often struggle to match.

We will continue to monitor compensation trends and practices for AI/ML engineers and publish updated reports to keep the total rewards community as informed as possible as data on this subject gets richer and richer.

ABOUT PAVE'S DATASET

Real-time market data from Pave

Pave's real-time dataset includes information from more than 8,500 customers who have connected their HRIS, ATS, and equity management systems to our platform. Currently, our dataset skews heavily toward the technology sector, especially software companies.

Offer Insights data

Pave has partnered with a leading ATS provider to bring real-time offer and hiring insights to our customers. Currently, Pave ingests de-identified data from more than 7,500 companies and performs analytics on this data to incorporate high-level insights into the Pave platform.

GENERAL DISCLAIMER

We've taken care in preparing this document. However, please note we can't guarantee that the information contained herein will fit the specific needs for your organization. Pave makes no representations or warranties on this document, so please keep that in mind when relying on the information contained herein. Please also note that this document and the information contained herein are Pave's proprietary information.

SELECT PARTICIPANTS IN PAVE'S REAL-TIME DATASET

PUBLIC

Affirm	AppFolio	Asana	Atlassian
Aurora	Bill	Block	Braze
Cloudflare	Coinbase	Confluent	CSL Behring
Doordash	Dropbox	Elastic	Electronic Arts
FanDuel	Five9	GitLab	Grubhub
HashiCorp	Hims & Hers	HubSpot	Instacart
Joby	Klaviyo	Lucid Motors	MongoDB
Nutanix	Okta	On Holding AG	Procore
Remitly	RingCentral	Rivian	Roblox
Samsara	Silicon Labs	Smartsheet	Snowflake
Sofi	Squarespace	Sweetgreen	The RealReal
Twilio	Unity Technologies	Zip Co	Zoom
Zscaler	Zuora		

PRIVATE

Anduril Industries	Auctane	Automation Anywhere	BAL
BambooHR	BitSight Technologies, Inc.	Boomi	Clio
Commonwealth Fusion Systems	Credit Karma	Databricks	Dialpad
DocPlanner	Emburse	Extreme Reach	fivestars
GLOBO	GOAT	Gusto	Industrious
Mews	Monzo	Motive Technologies, Inc.	Navan
Notion	OpenAI	OutSystems	Pax8
PayJoy	Qualtrics	Relativity Space	Remote
Rohlik Group a.s.	Scale AI	ServiceTitan	Solidcore
Sonder	SpotOn	Stripe	SymphonyAI
Tarro	Tekion	Tide	Undefeated Tribe Operating Co. LLC
Verkada	Virtual Business Partners	Wealthsimple	X
Zinnia	Zipline		

02

Introduction to AI & Machine Learning





A New World

2024 was a big year for the technology sector. And no trend was buzzier than the advent of everyday use cases for Artificial Intelligence (AI) and Machine Learning (ML). Naturally, this trend also had a significant impact on AI/ML hiring and pay.

As companies race to include AI/ML capabilities in their products and business processes, demand for skilled AI/ML engineers is booming. As a result, organizations are struggling to keep pace with compensation trends and practices across the market.

In this report, we dive into findings from Pave's real-time dataset to explore compensation levels, pay practices, and hiring trends for AI/ML roles. We also partner with Nua Group to examine how organizations can incorporate these findings into their go-forward job architecture and compensation strategy.

But first, let's get oriented on the differences between AI and ML roles, and how they differ from general software engineers (SWE).

What is the difference between AI & ML talent?

AI and ML Engineers have different skill sets than Software Engineers. But what do they do exactly, and what are the differences between AI and ML Engineers?



AI Engineers are the people building complete artificial intelligence systems that can mimic human intelligence and decision-making (think Claude or ChatGPT, for instance). They often work on complex problems like natural language processing, and focus on the theoretical aspects and architecture of AI systems.

There is another job title within this category worth noting: **AI Researcher**, sometimes also called an AI or ML Research Scientist. In general, these people represent the extreme outlier engineers or researchers building new types of models and furthering the future of AI. They typically have PhDs, or in some rare cases, have published research without a PhD that has moved the industry forward. Because of their unique skill set, folks in this 1% demand a very high premium when it comes to compensation.



ML Engineers adopt AI and ML best practices and research to solve business problems. They specialize in developing systems that can learn from and improve with data. They typically focus more on statistical modeling, data processing, and optimization algorithms, and are concerned with data pipeline engineering and model deployment. These tend to be the roles that the bulk of companies outside the AI field are looking to hire.

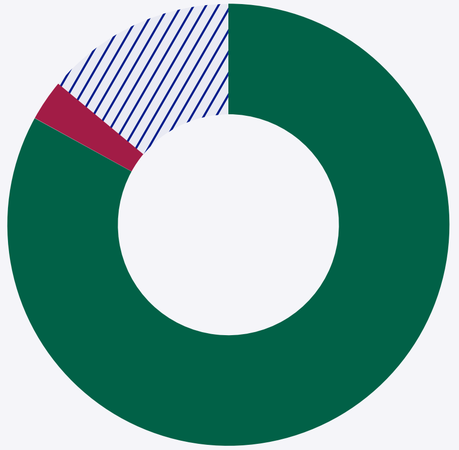
All this said, the terms AI and ML are often used interchangeably when it comes to job titles, as machine learning is used to build AI tools. A Machine Learning Engineer at one company might be considered an AI Engineer at another company. Currently, Pave's Market Data product puts AI and ML engineers together in our "ML Engineering" job. We expect this approach to evolve as companies get better at differentiating AI and ML roles.

While the terms AI and ML may be used interchangeably, the data shows a clear trend. “ML Engineer” is the most common job title for these roles in Pave’s dataset.

The majority of AI/ML Engineers have “ML” or “Machine Learning” in their title, and only 3% have “AI” or “Artificial Intelligence.” Other common titles include “Applied Scientist,” “Computer Vision Engineer,” “Algorithm Engineer,” and “Research Scientist.”

Further, most AI/ML roles are not specialized when it comes to their titles. Only 17% include a specialization such as “Research,” “Applied,” or “Infrastructure.” The majority of AI/ML Engineers in Pave’s dataset are generalists, at least in title.

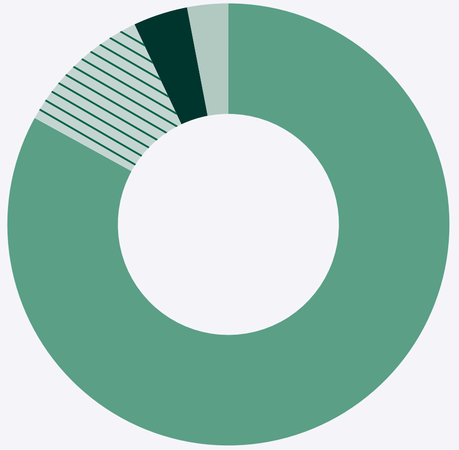
Percentage of roles containing “ML” or “AI” in their job title



- “ML” or “Machine Learning” **(83%)**
- “AI” or “Artificial Intelligence” **(3%)**
- Other **(14%)**

Sample size: 3.9K employees mapped to the ML job family

Percentage of AI/ML employees by job type



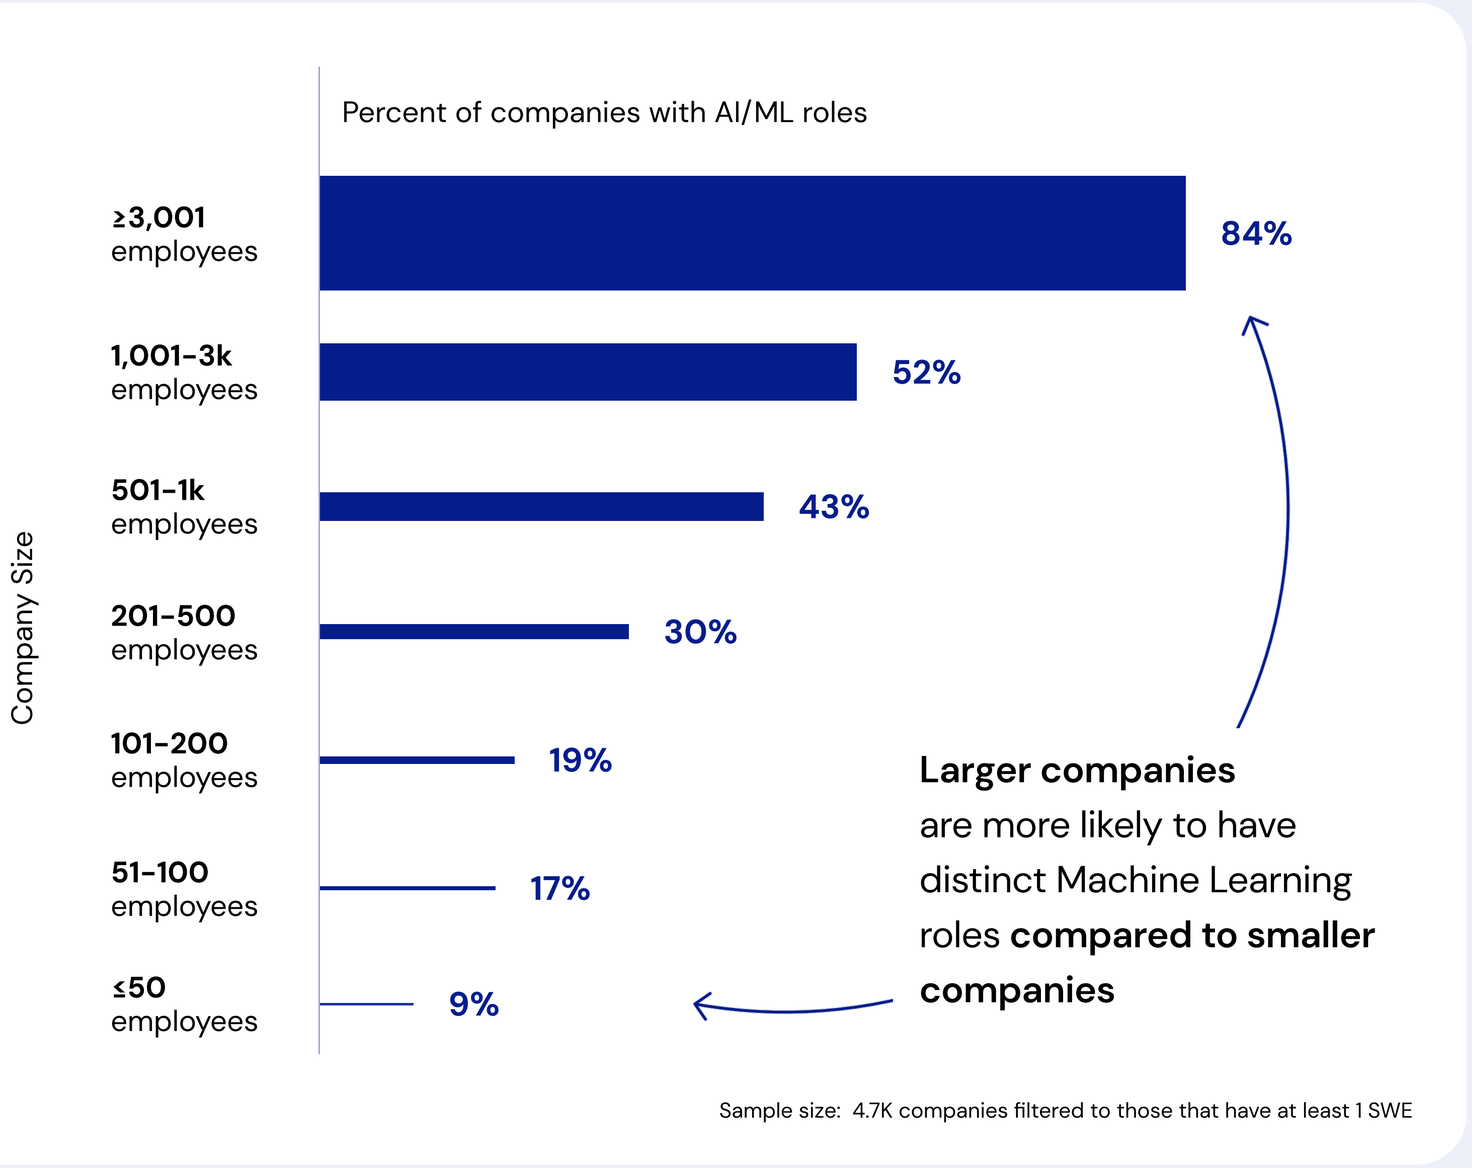
- Generalist **(83%)**
- Applied **(10%)**
  (e.g. applied, NLP, deep learning, computer vision, algorithm, perception)
- Research **(4%)**
  (e.g. research, ethics, statistics)
- Infrastructure **(3%)**
  (e.g. platform, ops/operations, infra/infrastructure)

Sample size: 3.9K ML employees mapped to ML job categories of research, infrastructure, applied, and generalist

Who is Hiring AI & ML Engineers?

Larger companies are more likely to hire for AI/ML roles — 84% of companies with 3,001+ employees have AI/ML talent on staff, compared to 9% at companies with 50 or fewer employees.

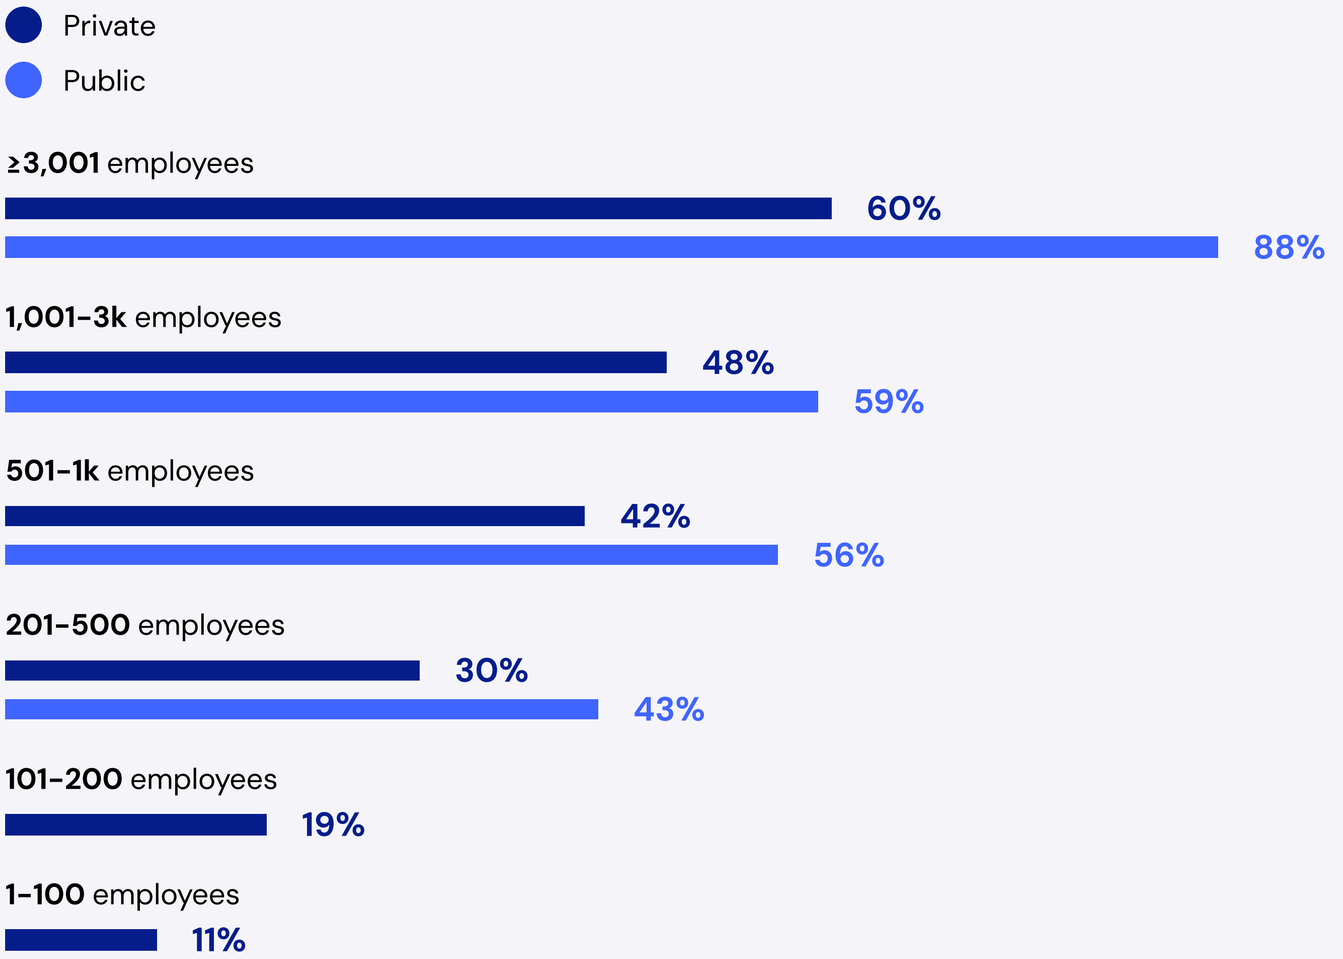
It's important to note that this analysis only includes companies with AI/ML talent who also have a distinct job family or title for those roles. This speaks to the maturity of larger organizations around managing compensation and job architecture. It is likely the case that more companies are leveraging AI/ML talent within their business, but these tasks fall to SWEs whose titles don't reflect their specific skill set.



Sample size: 4.7K companies filtered to those that have at least 1 SWE

When comparing public companies to private companies, the data shows public companies are more likely to have distinct AI/ML job titles than private companies, especially the larger they are.

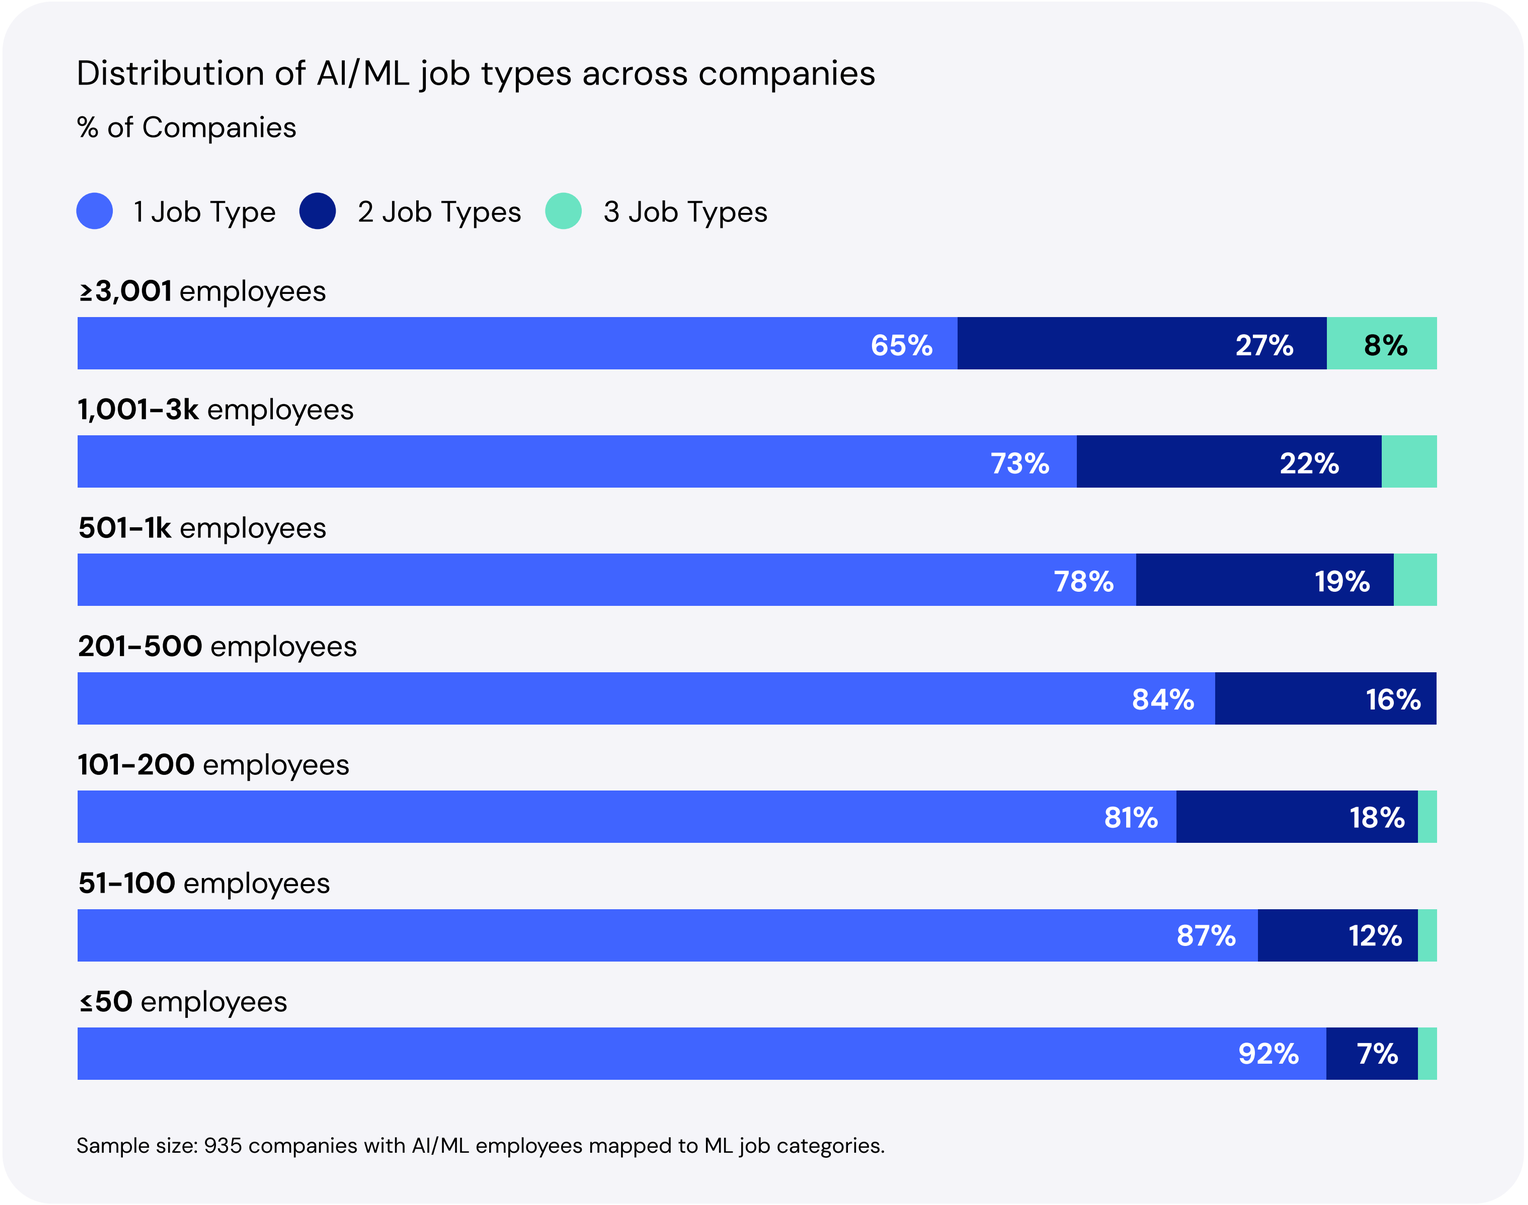
Prevalence of distinct AI/ML roles among public and private companies



Sample size: 4.7K companies.

Among companies with distinct AI/ML titles or job families, we can go even deeper and look at job types.

35% of businesses with 3001+ employees have 2 or 3 job types within their AI/ML community (e.g., ML Research Scientist, ML Infrastructure Engineer, and ML Engineer), which again speaks to the maturity of job architecture at larger companies. The number of AI/ML job types is far lower at small companies, who may not yet need that level of granularity within their job architecture.



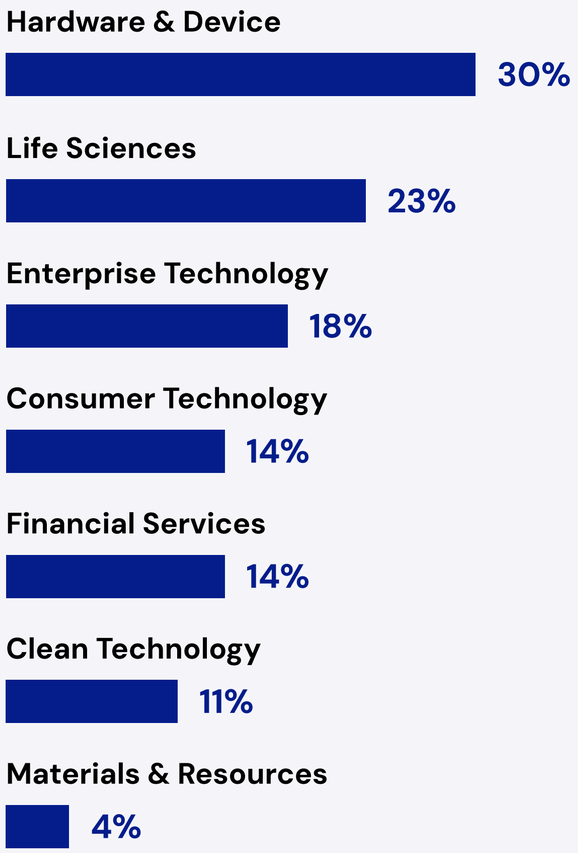
Sample size: 935 companies with AI/ML employees mapped to ML job categories.

Let's zoom in on the makeup of the AI/ML talent pool. In what industries do these employees work, and what levels are most common?

Looking at industries, our data show Hardware, Life Sciences, and Enterprise Technology companies are most likely to have AI/ML talent. Those first two results might come as a surprise, as AI/ML engineering is typically perceived as a software-dominated field. However, this data means SaaS companies are competing for AI/ML talent across industries, and those industries may also be more advanced in managing AI/ML talent in their job architectures.

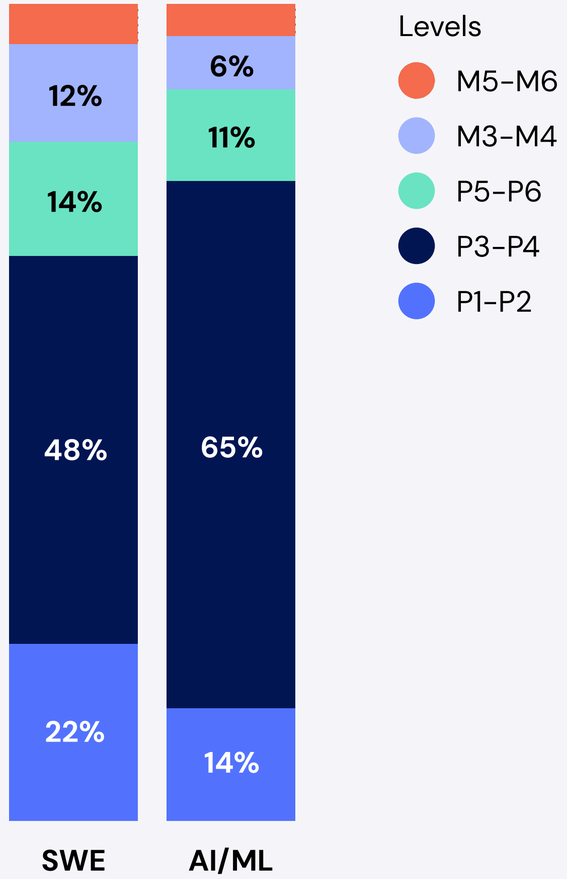
In terms of job level, Pave data shows the overwhelming majority of AI/ML engineers are at levels P3 and P4, which demonstrates that there are far fewer entry-level AI/ML Engineers than SWE generalists. Meanwhile, AI/ML managers at the M5 and M6 level are least represented. This may indicate that AI/ML engineers report to SWE managers rather than directors who specialize in this area.

Percent of companies with AI/ML talent by industry



Sample size: 3.9K companies.

Distribution of SWE and AI/ML employees by job level



Sample size: 118K employees, filtered to companies with more than 100 employees.

03

Pay Levels & Premiums





The Premium for AI & ML Talent

With demand for AI/ML talent so high, it naturally follows that compensation for these roles is also high relative to their software engineering counterparts.

The hype around AI has fueled dramatic trends in compensation. The top 1% of AI/ML engineers, and researchers in particular, are being recruited by very well-known companies and are being compensated at extreme multiples of a typical engineer's salary. You've likely heard anecdotes of tech leaders offering AI/ML researchers $1M+ compensation packages.

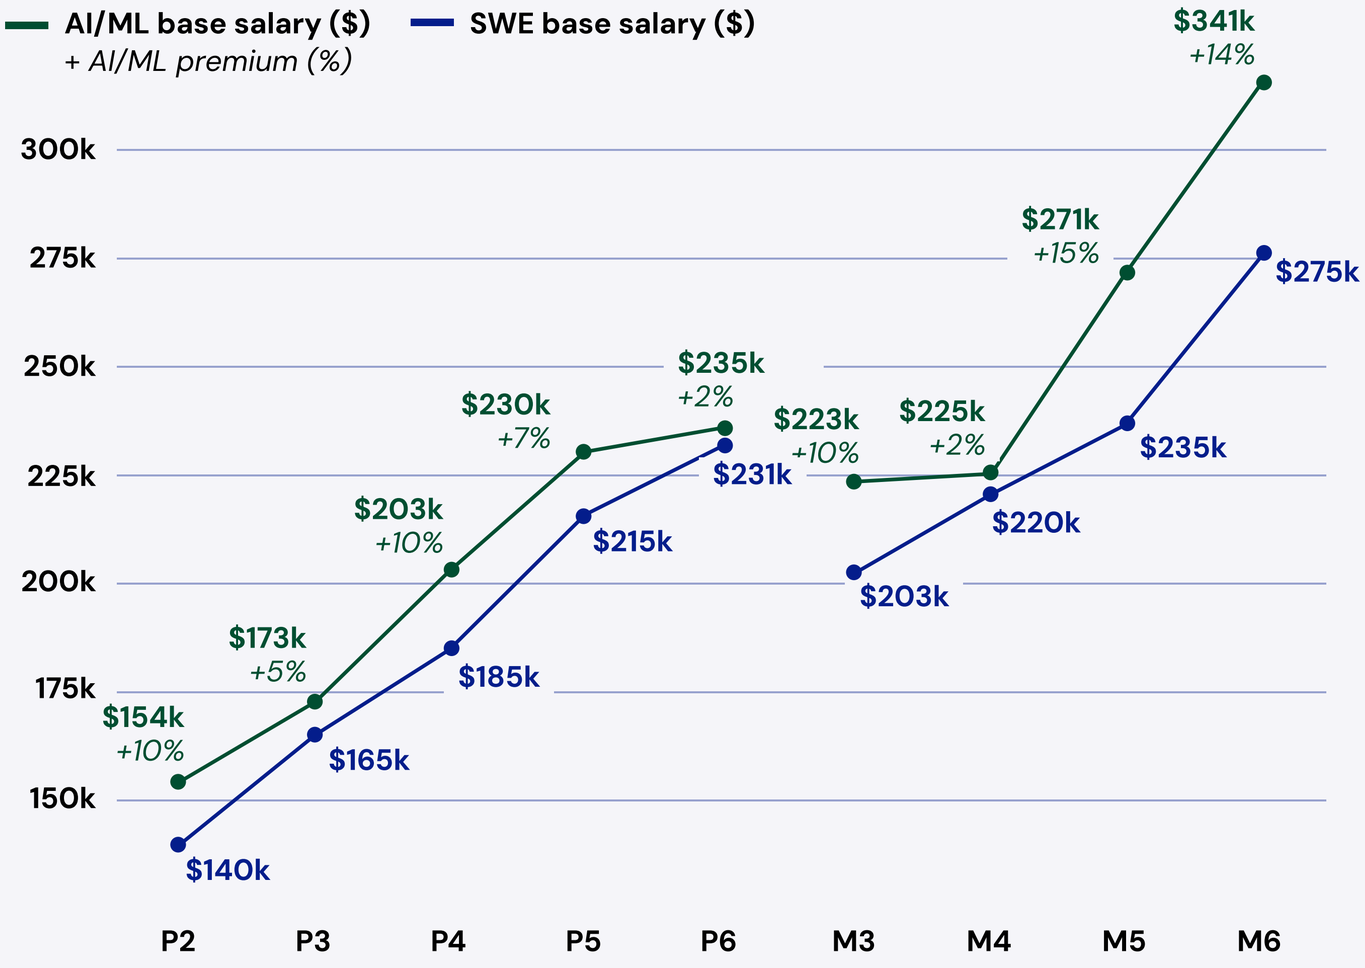
While those are exceptional cases, which are not reflected in the median pay levels reported here, Pave data still shows a clear premium for AI/ML talent across levels. On the following pages, we compare AI/ML talent to software engineers to provide context, and separate private and public company data to show any discrepancies among these two groups.

AI/ML Engineers enjoy a base salary premium at all levels.

At private companies, the base salary premium for AI/ML roles relative to SWEs is typically 10 to 15%, but can dip to as low as 2% at some job levels (e.g., P6 and M4). The highest premiums are found at the M5 and M6 levels, which are ~15% and ~14%, respectively.

While these benchmarks are for all private companies, organizations should filter to specific company stage and/or geographic peer groups when conducting their own pricing exercise.

Median base salary by level at private companies (in USD)



Sample size: 1.7K AI/ML employees, 51K SWE employees, based in the US.

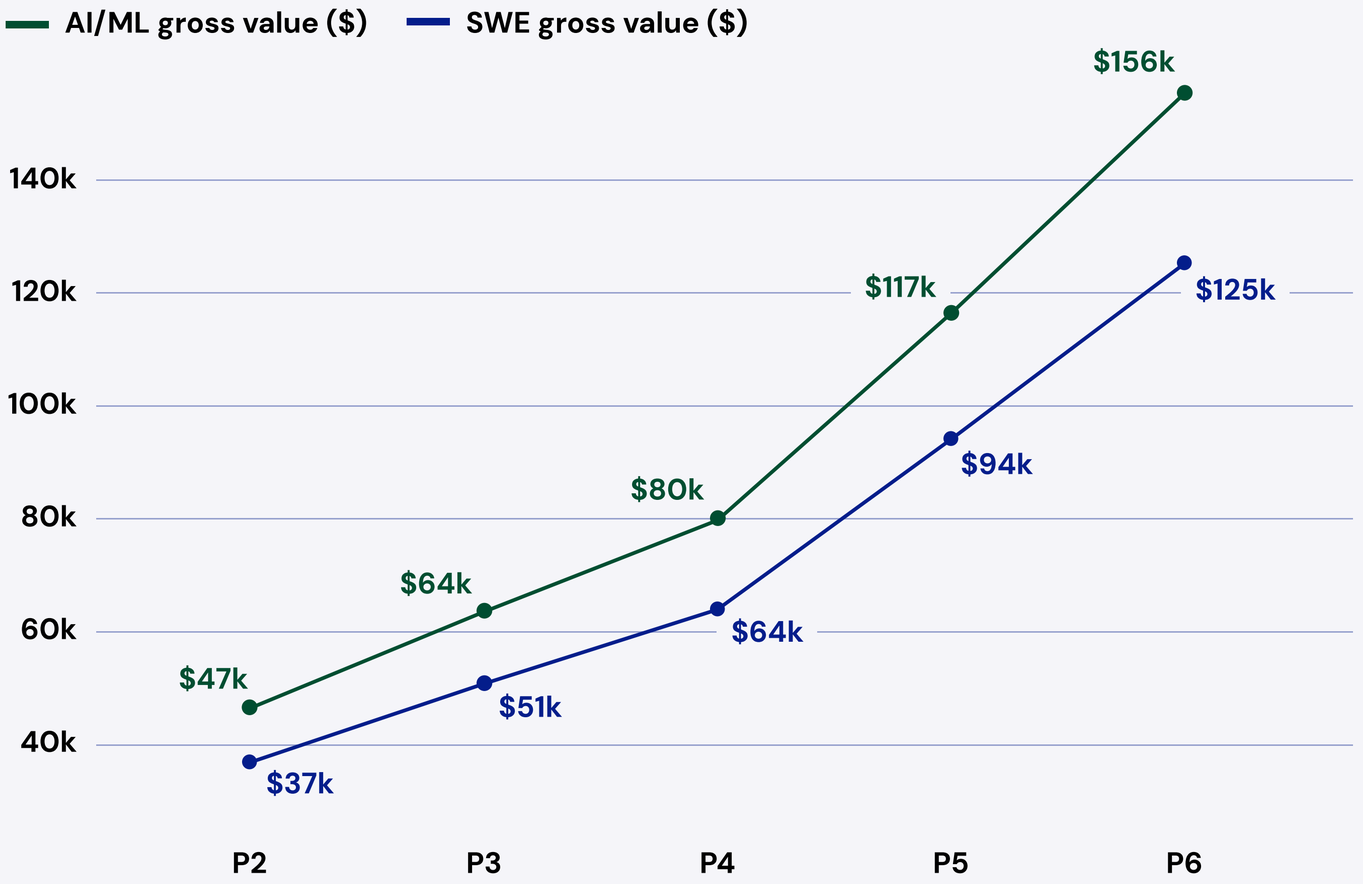
PRIVATE COMPANY BENCHMARKS

Pay premiums are most pronounced when looking at new hire equity awards.

On an annualized basis, private companies are providing more value to AI/ML engineers through new hire grants. This indicates firms are leaning on equity to give them a competitive edge when hiring key talent.

Of note, total new hire equity grant values are typically 3x to 4x larger than the values reported here. We show annualized grant values to normalize benchmarks across varied vesting schedules.

Median new hire equity grant by level at private companies (annual intended gross value in USD)



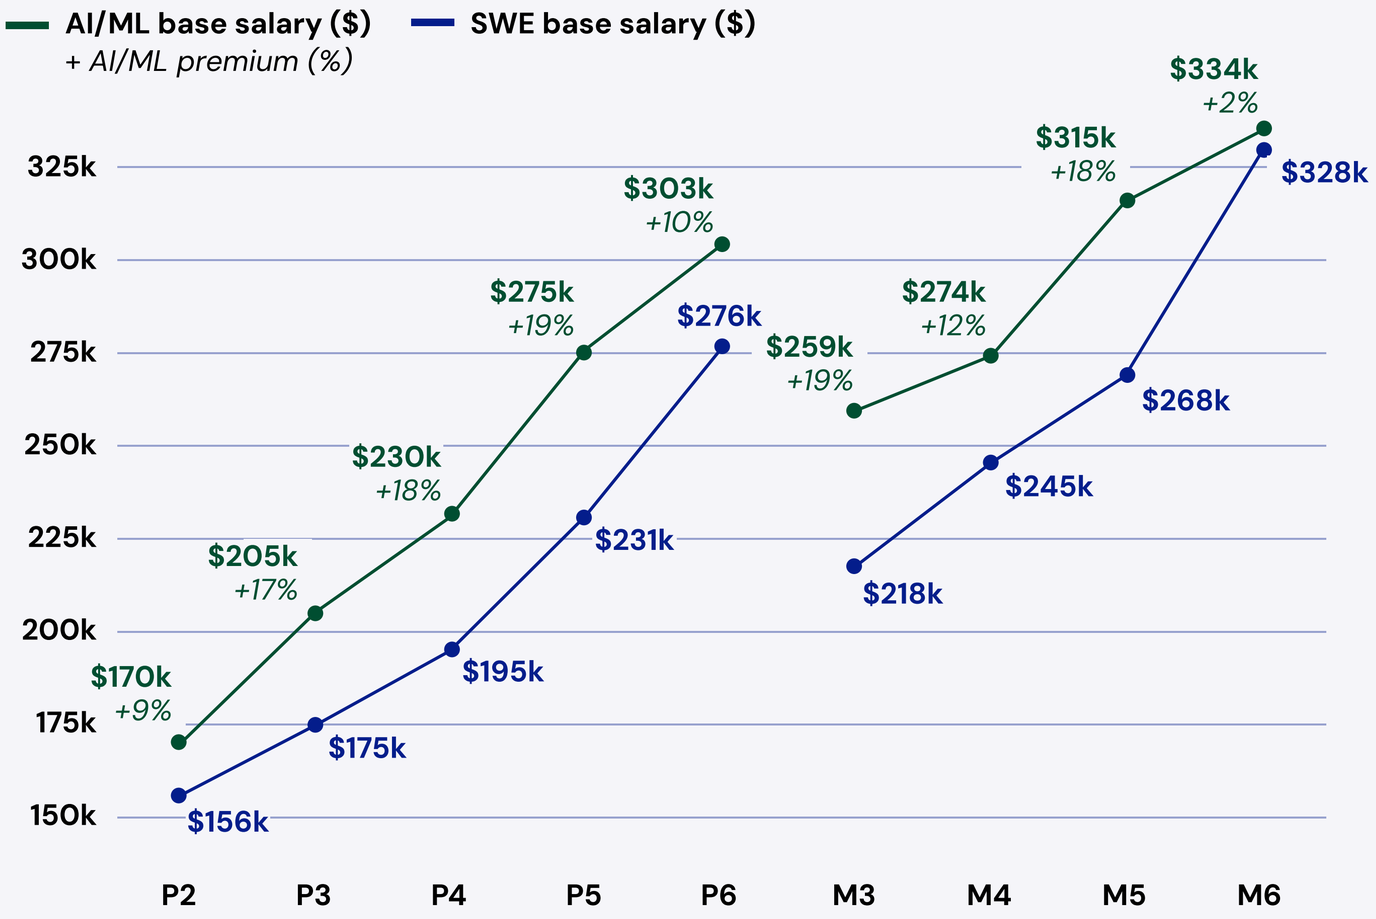
Sample size: 216 AI/ML employees, 4.6K SWE employees, based in the US. These results leverage Pave's Calculated Benchmarks.

PUBLIC COMPANY BENCHMARKS

Public companies follow a similar pay premium path.

In fact, base salary premiums for AI/ML talent tend to be even more dramatic at public companies than at private firms. The highest premiums for AI/ML engineers at public companies are at the P5 and M3 levels and approach 20%. We see the lowest premium at the M6 level, where Directors/Sr Directors in AI/ML make only 2% more than their SWE counterparts in base salary.

Median base salary by level at public companies (in USD)



Sample size: 646 AI/ML employees, 21.5K SWE employees, based in the US.

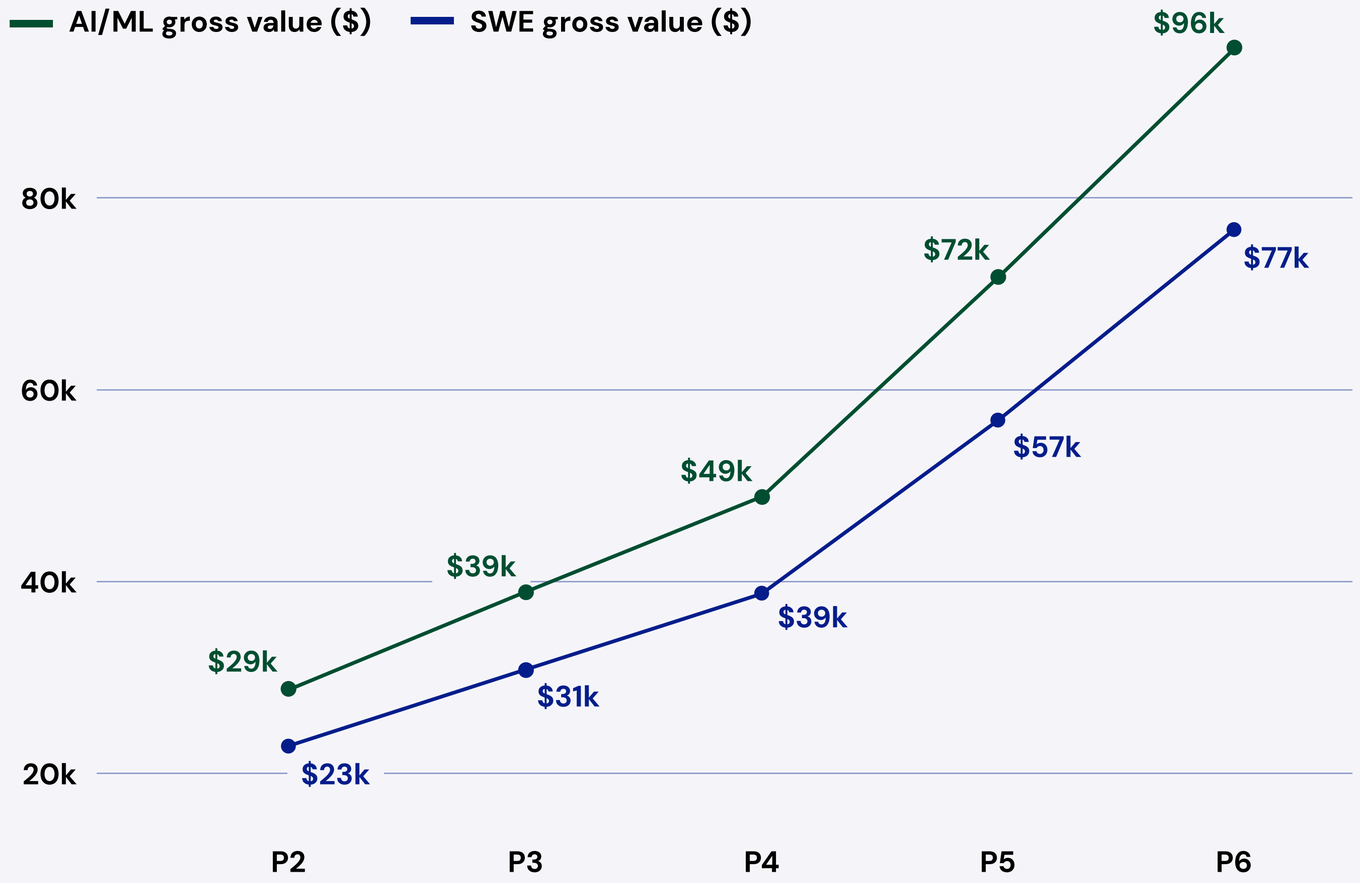
PUBLIC COMPANY BENCHMARKS

Public companies also use new hire equity as an incentive in a competitive talent market.

Once again, we see a clear premium for AI/ML talent when looking at annualized new hire equity grant values. Clearly, competition is fierce, and public companies are also pulling the equity lever to attract top talent.

Again, total new hire equity grant values are typically 3x to 4x larger than the values reported here. We show annualized grant values to normalize benchmarks across varied vesting schedules.

Median new hire equity grant by level at public companies (annual intended gross value in USD)



Sample size: 173 AI/ML employees, 3.6K SWE employees, based in the US. These results leverage Pave's Calculated Benchmarks.

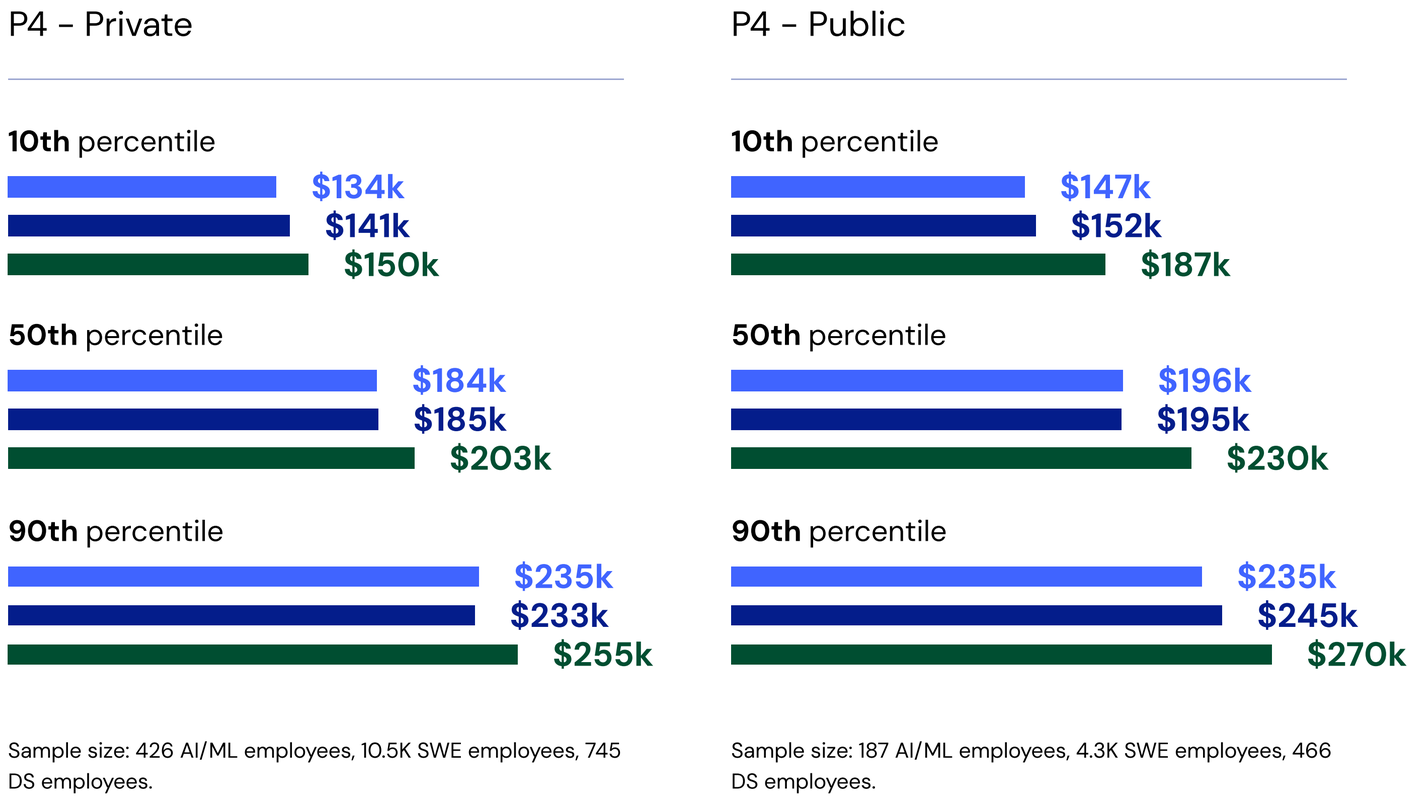
At All Levels of Pay, How Expensive is AI/ML Talent?

These benchmarks show what employers at both private and public companies are paying their AI/ML engineers and SWEs in base salary at the 10th, 50th, and 90th percentiles. In this and future sections, we've also included information on Data Scientists, as this role is closely related to AI/ML and can provide additional context. Regardless of where companies target pay, premiums for AI/ML talent exist.

Again, while this data is illustrative of broader trends, it's important for organizations to specify company stage and/or geographic peer groups when doing their own benchmarking and pricing exercises.

Base salaries at the bottom, middle, and top of the market for the P4 level (in USD)

- Data science
- Software engineering
- AI/ML



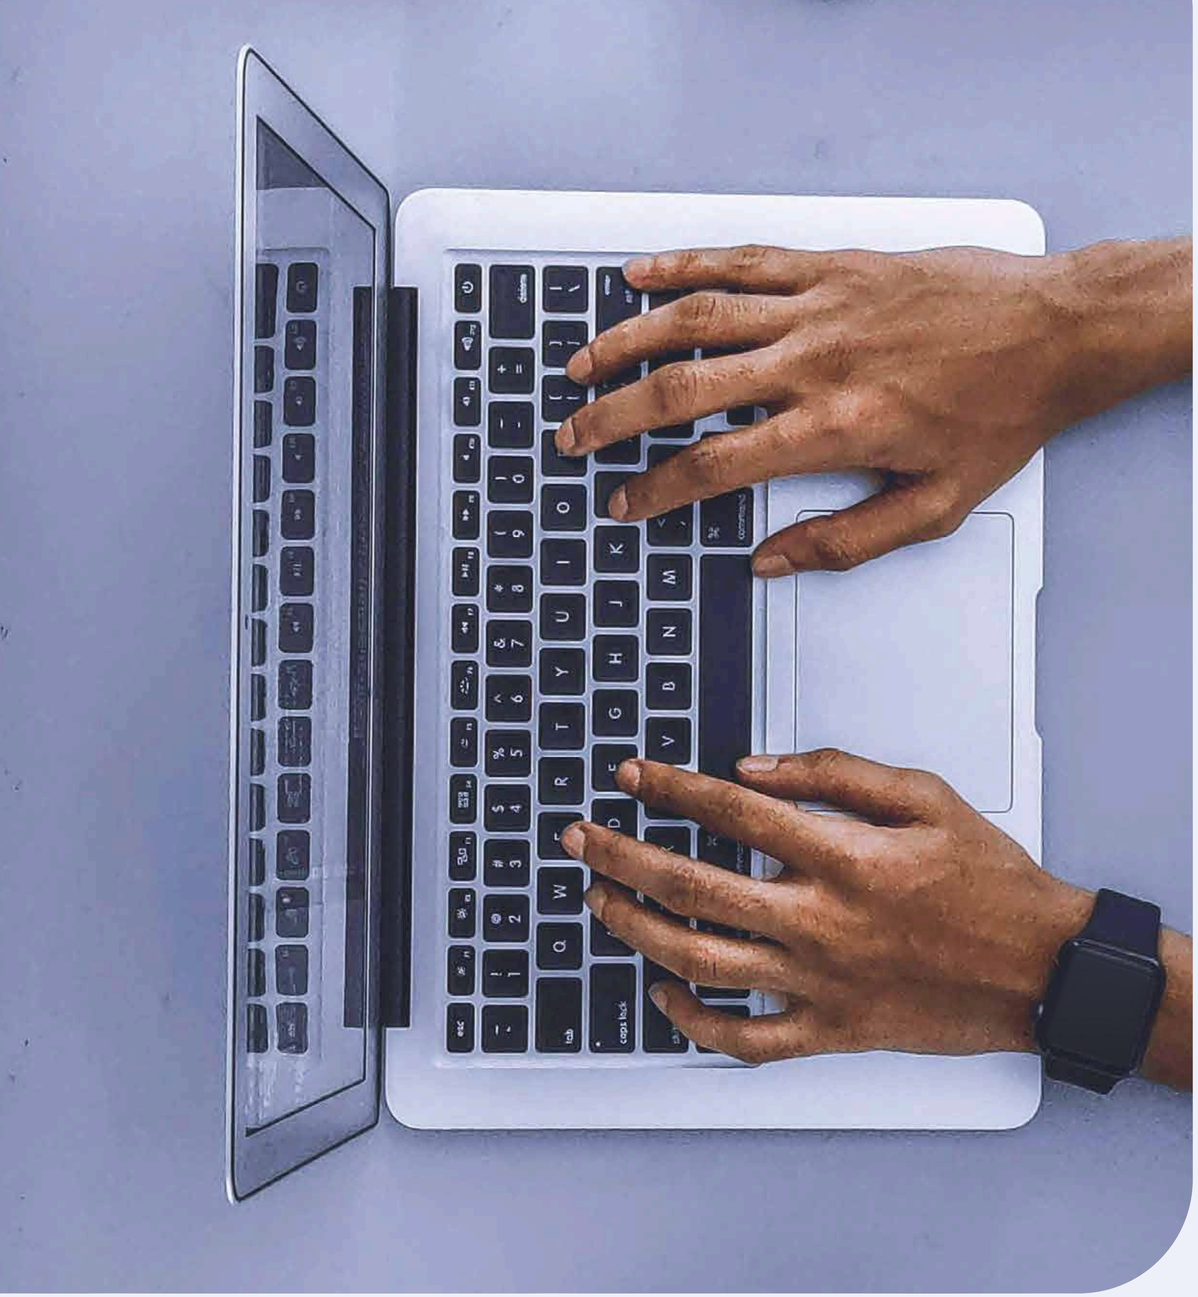
Sample size: 426 AI/ML employees, 10.5K SWE employees, 745 DS employees.

Sample size: 187 AI/ML employees, 4.3K SWE employees, 466 DS employees.

04

Digging Deeper: Attrition, Sign-on Bonuses & Vesting





Digging Deeper

Going beyond base salary and equity is key, as there are other levers companies can pull to help them attract and, crucially, retain AI/ML talent.

Sign-on bonuses are becoming an increasingly popular tool in an organization's recruiting toolkit to get AI/ML talent on board. Additionally, equity grant structure is another method that companies can leverage, both for new hire and ongoing (or refresh) grants.

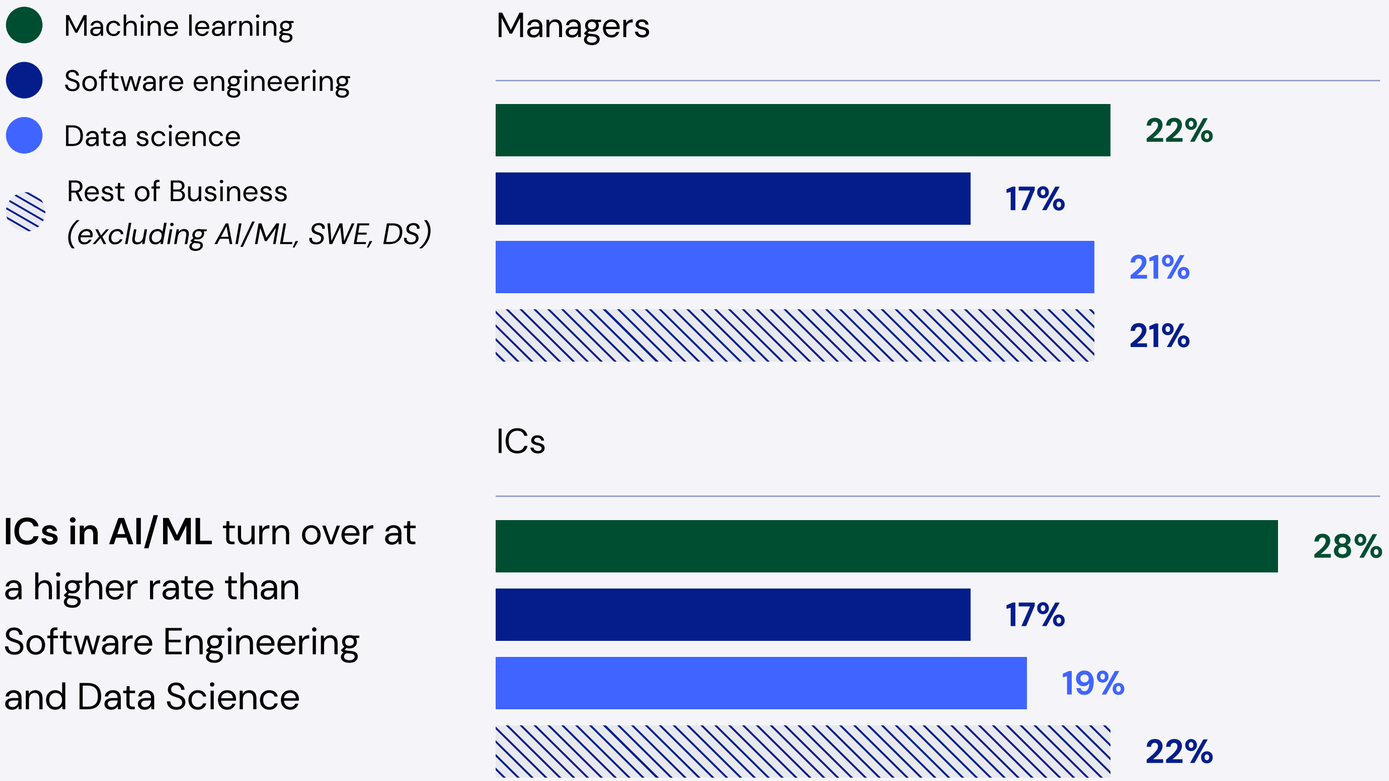
Let's take a closer look.

Attrition

Unsurprisingly, AI/ML engineers leave their jobs at a faster rate than others.

This trend is driven by rising demand for AI/ML talent, especially at individual contributor (IC) levels. When looking at IC roles, we see a clear difference across AI/ML, SWE and data science roles.

Annual attrition rates by engineering role



ICs in AI/ML turn over at a higher rate than Software Engineering and Data Science

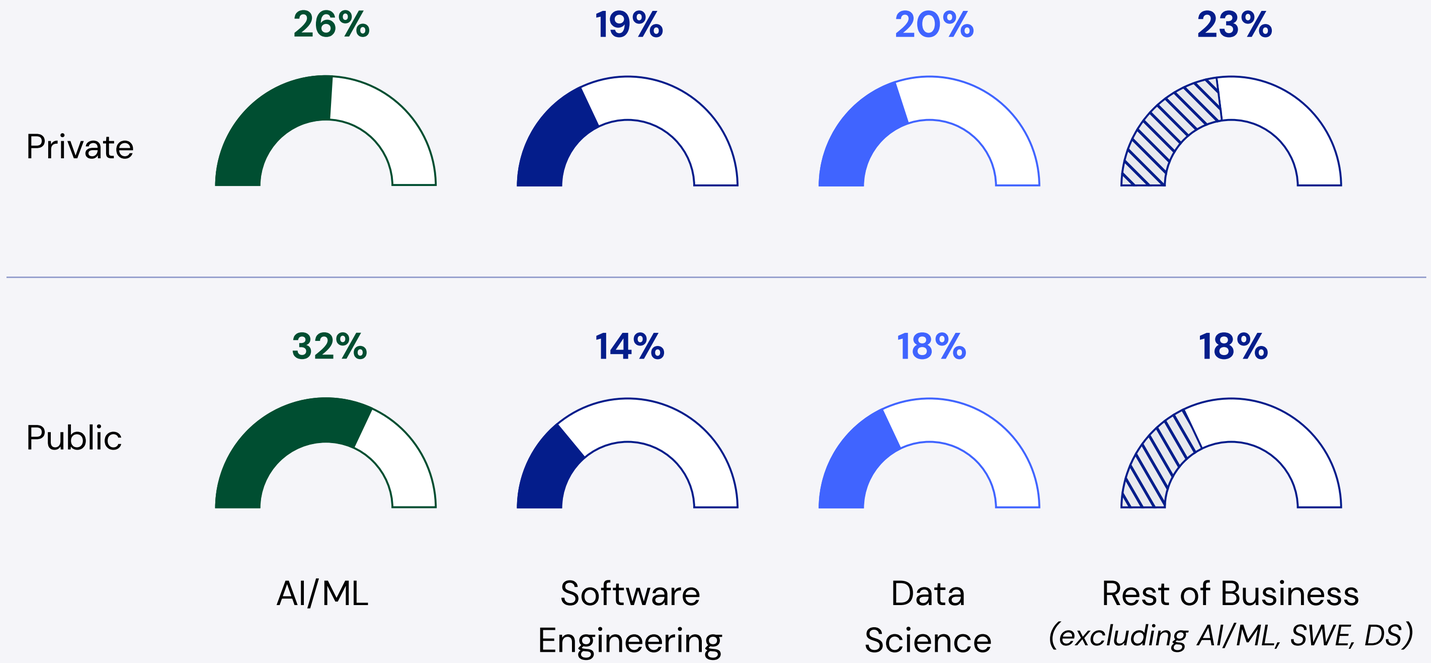
Sample size: 1.3K AI/ML employees, 62K SWE employees, 5.3K DS employees, 221K other employees. Filtered to companies with 100+ employees. Annual attrition rate is calculated as the % of employees in-seat 12 months ago that are still at the company today, and includes both voluntary and involuntary turnover.

Digging deeper into attrition rates, we find big differences between public and private companies.

One hypothesis is that some AI/ML talent moves from public companies to private companies, betting on the chance that equity packages will pay off in the long run. For these in-demand employees, now is their time to gamble on the private company lottery ticket, or to simply explore a work environment that's more entrepreneurial or on the cutting edge.

Another hypothesis is that the large and varied nature of the public market makes it easier for AI/ML talent to move to new roles or industries at public companies, whereas the private market offers a smaller pool of opportunities.

Annual attrition rate across engineer types for private and public companies



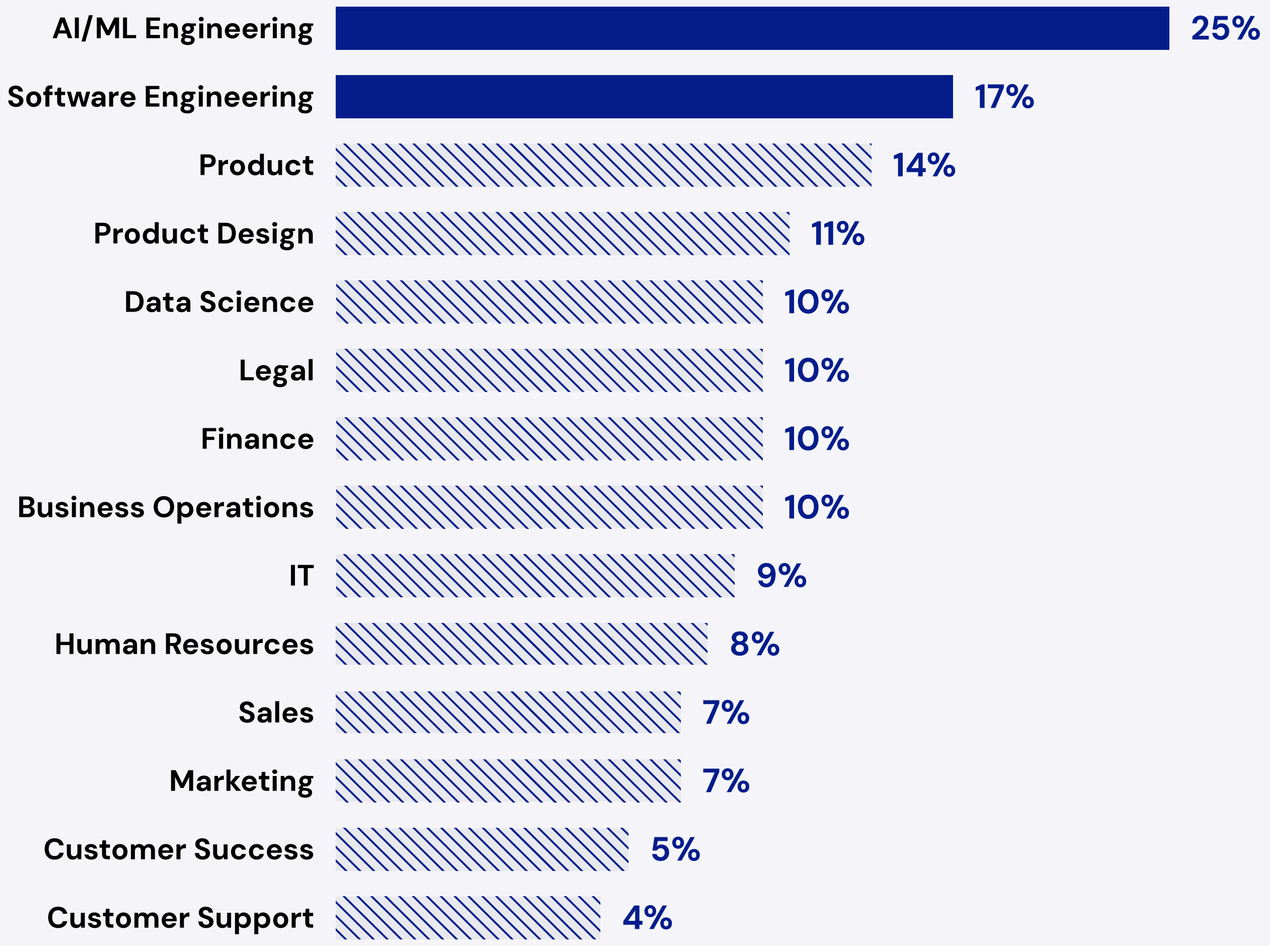
Sample size: 1.3K ML employees, 62K SWE employees, 5.3K DS employees, 221K Other employees. Filtered to companies with 100+ employees. Annual attrition rate is calculated as the % of employees in-seat 12 months ago that are still at the company today, and includes voluntary and involuntary turnover.

Sign-on Bonuses

Over the past 18 months, AI/ML Engineers received sign-on bonuses ~2.5x as frequently as candidates in other job families, according to data from Pave's Offer Insights.

Companies may use this tactic to help offset the opportunity cost of a new hire, and buy a candidate out of the equity, perks, or other compensation that they might be leaving on the table. Sign-on bonuses can be helpful to sweeten the deal in hiring efforts—especially when competition is so fierce.

Percent of offers that include a sign-on bonus in the past 18 months

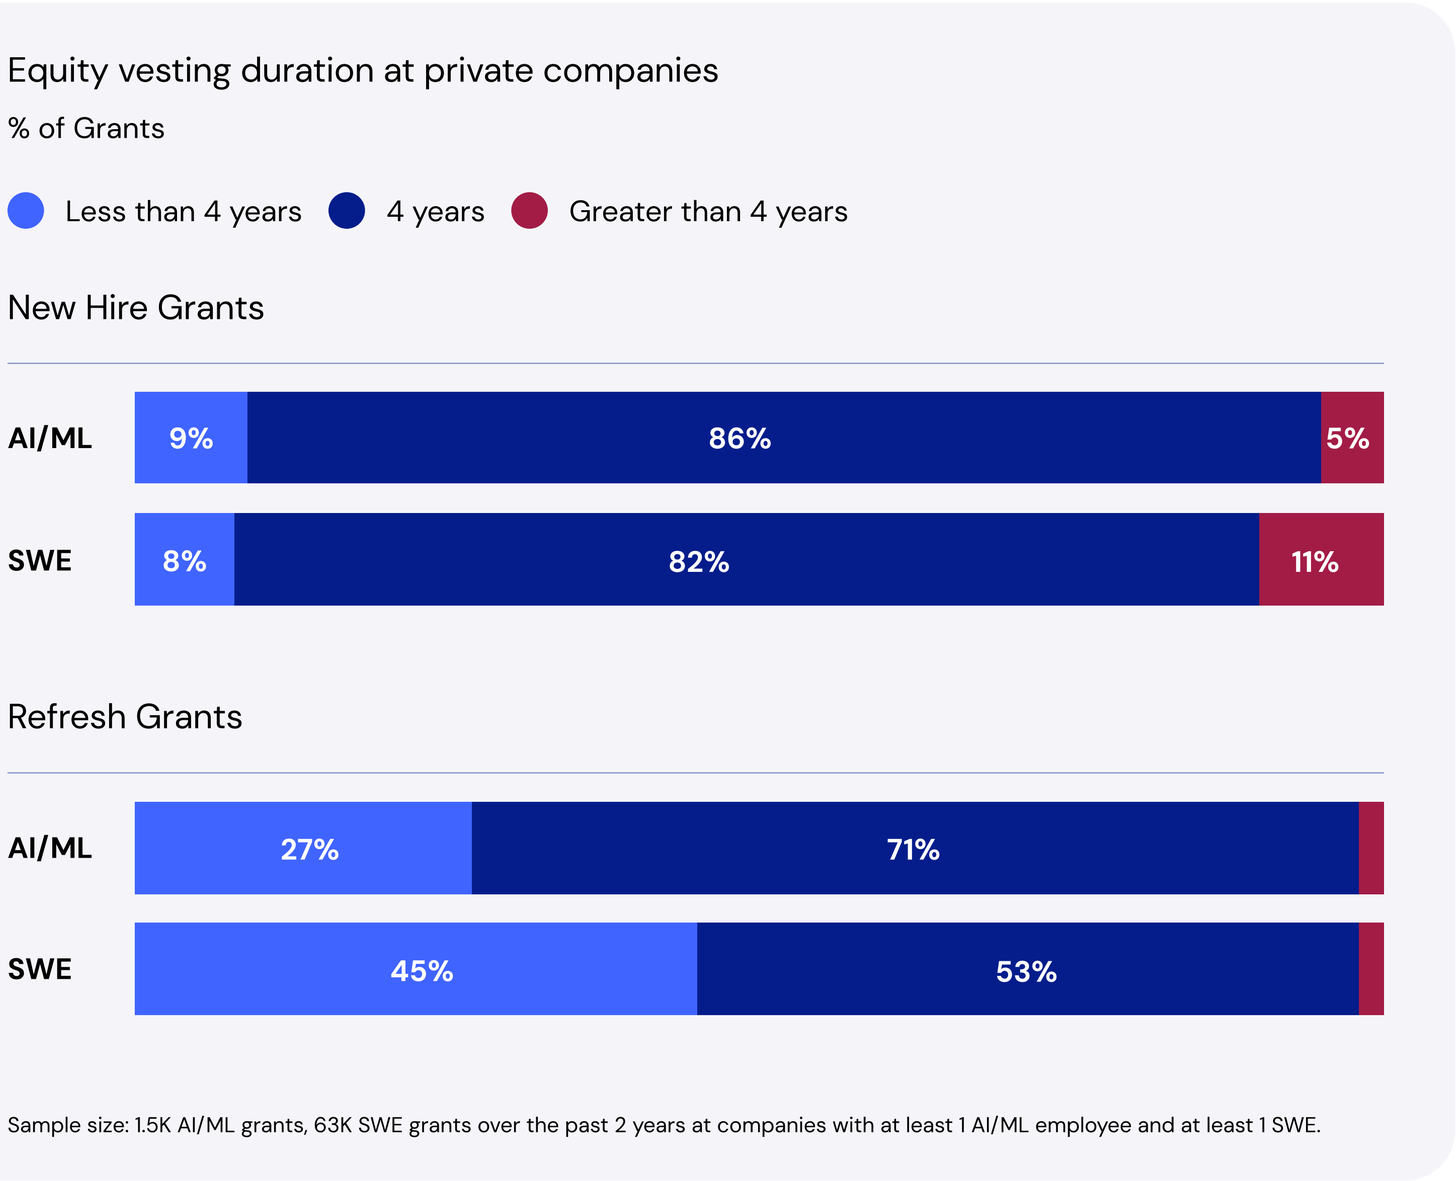


Sample size: 140K accepted offers over the trailing 18 months from 3K companies, powered by Offer Insights.

Equity Vesting Duration

At private companies, new hire equity vesting duration is similar across AI/ML and SWE roles. When it comes to refresh equity grants, however, AI/ML roles tend to receive refresh grants with a longer vesting period, with 71% receiving a 4-year vest compared to 53% of SWEs.

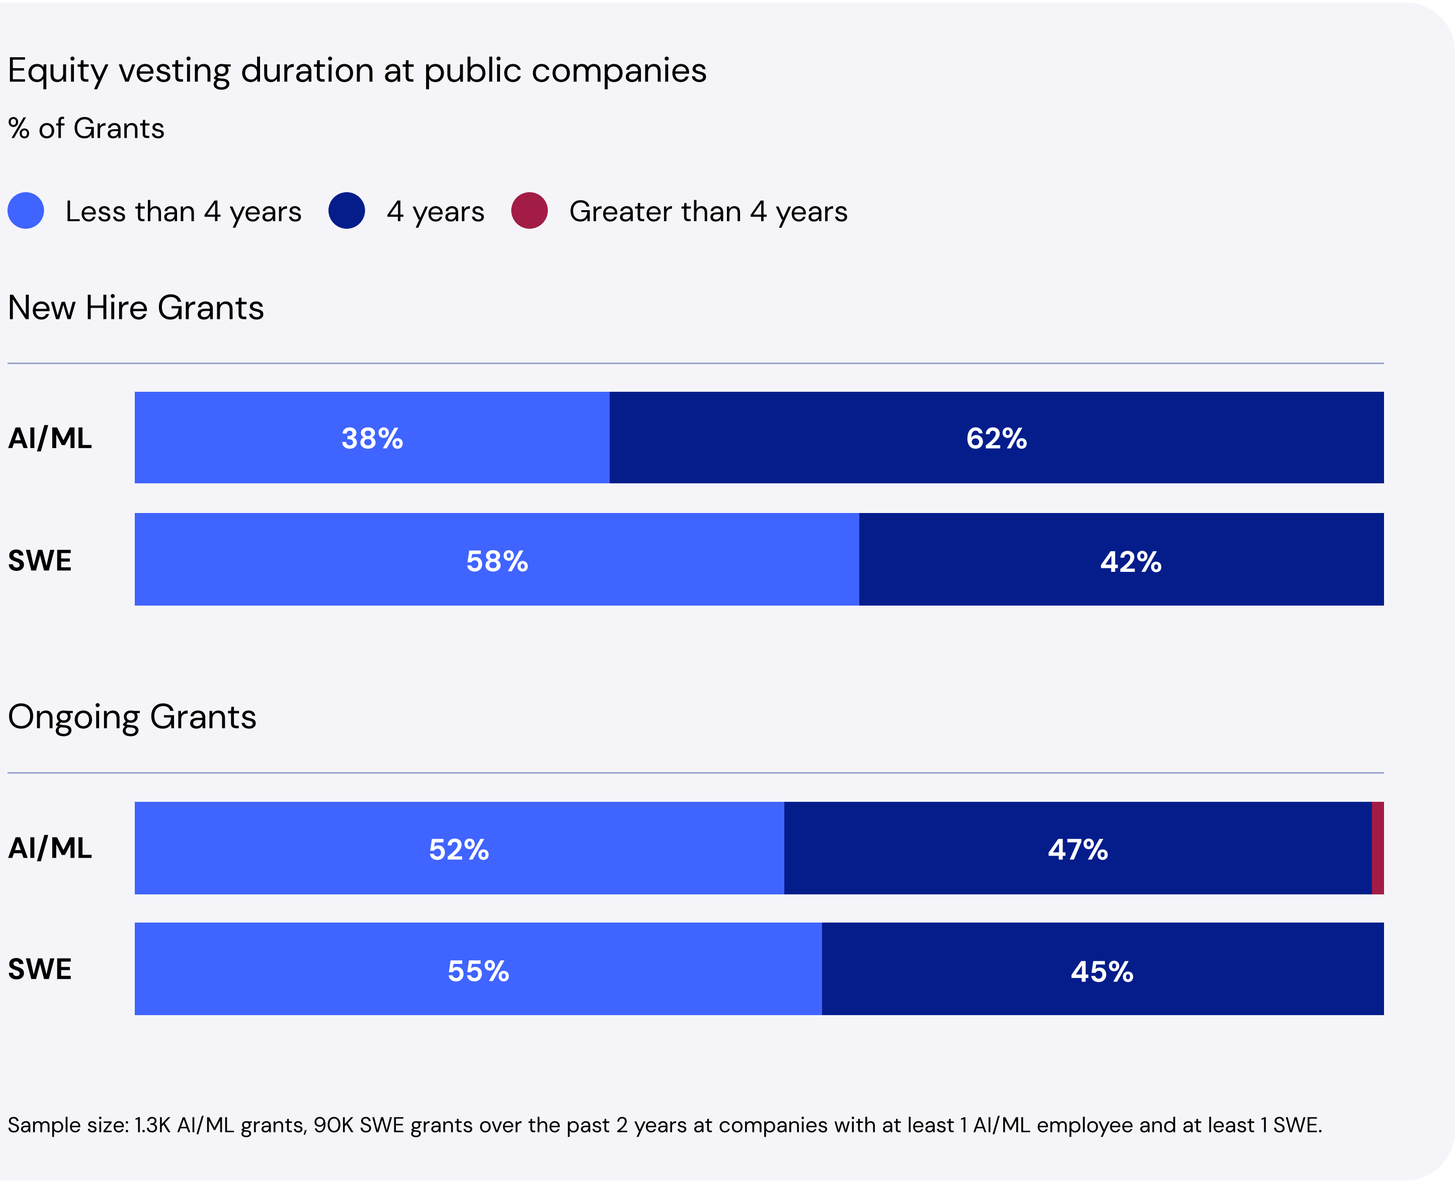
While private market practices can vary, this may indicate that private companies are leveraging longer vesting schedules as a retention play to keep these valuable employees within the organization.



Sample size: 1.5K AI/ML grants, 63K SWE grants over the past 2 years at companies with at least 1 AI/ML employee and at least 1 SWE.

At public companies, the data highlights another trend.

AI/ML roles tend to receive new hire grants that have a longer vesting duration relative to SWE, while the vesting duration of ongoing equity grants are more evenly distributed. Public companies may be leveraging these shorter vesting schedules for new hire grants as a tool to get AI/ML talent in the door. AI/ML talent also tend to receive larger grants, which are more likely to have a longer vesting duration.



Sample size: 1.3K AI/ML grants, 90K SWE grants over the past 2 years at companies with at least 1 AI/ML employee and at least 1 SWE.

05

Offer & Hiring Trends





Trends in Hiring for AI & ML Engineers

The boom in hiring for AI/ML roles accelerated rapidly after the launch of ChatGPT in late 2022. **Since then, offer and hiring data shows a clear growth trajectory for AI/ML talent.**

Let’s dive into the hiring trends, as well as location data for people in the AI/ML field.

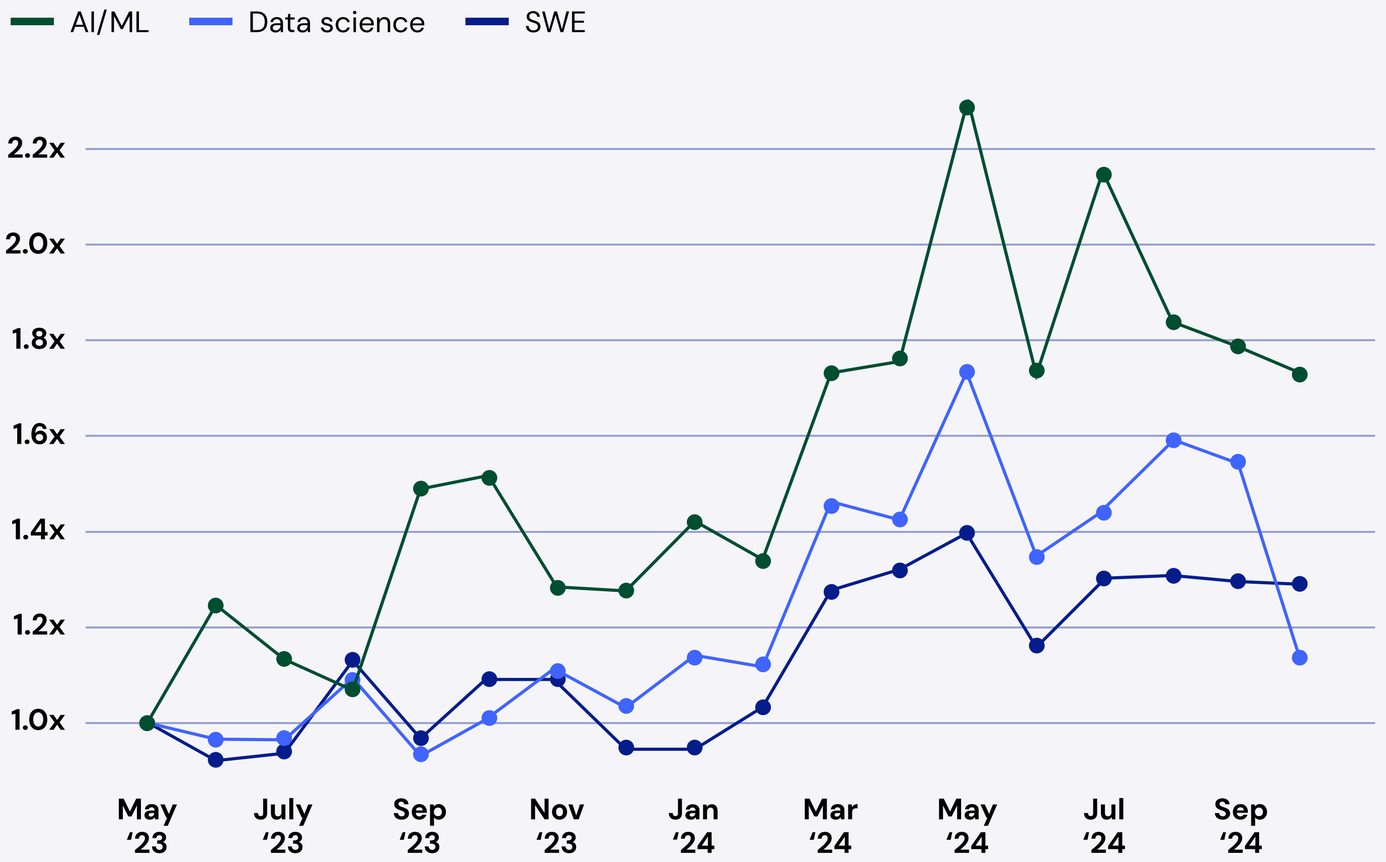
Offer Volume

Data from Offer Insights shows that the number of accepted offers for AI/ML roles has doubled in the last 18 months, outpacing Data Science and SWE.

That said, Data Science has also seen a lift, followed closely by Software Engineering.

It's important to note that what a job is called in an offer doesn't necessarily map to what the title will be when the employee is hired. Hiring data can fill in this other crucial part of the equation, as we see on the next page.

Number of accepted offers relative to May 2023 by month



Sample size: 2.7K AI/ML offers, 24K SWE offers, 6.9K DS offers, powered by Offer Insights.

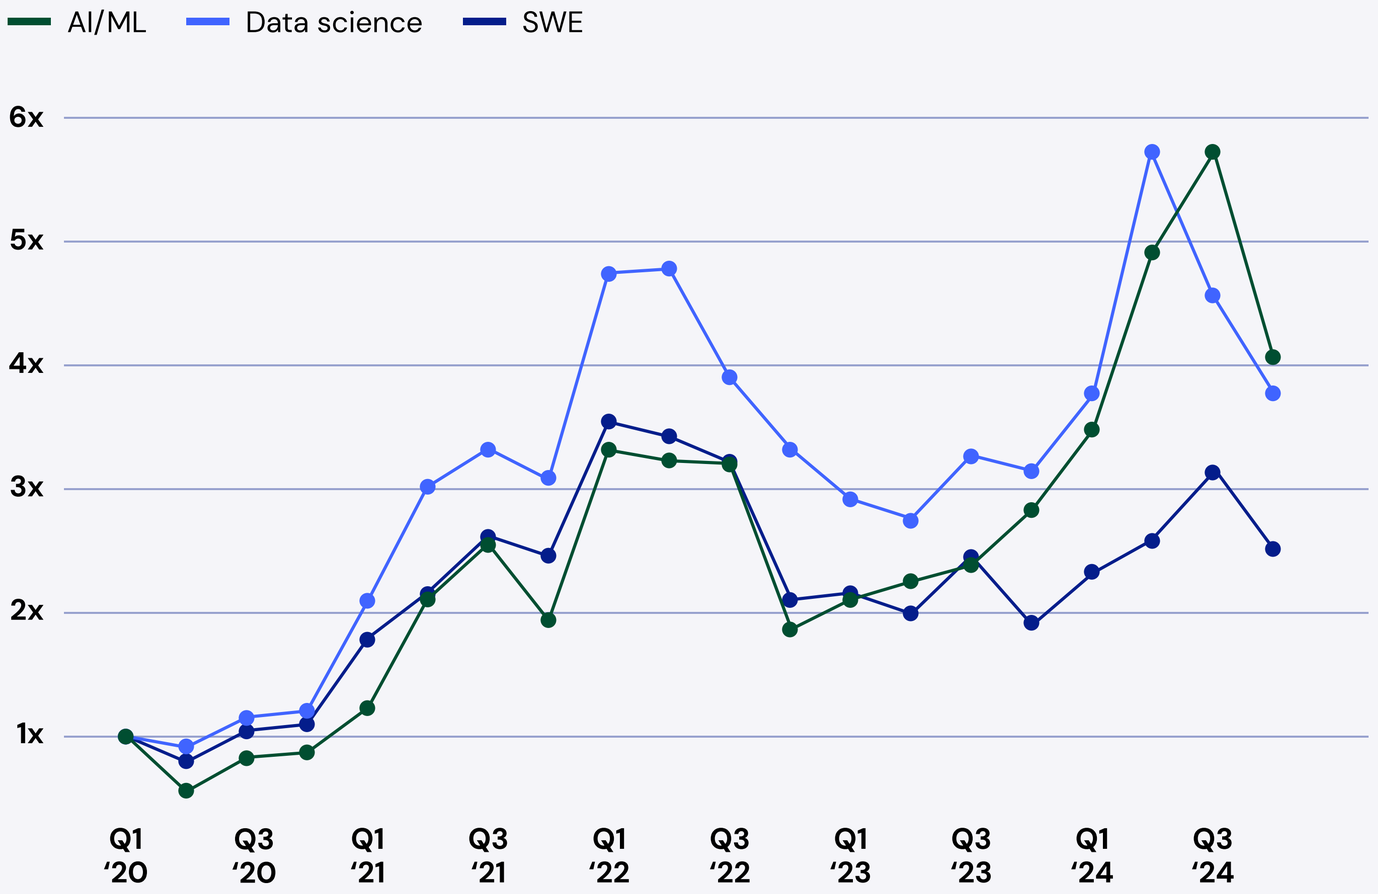
Hiring Volume

Data from Pave shows hiring for AI/ML, SWE, and Data Science talent followed a similar trajectory until the launch of ChatGPT. Then things changed.

In Q4 of 2022, Chat GPT launched and took the world by storm, prompting a surge in hiring for AI/ML engineering. Since then, the hiring rates for these roles more than doubled compared to SWE hiring rates.

However, the market is nowhere near bleak for SWEs. Hiring for these roles has bounced back well above the pandemic slump of 2020. In terms of sheer numbers, SWEs still make up a much larger pool than AI/ML.

Number of new hires relative to Q1 2020 by quarter



Sample size: 4.1K AI/ML employees, 137K SWE employees, 13.2K DS employees. Dataset: Pave HRIS integrations. Results only include current employees.

Where is AI & ML Talent?

Globally, the US is home to nearly 60% of AI/ML talent.

Among US-based employees, AI/ML talent is more likely to be in US Tier 1 cities than SWE, including metros like New York and Seattle.

Outside the US, these employees are more likely to be based in Canada and the UK than their SWE counterparts.

Top 10 US locations for AI/ML employees



1. **SF Bay Area** (CA)
2. **NYC Metro** (NY/NJ/PA)
3. **Seattle Metro** (WA)
4. **Boston Metro** (MA)
5. **LA Metro** (CA)
6. **Austin Metro** (TX)
7. **Washington DC Metro**
8. **San Diego Metro** (CA)
9. **Denver Metro** (CO)
10. **Chicago Metro** (IL)

Sample size: 2.2K AI/ML employees in the US.

Top 10 global locations for AI/ML employees



1. **United States**
2. **United Kingdom**
3. **Canada**
4. **India**
5. **Australia**
6. **Germany**
7. **Israel**
8. **China**
9. **Brazil**
10. **Poland**

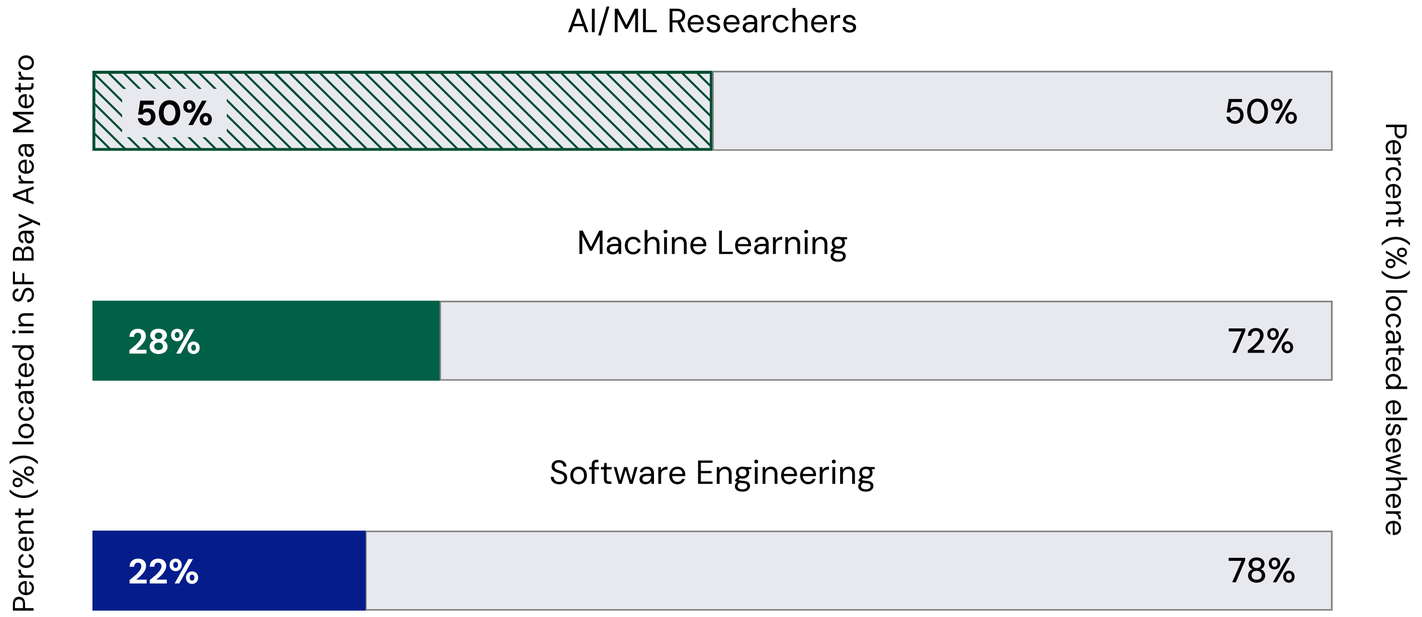
Sample size: 3.8K AI/ML employees globally.

Honing in on the Bay Area

AI/ML talent is heavily concentrated in the San Francisco Bay Area metro, especially highly technical AI/ML Researchers. In fact, 50% of the AI/ML Researchers in Pave's real-time dataset are based in this area.

Percentage of employees located in the SF Bay Area Metro by job type (among all employees in the same job type)





Sample size: 94K employees.

06

WITH NUA GROUP

Bringing AI & ML into Your Job Architecture



Job Architecture

Given the findings of this report, how do we see companies treating AI/ML roles from a job architecture perspective?

Career Pathways

Larger companies, particularly public ones, are more likely to have a separate career track for AI/ML engineers, at least at the individual contributor, manager, and senior manager levels. At smaller companies where AI and ML are a key strategic imperative, we may see AI/ML roles broken out in job architectures, but this is generally less common.

Among companies with a separate career path for AI/ML roles, titling conventions will distinguish these roles from traditional SWE roles. In companies that do not differentiate, the SWE career path will be used to accommodate AI/ML roles, though these roles may still be treated differently.

Given both the real and perceived pay premiums AI/ML roles enjoy, we expect general SWEs to be attracted to AI/ML career paths and work to acquire the requisite skills.

Cash Compensation

Getting accurate, up-to-date, and statistically robust salary data for AI/ML roles remains a challenge for many organizations. When using separate career paths, companies often anchor to SWE pay levels and apply a 15% to 20% premium.

When companies try to hire AI/ML talent using existing SWE pay scales, they often target the 75th percentile or higher. For advanced AI/ML talent (e.g., Research Scientists), cash compensation is often defined outside job architectures, even when adding a premium to SWE.

Pave is among the first data providers to publish benchmarks confirming base salary premiums of 15% or higher.

Equity

Similar to cash compensation, accessing robust data when developing equity guidelines is often a challenge. Using SWE as the anchor but applying a differential is the most typical approach. However, the premium is likely higher—perhaps up to 30% or higher for public companies.

We see new hire equity awards for AI/ML engineers are often increased compared to a company's normal practice to provide an even greater transfer premium. Similarly to cash compensation, it's often the case that equity compensation for advanced AI/ML talent is defined outside of job architecture.

Again, Pave is among the first to publish data highlighting clear equity premiums.

Management Roles

Management roles for AI/ML should show premiums on compensation and equity similar to the individual contributor track.

Higher pressure may come in some instances when hiring a "key" leader, i.e. someone who will attract other talent. That will be an important value to recognize in both cash compensation and equity.

While most executive roles still appear to be SWE based, the skills are likely evolving, and the importance of the role AI plays in strategy should be reflected in the premium given to the executive roles, if any. We may even see some companies go a step further and add a Chief AI Officer role.

Specialized Roles

For more specialized AI/ML roles, in particular AI or ML Research Scientists, the pay differences may be more substantial in some cases, recognizing the different work involved and the requirements of the role—typically needing a PhD. This would imply dealing with these roles outside normal pay structures, whether it is SWE or AI/ML.

Geographic Differences

Within the US, Pave data currently suggests using the SF Bay Area as a baseline and not adjusting for geography. This may change over time. Outside the US, organizations will likely use local SWE data and add a premium for AI/ML roles.

Thoughts for the Future

To maximize the value of AI/ML engineers, organizations must rethink their job architecture and pay structures to accommodate the unique demands of this talent pool.

It's important to note that companies will need to plan to regularly review these strategies, as "set it and forget it" won't work. In some ways, the market data is only catching up to actual practices and businesses need to stay tuned in to how the trends continue to evolve from here. We are by no means at a status quo.

Flexible Role Definitions

AI/ML engineers often work across disciplines. Job families should allow for cross-functional collaboration with product management, data science, and engineering, while also supporting deep specialization where needed. Longer term, it might be expected that a core requirement of SWEs will include AI/ML skills. Over time, this could shrink the premium we are seeing now for this skillset.

Dynamic Career Ladders

Establishing clear but flexible career ladders, including technical and managerial tracks, will be key. This might include roles like "AI Research Fellow" for senior technical experts or "AI Strategy Lead" for those bridging business and technology.

Leadership Development

Building leadership capabilities in AI/ML engineers will require targeted mentorship, exposure to cross-functional projects, and participation in strategic decision-making.

Retention-Focused Compensation

Beyond higher base pay and bonuses, companies should explore creative retention mechanisms, such as flexible vesting schedules, learning opportunities, and project-based incentives that align with AI/ML engineers' intrinsic motivations. Strong benefit packages with 401(k) match are positively viewed.

Global Talent Strategies

As AI/ML hubs emerge globally, companies should diversify their hiring and consider remote or hybrid models to access untapped talent pools. In addition, we expect competition for talent to broaden across industries as non-traditional employers of engineering talent are more likely to need and want AI talent.

pave + nua group

An Evolving Landscape

AI/ML engineers are reshaping the business and technology landscape, and organizations are evolving to fully harness their potential.

By developing robust job architectures that emphasize clarity, flexibility, pay equity, and pay transparency, companies can ensure their AI/ML engineers thrive—not only as contributors but as strategic drivers of innovation. **Addressing questions around their compensation, career development, and organizational integration, as well as keeping on top of evolving trends, will be pivotal in staying competitive in this rapidly advancing field.**



ABOUT PAVE



Pave simplifies compensation for everyone in your business.

Purpose-built for total rewards professionals, our complete compensation management platform, including real-time benchmarking data, ensures all your programs deliver for leaders, people managers, employees and candidates.

Learn more at Pave.com

ABOUT NUA GROUP



Nua Group is an independent broad-based human resource consulting firm specializing in compensation, benefits and HR optimization, focused on providing companies with unbiased, holistic and expert advice.

Learn more at NuaHR.com