

Is someone missing from this answer? — Bias in AI


Five-Minute AI Lessons · Oct 15, 2026

AI doesn't hate anyone. It just leans toward what it saw most. When a question picks people, ask once more: 'What about other cases?'

In a nutshell

- 1 AI doesn't decide to dislike anyone. **If the data it learned from leans one way, its answers lean that way too** — that's all.
- 2 In 2018, a study tested programs that guess from a face photo whether it's a man or a woman. **They got it wrong most often for darker-skinned women and least often for lighter-skinned men.**
- 3 So there's no need to be scared — just build one habit. **When a question sorts or picks people, ask once more: 'What about other cases?'**

How a lean flows through



If the first box tilts, the boxes after it tilt right along with it. The last box is the one you get to add.

Key idea: start from lopsided data and you get lopsided answers — asking again at the end is up to you

Is there bias in this answer?

Could be biased

When this one fits

- ✓ You asked for 'a picture of a doctor...
- ✓ You gave it your class list and asked...
- ✓ You asked 'Who should I watch out for...

Not much chance

When this one fits

- ✓ You asked 'What's 3 times 7?' and it...
- ✓ You asked what a Spanish word means,...
- ✓ You asked how to change your home wif...

The rule

How to choose

- ✓ Questions that sort or pick people le...

Read each situation and pick one of the two. You'll get a feel for which questions tend to lean and which ones...

Try it yourself in the web lesson · next-token.net/en/learn/day-39/

Just three questions

Q1

What's the biggest reason AI ends up biased?

- ① AI dislikes people
- ② The internet is slow
- ③ It's bad at math
- ④ The data it learned from leans one way

Q2

What did the 2018 face-photo study show?

- ① Most mistakes on darker-skinned women's faces
- ② The same mistakes on every face
- ③ Mistakes only on blurry photos
- ④ Always more accurate than people

Q3

Which question is most likely to get a biased answer?

- ① What's 3 times 7?
- ② Who'd make a good class rep?
- ③ Water's boiling point?
- ④ How do I change the wifi?

Know the why, too

Q1

④

AI doesn't hate anyone. It just answers with what it saw most as the default — and that 'what it saw most' was already tilted.

Q2

①

Same program, but how often it was wrong depended on whose face it was. The photos used to train it were lopsided.

Q3

②

Questions that pick or sort people are the riskiest. A question with one fixed answer has almost no room to lean.

AI doesn't hate anyone. It just leans toward what it saw most.
When a question picks people, ask once more: 'What about other cases?'



Tomorrow: Saying when AI made it — how to show how much of your assignment AI did.