

For this interview, I prepared a similar topic, but different in core. I interviewed someone with a stronger background in the AI industry about the baseline for AI. I'm going to dive into some details, trying not to be too technical. Generally, when you train an AI model, you need "ground truth" so that from what's right and wrong, the model knows how to judge. But how do those companies define ground truth? What I'm doing in research is easy: the amount of precipitation, which is a certain number. But how about HireVue, the interview AI that tortured us? Many AI systems are not predicting something as clean as rainfall. For example, in hiring tools such as HireVue-style interview platforms, the question becomes much harder to answer. If an AI system is evaluating a candidate, what is the "right answer"? Is it whether the candidate later becomes a high-performing employee? Whether a human interviewer liked them? Whether they stayed at the company for several years? Whether they matched the traits of previous successful hires? This became the main theme of my interview: AI does not simply learn truth; it learns whatever humans choose to define as truth. That choice can be useful, but it can also carry bias, business incentives, and hidden assumptions.

The first formal question I asked him, is that how the whole process is being monitored. Medical AI, are monitored through a relatively rigorous process demanded by FDA. However, those AI general products, are not being monitored at all, bringing them larger grey space for their actions. However, the competition within the market may bring some balance. But if there's something that bothers most in the industry, it is likely that competitors would not reveal each other.

The general practice, is what I asked next. How do we set the ground truth? For Hirevue, it is said that they used several experts to give ratings. There is no perfect ground truth for "good candidate." Companies usually choose a criterion, such as job performance, retention, or manager ratings. But each of these choices has limitations. Past hiring decisions may reproduce historical bias. Manager ratings may be subjective. Retention may reflect organizational fit or workplace conditions, not just candidate quality. It depends on the task. For objective prediction problems, ground truth may come from measurement. For subjective or human-centered problems, companies often rely on expert labels, consensus review, or proxy outcomes. The important thing is that ground truth is not always naturally given. It is often designed.

Finally, some more related to myself, that later, would he believe that those applied industry of ai would be good for career path. His answer was more cautious than I expected. He said he was not fully negative about AI, but he did not think I should define my future too narrowly around the current version of applied AI. The field changes too quickly. The dominant modality keeps shifting from text to images, audio, video, multimodal systems, agents, and workflow automation. Because of that, a person who only becomes good at one specific AI tool or application layer may become outdated quickly.

Instead, he advised me to build skills that can survive those changes: problem definition, data judgment, model evaluation, domain expertise, and the ability to explain uncertainty clearly. This made me reflect on my own career goals. I am interested in AI because I like applying technical systems to real-world problems, but I should not confuse tool fluency with long-term expertise. His advice helped me see that the safest foundation is not chasing the

newest AI trend, but becoming someone who can understand a problem deeply, decide whether AI is actually useful, and communicate both its power and its limits.