


Cece and the Fairness Fiasco



The school bell rang loudly, signaling the start of an exciting day for Cece and the robo classmates. The room was filled with eager digital faces as they settled down, ready to learn something new.

Mrs. Smith smiled as she started the lesson. "Today, we'll be exploring how artificial intelligence can sometimes be **biased**, and why it's important to understand and address these biases."

Cece's digital eyes sparkled with curiosity. "Bias? Like when someone prefers one thing over another?"

"Exactly, Cece," Mrs. Smith nodded. "But in AI, bias can lead to unfair outcomes. Let me show you how."

She introduced a new challenge where Cece and Max would train virtual robots to recommend meals for the class. The robots had to consider everyone's preferences and suggest the best meals. Cece's robot, named FairBot, and Max's robot, named Speedy, were ready to take on the challenge.

The students eagerly input their favorite meals into the system. Unsurprisingly, many chose pizza and chicken nuggets. FairBot and Speedy began processing the data, ready to make their recommendations.

Max smirked confidently. "Speedy will definitely come up with the best meals. It's fast and efficient."

In the quiet corners of the library, Cece stumbled upon a dusty old book titled "Unveiling Bias in Knowledge." Intrigued, she flipped through its pages and discovered stories of how biases had shaped historical narratives and scientific discoveries. Cece realized how crucial **fairness aware learning** would be in addressing these biases, and how Fairbot could prevent bias in its model by incorporating this learning type.

As the robots processed the data, Mrs. Smith explained more about bias in AI. "Bias in AI occurs when the training data doesn't represent everyone equally. This can happen if we have **selection bias**, where certain groups are overrepresented, or if we have **confirmation bias**, where we only see data that supports our expectations."

Name:

Date:

Cece nodded. "So if we only have data on pizza and chicken nuggets, the robots might think those are the only good meals."

"Exactly," Mrs. Smith said. "Let's see what happens."

Speedy quickly suggested pizza and chicken nuggets for every meal, while FairBot took a bit longer to make its recommendations. When FairBot finally spoke, it suggested a more balanced menu that included salad and spaghetti as well.

"FairBot considered everyone's preferences, even if they were in the minority," Cece explained proudly. "It's using a **diverse dataset** and avoiding bias."



Max frowned. "But why does it matter if most people want pizza and chicken nuggets?"

Mrs. Smith chimed in. "It matters because diversity is important. If we only recommend the most popular choices, we might miss out on other great options. Plus, fairness ensures that everyone's preferences are valued."

Max's eyes widened with realization. "So Speedy was biased because it overrepresented the most popular choices?"

"Exactly," Mrs. Smith affirmed. "FairBot used techniques to reduce bias, such as rebalancing the dataset and ensuring that less popular choices were still considered."

Cece added, "FairBot also used **probability distributions** to compare different groups' outcomes. It checked if the probability of recommending certain meals was unfairly different between groups."

Max nodded thoughtfully. "I see now. Bias in AI can make things unfair. We need to be careful about how we train our models."

Mrs. Smith continued, "Another important concept is **conditional probability**. This helps us understand how the probability of an outcome changes given certain attributes. For example, if FairBot noticed that the probability of recommending pizza for boys was much higher than for girls, that could indicate gender bias." Cece explained, "FairBot used conditional probability to ensure its recommendations were fair for everyone, regardless of their preferences. It adjusted its **learning process** to avoid overfitting and ensured a fair representation."

Max was impressed. "Wow, Cece! So by understanding probability distributions and conditional probability, we can detect and reduce bias in AI?"

Name:

Date:

"But I still don't understand what probability distributions are, or how we use them in our models," Max frowned.

"That's okay, I'll explain," Mrs. Smith said with a smile. She described how in the world of computers, these distributions are like graphs that show us how different factors, such as age or favorite hobbies, can affect the decisions the computers make. Then, Mrs. Smith displayed a large chart on the board showing everyone in the class's favorite foods. Cece leaned forward, curious about the colorful bars on the chart. Mrs. Smith helped them understand: to analyze these graphs, we closely look at how often each food choice appears. We're checking to see if some choices happen more often than others, which helps us find out if there's any unfairness. Cece realized how these ideas could help computers be fairer when making decisions. It amazed her how something as simple as counting and comparing could ensure technology is fair for everyone.

"It's all about making sure everyone is treated fairly and that the AI system learns from diverse data", Cece added.

The class cheered as Cece and Max shared a high-five. They had learned an important lesson about AI and bias, and understood the importance of fairness in their work. As they left the library, Cece and Max excitedly discussed their next project. They knew that by combining their strengths and addressing biases, they could create even better AI systems. With a newfound respect for fairness, they embarked on their next adventure, eager to tackle any challenge that came their way.

Glossary

- 1.**Bias:** Unfair preference or prejudice in AI models.
- 2.**Selection Bias:** When the training data overrepresents certain groups.
- 3.**Confirmation Bias:**When only data supporting a pre-existing belief is considered.
- 4.**Probability Distributions:** Comparing outcomes of different groups within the dataset to identify bias.
- 5.**Conditional Probability:** Understanding how the probability of an outcome changes given certain attributes.
6. **Fairness-Aware Learning:**Techniques to minimize bias and ensure fairness in AI.
- 7.**Diverse Dataset:** A collection of data that represents a wide range of preferences and characteristics.