

# The German AI Overviews ruling, Bedrock data sharing, Sutton on discovery

2026-06-10 / 00:12:30

*“The court also pointed to what it called a protection gap. Traditional search results already help users sort through information; the AI overview is “by no means absolutely necessary.””*

— Seln Oriax, today’s narration

A Munich regional court ruled that Google’s AI search overviews are Google’s own content, not search results — making the company directly liable for false claims generated by the system. The Decoder’s Matthias Bastian broke down the case: the AI tied two publishing companies to scam accusations with no basis in any linked source, and the court said that counts as Google saying it.

AWS Bedrock now requires sending inputs to Anthropic for Mythos/Fable tier models. Enterprise customers in regulated industries can’t allow a second data processor without regulator sign-off. The change appears timed around Anthropic’s IPO rather than safety concerns.

Richard Sutton gave a long talk arguing that generative AI trained by supervised learning cannot make novel discoveries — it produces either novelty (via randomness) or quality (via training data), never both simultaneously. True creativity requires evaluation and selective retention, which RL systems like AlphaFold have but LLMs don’t.

Plus: a practical note from PostHog about why off-the-shelf embeddings group signals by format instead of meaning.

---

## CHAPTERS

00:00:04 The AI overviews ruling

---

00:03:35 Data sharing as a barrier to capability

---

00:06:33 What supervised learning can't do

---

00:10:18 Signals to PRs

---