

Preface:

Before we begin, I'll assume you're a sentient individual capable of self-discovery. Please don't just take my word for anything—verify it yourself. I'm not an “expert” or “guru,” and this is not financial advice.

I'm simply a tech professional connecting the dots.

Take what's valuable to you, and feel free to discard the rest.

However, if you find this information helpful, please let me know at *fido.selftaughthustle@gmail.com*. Your success would make all the hours and sacrifices that went into creating this report worthwhile.

Introduction:

America's economy faces more profound challenges than it appears at first glance. Wages are not keeping pace with rising living costs, and more workers are relying on debt “to get by.” Uncontrolled spending, combined with the absence of affordable essentials, often gives people the impression that they are just treading water without progressing.

Even technologists—like you, my beloved reader—who usually earn higher incomes, struggle in areas where major tech companies drive up housing prices and living expenses. Moreover, the “tech bro” boom following the 2008 financial crisis has faded as the industry experiences wave after wave of layoffs. Foreign labor, AI, and a growing supply of tech talent threaten today's American technologist's ability to earn a comfortable income.

In addition, the once-strong link between tech and remote work is disappearing. Even Zoom, the champion of remote work, has called its employees back to the office. For the first time in 17 years, the average American tech worker has little leverage in the job market.

After many hours and scanning the internet like a madman, I see no easy fix for technologists seeking career growth in 2025. There are no shortcuts or quick hacks, but there are solutions and paths forward if you're willing to work your butt off.

My research and professional experience show that preparing for the worst possible outcome should begin today. This market rewards those who solve problems methodically, yet identifying the correct issues and developing the skills to solve them requires a “moonlight” effort.

In short, crafting your ideal career—finding a new role or building a business—takes time.

This means you actively build your next career step after work ends between 5 and 9 p.m.

I've looked for work, added a new skill, and built businesses by choice when employed and out of necessity when laid off. Anyone sane who's gone through the same would say:
"Build your next move now while you have a job."

If you wait until you're let go, you will experience tremendous pressure. I wouldn't wish that stress on anyone, especially a father responsible for his family.

I don't mean to sound like an alarmist, but the industry is volatile—frankly, I assume you know this. So, if you have managed to avoid a layoff at this stage of the decade, I recommend remembering the following phrase,

"Act today so you won't be caught unprepared tomorrow."

In the following chapters, I'll explore the income opportunities available to technologists in this rapidly shifting world and show how these opportunities may affect your quality of life. By the end of this report, you'll have a clearer view of your best next steps in a changing career landscape.

Let's not waste any time and get started.

Chapter 1: Income and Career Opportunity in AI

Demand and Growth Potential

Having lived through recent technology hype cycles—crypto, the metaverse, virtual reality—I've become skeptical toward tech "innovation." Each of these technologies once boldly claimed to be the "next big thing," yet by 2025, none have meaningfully reshaped daily life.

I don't assume AI will achieve the same "lukewarm" outcome, and I think it's a technology that should be taken seriously. However, because of my experience with "pump and dumps," I also believe it is important to proceed cautiously when assessing anyone's viewpoint on the technology, including my own.

Right now, AI's actual economic contribution is still evolving day by day.

One sobering test for AI's actual value is its impact on global GDP growth. So far, AI hasn't produced substantial, large-scale economic changes akin to the Industrial Revolution. This trend generally aligns with digital products of the past. Accordingly, wealth disparity is historically high despite software-enabled conveniences.

AI's growth potential is also a double-edged sword for tech workers since many applications explicitly aim to replace human labor. It remains unclear whether AI creates as many jobs as it displaces. It's not evident that AI, as it stands today, has led to actual job displacement.

In 2024, the United States saw [81.298 new job openings related to AI](#), while [150.000 tech workers were laid off during the same period](#).

Side note: Regarding job openings, the data here is ultimately derived from the US Bureau of Labor Statistics, so please take its findings with a grain of salt. Moreover, the 150k tech workers laid off are global numbers, not just the United States.

Notably, the number of AI-related job openings increased by nearly 25% between 2023 and 2024. Moreover, there has been no documented correlation between the mentioned layoffs and AI.

As it stands today, I'd argue that the current demand for "AI jobs" seems to be propped up by investment capital rather than solid, profitable business models.

OpenAI is a prime example: despite rapid revenue growth—from \$3.7 billion in 2024 to an estimated \$11.6 billion in 2025—it remains unprofitable, reporting losses of \$5 billion in 2024

High operating costs and prioritization toward market expansion mean profitability is still uncertain. However, large investors like Microsoft (which has invested \$14 billion) anticipate and seem eager to see future returns on their spending.

At this point in its history, AI is closer to a speculative bet for all its stakeholders than a mature industry.

Accordingly, major players like [Meta, Amazon, and Microsoft collectively invested \\$246 billion in AI and data center expansions in 2024](#), fueling this speculation. But significant funding alone doesn't guarantee winners—technology has always had more failures than successes.

Moreover, the AI investment boom mirrors the 2000 dot-com bubble, where media hype overshadowed business fundamentals like profitability. In those days, once speculative capital dried up, panic ensued.

Therefore, we must avoid confusing the ample supply of general-purpose technology (AI) with genuine market demand. Adoption exists but remains clumsy. Initially, heavy investments focused on building large AI models, yet attention has recently shifted to the belief that value resides at the application layer.

At this stage, genuine business demand is possible—but speculative. Investors will eventually seek concrete returns on these vast investments.

In summary, my analysis reveals AI's substantial growth potential and investor enthusiasm yet emphasizes that current industry dynamics are primarily speculative, lacking clear evidence of sustainable profitability or broad economic impact. While AI adoption and funding remain high, actual demand and value generation remain uncertain, urging cautious optimism rather than blind enthusiasm.

Quality of life

Quality-of-life considerations in AI-specialized roles vary widely, mirroring broader trends across the technology industry. Work arrangements span from fully remote and hybrid to entirely in-office models. [Notably, technology employers had the highest rate of fully remote job postings \(18%\) compared to other sectors, such as legal \(8%\) and healthcare \(9%\). Beyond fully remote roles, hybrid work arrangements have also become increasingly common.](#)

However, the number of tech employees returning to office-based work has surged recently. The pandemic-induced era of remote work as a widespread norm across tech companies is effectively over—even [Zoom has called its employees back to the office](#). Ultimately, remote or hybrid work options are evaluated individually, reflecting broader industry patterns rather than unique policies specific to AI roles.

Furthermore, working hours in AI-related positions significantly differ depending on company size and culture. Early-career professionals might accept long workdays at startups, viewing it as a necessary career investment. Conversely, more experienced professionals with outside priorities often prefer stable, predictable hours typically offered by established corporations.

Compensation:

Despite uncertainties regarding AI's long-term economic impact, current data highlights clear compensation advantages for tech workers specializing in AI. For instance, senior-level AI engineers earn approximately 19.5% more on average than their traditional software engineering counterparts, with even more significant discrepancies observed in tech hubs like San Francisco (Coursera, 2024; ZipRecruiter, 2024; Dice, 2025).

Considering the AI market is projected to reach nearly \$2 trillion by 2030, this salary premium might expand further, indicating continued demand for specialized AI skills (Robert Half, 2025; LifeHealth, 2025). However, given the speculative nature of the AI sector, these increased earnings could also be accompanied by greater volatility.

As with most compensation scenarios, accurate salary expectations significantly depend on the specific company and geographic location. Earning nearly \$200,000 annually may initially appear lucrative until accounting for the high cost of living in cities like San Francisco, where median monthly rents exceed \$4,271. Conversely, securing a similar salary in a region with lower living costs and a remote working arrangement might "feel" substantially more valuable.

Before conducting this analysis, I assumed the compensation gap between AI specialists and general tech professionals would be even more significant. A 19.5% salary increase is certainly notable; however, for someone already earning upwards of \$140k annually, especially with disciplined spending and modest living costs, this increase might feel negligible in terms of improving overall quality of life.

Given this modest premium, retraining to specialize in AI or related skills may represent career longevity and market relevance more than purely financial motivation. Of course, this is anecdotal.

Summary:

In conclusion, AI specialists face workplace realities aligned closely with broader tech-industry standards, with remote and hybrid flexibility determined individually rather than uniquely associated with AI roles. Compensation data demonstrates an approximate 19.5% salary advantage for senior AI engineers compared to traditional software engineers, particularly in tech-centric locations like San Francisco. Yet, when factoring in living costs and market volatility, professionals likely pursue AI specialization primarily for career stability and relevance rather than solely for higher financial gain.

The following section will evaluate specific AI-related skills, learning curves, and educational resources required for successful specialization.

Below is a chart to visualize the compensation across different roles.

Role	Average Salary (USD)	Entry-Level (0–1 years)	Mid-Level (4–6 years)	Senior-Level (10–14 years)	Source
AI Engineer	\$134,188	\$103,140	\$138,301	\$172,468	1,4
Machine Learning Engineer	\$123,117	\$98,798	\$122,505	\$153,286	1,3
AI Researcher	\$99,578	\$94,972	\$114,931	\$142,511	1
Software Engineer	\$118,366	\$99,438	\$121,118	\$144,306	1,6
Data Scientist	\$118,399	\$110,720	\$127,098	\$145,724	1
Generative AI Engineer	\$112,693–\$155,695	\$91,096	\$117,680	\$146,030	3

Skills, Learning Curves & Resources:

Titles and terms in AI are notoriously confusing. Industry naming conventions are inconsistent and ephemeral and have even become memes. Shout out to all the professional “prompt engineers” and “vibe coders.”

Approaching what to learn in AI solely based on job titles, at least from a technical perspective, is somewhat “shaky” from a technical standpoint. Titles aren’t entirely insignificant, but understanding the AI technology stack will more effectively focus your time and effort—assuming you've already decided to specialize in AI.

As with any programming specialty, market fit for your skillset is crucial. You should ideally acquire skills that align closely with current market demands. Becoming a chess grandmaster, for instance, is impressive but holds limited market value regarding employment opportunities. There aren't substantial market demands uniquely solved by expertise in chess.

Experienced engineers likely understand this principle, but for those earlier in their careers, it's critical to grasp this to avoid wasting valuable time, energy, and resources.

For clarity, I'll anchor my evaluations of learning curves to a hypothetical mid-level software engineer. If you aren't currently coding, this information can still provide valuable context, but framing it this way allows for more precise analysis.

Additionally, I'm an engineer and not an AI expert. Please don't take my word alone. For your career's sake, use this report as a starting point—not as definitive guidance. Always independently verify that a career in AI suits your situation and aspirations.

The easiest way to understand AI is to analyze its technology stack practically. Andrew Ng, a highly respected pioneer in AI, originally conceptualized this stack. He brings an academic yet refreshingly practical and sober perspective to a hype-filled field. If you're interested, please check out his work; he maintains an impressively grounded approach to AI technology.

Below is the *The AI Stack*:

Layer	Description	Examples/Tools	Notes & Commentary from Andrew Ng
Application Layer	Specific AI applications solving real-world problems.	Amorai (relationship coaching), Bearing AI (shipping optimization), visual AI apps, health care diagnostics	The biggest long-term opportunity; applications must generate value to sustain the lower layers.
Agentic Orchestration Layer (Newly emphasized)	Frameworks and workflows that facilitate complex, iterative, agent-driven AI tasks.	LangChain, Landing AI's Vision Agent, HuggingGPT, LlamaIndex, AutoGPT	Enables advanced AI workflows (reflection, tool use, planning, multi-agent collaboration), significantly improving AI performance and ease of development.
Foundation Model Layer (Explicitly clarified)	Large-scale AI models that power generative AI capabilities.	OpenAI GPT-4, Anthropic Claude, Google Gemini, Meta Llama-3	Receives significant media attention; critical for generative and multimodal AI, optimized increasingly for agentic workflows and tool integration.

Cloud Infrastructure Layer	Infrastructure providing compute and storage resources for AI workloads.	Snowflake, AWS, Microsoft Azure, Google Cloud Platform	Capital-intensive, essential for scaling AI, with established major providers.
Semiconductor (Hardware) Layer	AI-specific hardware, chips, and semiconductor technologies.	NVIDIA GPUs, Google TPUs, Cerebras, AMD MI-Series GPUs, SambaNova	Highly concentrated, extremely capital-intensive; fewer winners due to barriers to entry. Critical for speeding up AI inference and training.

Okay, the table might initially seem daunting, but bear with me.

In the table, darker shades indicate lower-level technologies requiring deeper specialization, whereas lighter shades are higher in the stack and generally more accessible for software engineers.

As a mid-level software engineer, several stack layers are accessible to you. However, before addressing these directly, let's briefly clarify learning curve expectations across the stack:

- High Difficulty (Long Learning Curve):**
 The Semiconductor (Hardware) and Foundation Model layers require extensive expertise, making them unrealistic targets within a reasonable learning timeframe. These areas typically require years of specialized training and experience. While nothing is impossible to learn, investing heavily in them might not be fruitful, given practical constraints.
- Moderate Difficulty:**
 The learning curve for the Cloud Infrastructure Layer is more manageable. Many software engineers already have experience in cloud environments. Familiarizing yourself with cloud services used specifically for AI training and deployment can significantly enhance your value to employers.
- Accessible and Immediately Relevant:**
 The Application Layer is the most natural fit for traditional software engineers. This layer is about building practical, real-world applications with specific AI integrations.

Ng emphasizes that the top two layers (Application and Agentic Orchestration) are especially crucial for industry sustainability since the revenue generated at these layers funds the lower layers.

Remember my earlier mention that OpenAI is not yet profitable? Their chat interface at the Application Layer is an example of revenue-generating activity at the top of the stack.

The good news for software engineers is that your existing skills significantly overlap with this layer—existing programming knowledge and familiarity with AI frameworks and libraries position you to build AI-integrated applications effectively. Although generative AI is inherently

non-deterministic, much industry effort is currently directed toward enabling AI to solve challenging problems reliably. This effort aligns perfectly with software engineers' traditional skills and may be worthwhile.

The one layer I'm less familiar with is the newly emphasized Agentic Orchestration Layer. While this layer received considerable hype in 2025, it's still evolving rapidly, and its practical applications are less stable and clear-cut. Exploring this area may be beneficial, but proceed cautiously due to its fluid state.

Of course, each layer requires some dedicated learning, regardless. You'd likely have to familiarize yourself with existing models, terminology, new tools, and concepts such as agentic AI and other emerging stack layers. However, the learning curve is quite reasonable for a motivated, astute learner who commits time and effort.

If you decide to specialize in AI, an excellent starting resource today is [DeepLearning.AI](#). Founded by Andrew Ng, this platform offers courses taught by industry practitioners. Courses are free or paid. The company also regularly updates its content to keep pace with industry developments.

I've personally visited the site periodically over the past year and found the content continually updated and relevant. To be clear, I'm not affiliated. However, it may be worth your time to dive into their materials and select the subjects and instructors best suited to your learning style.

Let's wrap up this section and move forward.

Summary:

AI holds significant potential for growth, attracting considerable investment and enthusiasm, yet currently lacks evidence of widespread, sustainable profitability or transformative economic impact. While compensation for AI specialists is higher than for traditional software engineers, especially in tech hubs, the modest salary premium and volatile industry conditions suggest that upskilling in AI may be driven more by career self-preservation than solely financial gain. Engineers seeking specialization can focus their learning on practical skills in the higher layers of the AI stack, leveraging their existing expertise while approaching this field with cautious optimism and critical evaluation at every step.

Sources for Profitability and Investment Claims

- [The New York Times \(2024\). *OpenAI—ChatGPT Investors Funding*.](#)
 - [Foundation Capital \(2024\). *Why OpenAI's \\$157B Valuation Misreads AI's Future*.](#)
 - [Salon \(2025\). *OpenAI Valued at \\$150 Billion Isn't Profitable Yet—Should That Be Normal?*](#)
 - [CMSWire \(2025\). *4 Truths About OpenAI's Wild Financial Position*.](#)
 - [Salon \(2025\). *OpenAI Valued at \\$150 Billion Isn't Profitable Yet—Should That Be Normal?*](#)
-

Sources for Chart Related to Salary

1. [Coursera \(2024\). *Artificial Intelligence Salary*.](#)
 2. [Mobility \(2024\). *AI Engineer Salary*.](#)
 3. [Coursera \(2025\). *Generative AI Salary*.](#)
 4. [Dice \(2025\). *Tech Salary Report: Salary Trends*.](#)
-

Sources for Claims Related to Compensation

- **Coursera. (2024).** *Artificial Intelligence Salary: Your 2024 Guide*. Retrieved from <https://www.coursera.org/articles/artificial-intelligence-salary>
 - **Dice. (2025).** *Tech Salary Report 2025*. Retrieved from <https://www.dice.com/technologists/ebooks/tech-salary-report/salary-trends.html>
 - **LifeHealth. (2025).** *Tech Salary Trends 2025*. Retrieved from <https://www.lifehealth.com/tech-salary-2025/>
 - **Robert Half. (2025).** *Technology Salary Guide 2025*. Retrieved from <https://www.roberthalf.com/us/en/insights/salary-guide/technology>
 - **ZipRecruiter. (2024).** *Artificial Intelligence Salary Data*. Retrieved from <https://www.ziprecruiter.com/Salaries/Artificial-Intelligence-Salary>
-

Sources for Charts Related to the AI Stack

- [YouTube: *AI Stack Overview 1*](#)
 - [YouTube: *AI Stack Overview 2*](#)
-