AIBridge

Lecture 4

Unsupervised Learning

Species #1?

Species #2?

Unsupervised Learning

Clustering

Dimension reduction

ChatGPT

Clustering: Google news

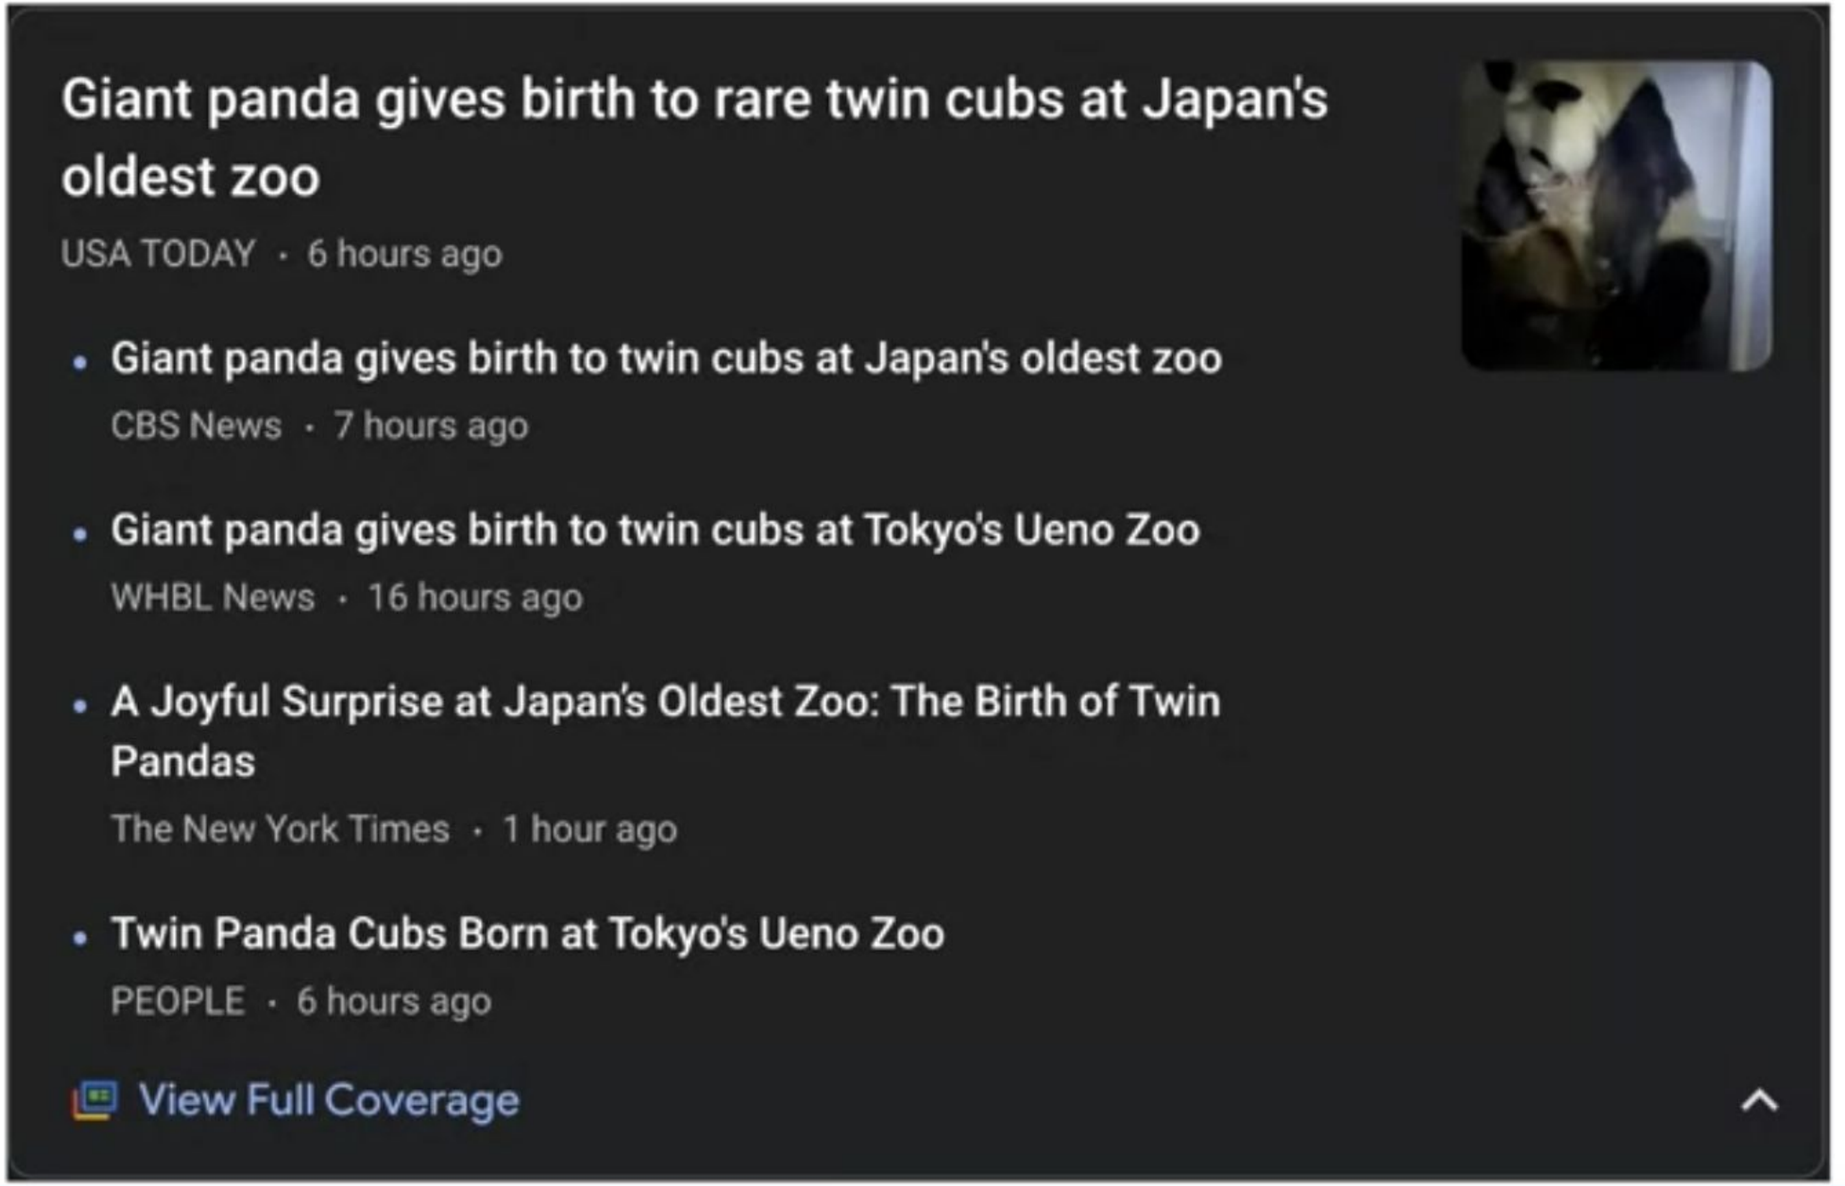

Credit: Andrew Ng, Machine Learning

k-means clustering

how do we find these clusters?

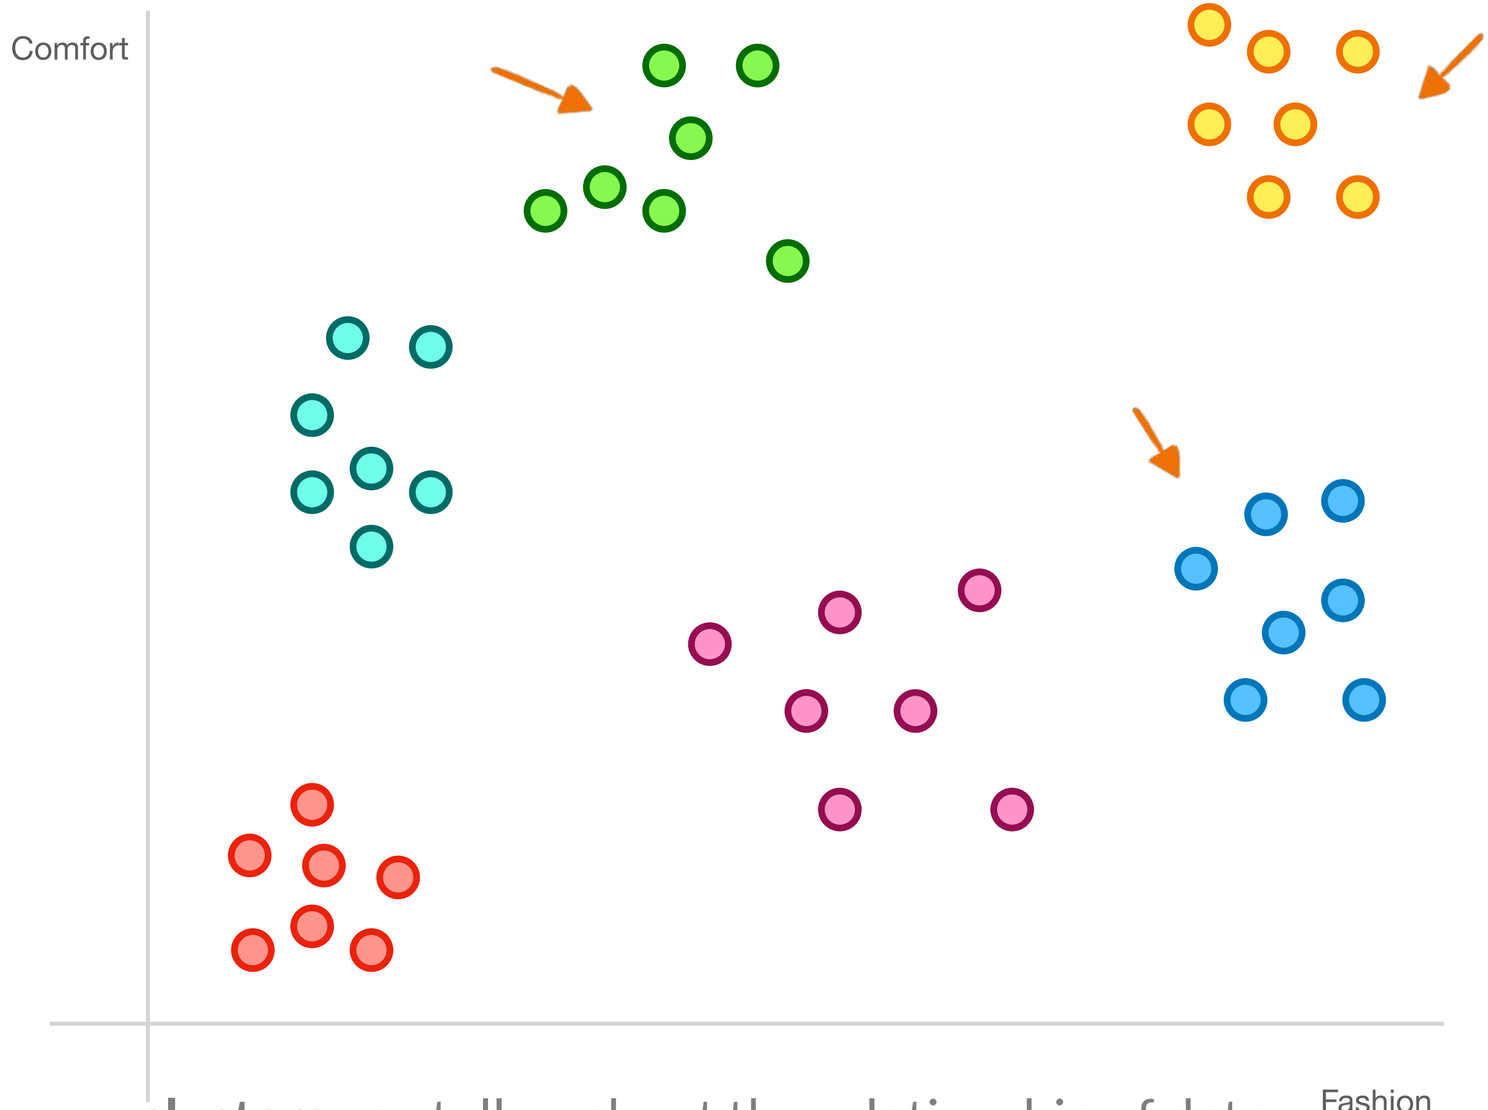

clusters can tell us about the relationship of data

...even if they are unlabeled!

→ **unsupervised learning!**

k-means clustering

1. pick a K-number of clusters
2. randomly pick a series of “**centroids**”
3. assign each particle to the **centroid** closest to it

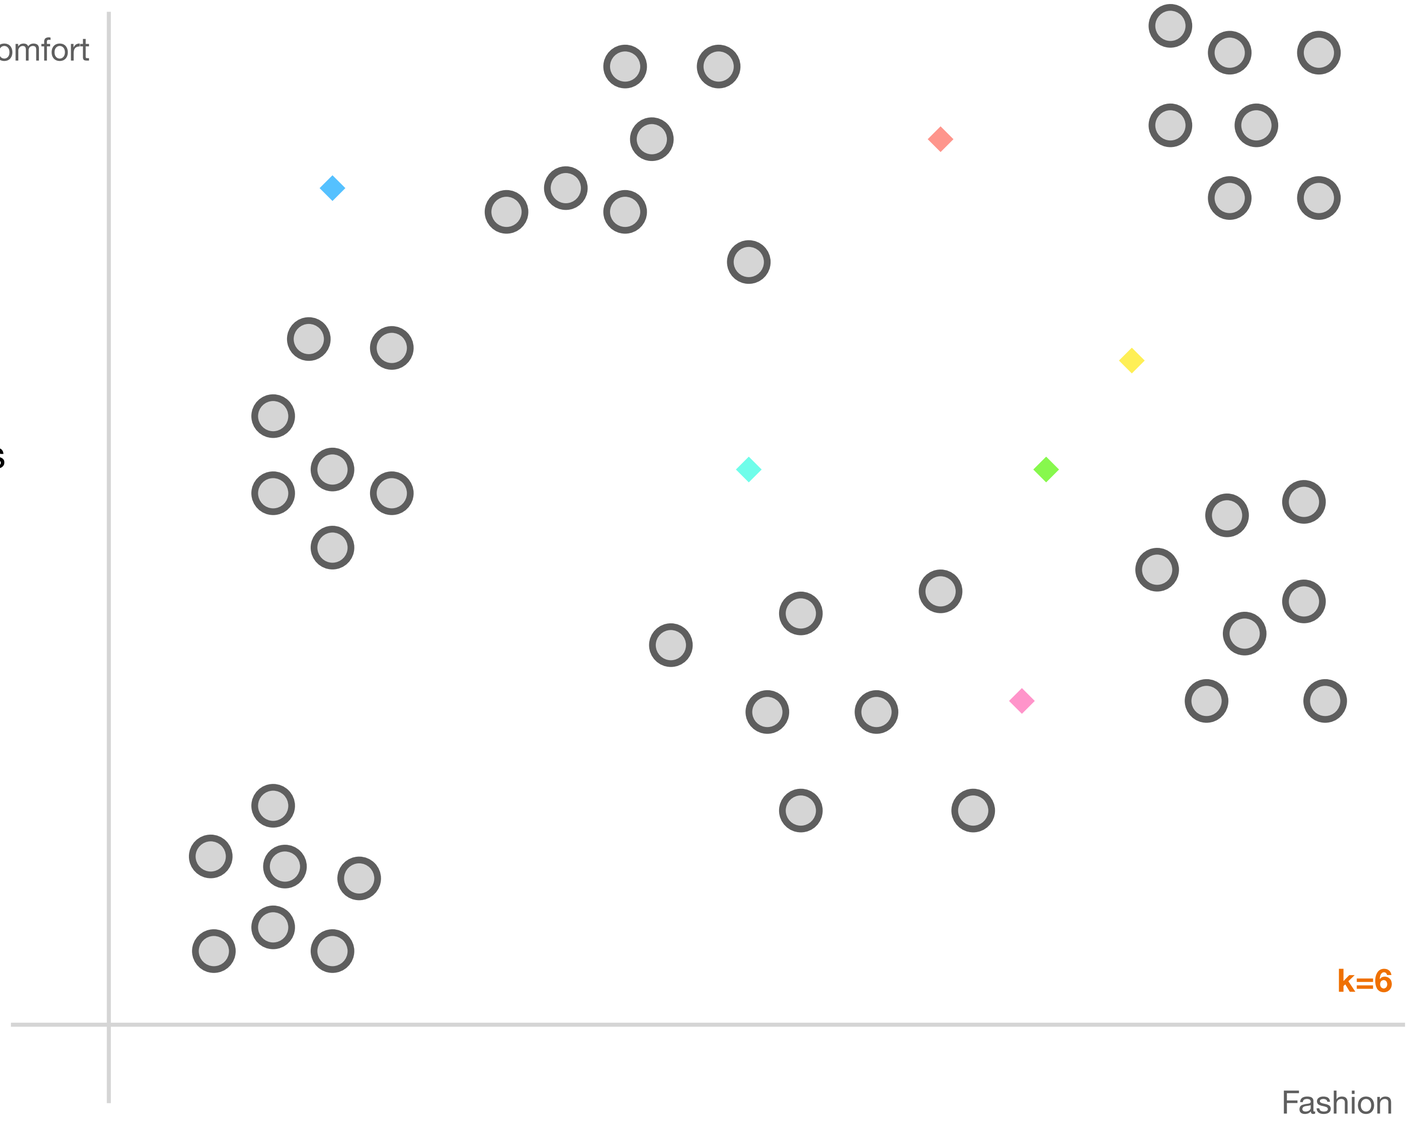

k-means clustering

1. pick a K-number of clusters
2. randomly pick a series of “**centroids**”
3. assign each particle to the **centroid** closest to it
4. move the **centroid** to the weighted geometric center of samples assigned to it

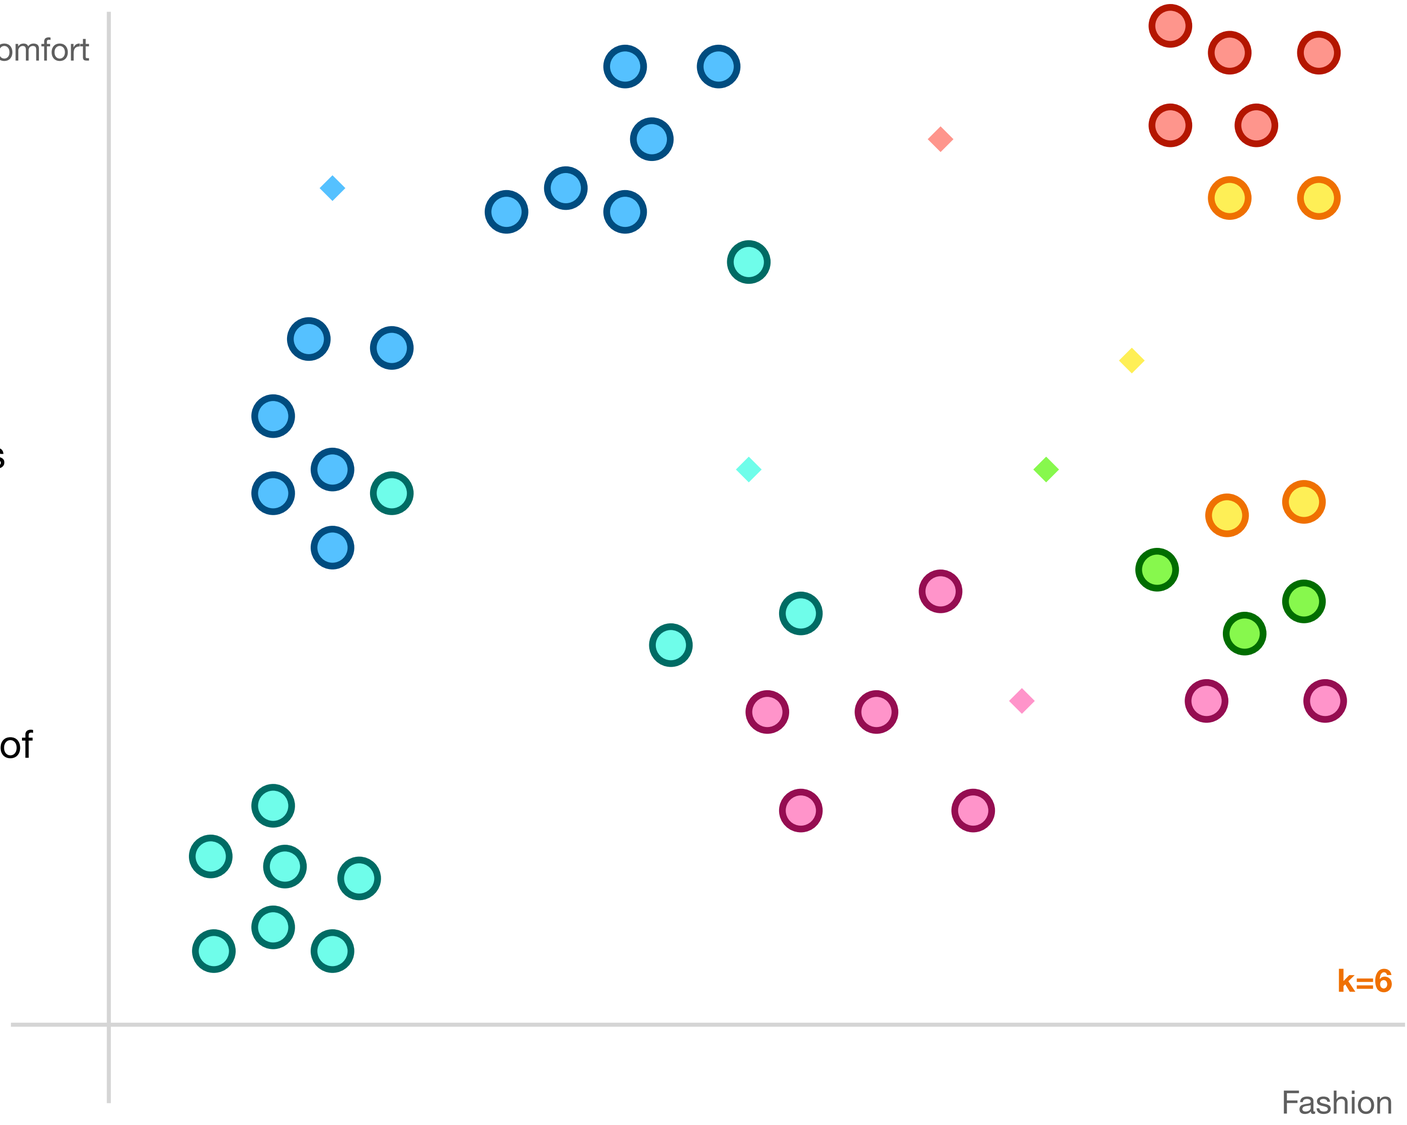

k-means clustering

1. pick a K-number of clusters
2. randomly pick a series of "**centroids**"
3. assign each particle to the **centroid** closest to it
4. move the **centroid** to the weighted geometric center of samples assigned to it
5. Repeat 3-4 until centroids stop moving!

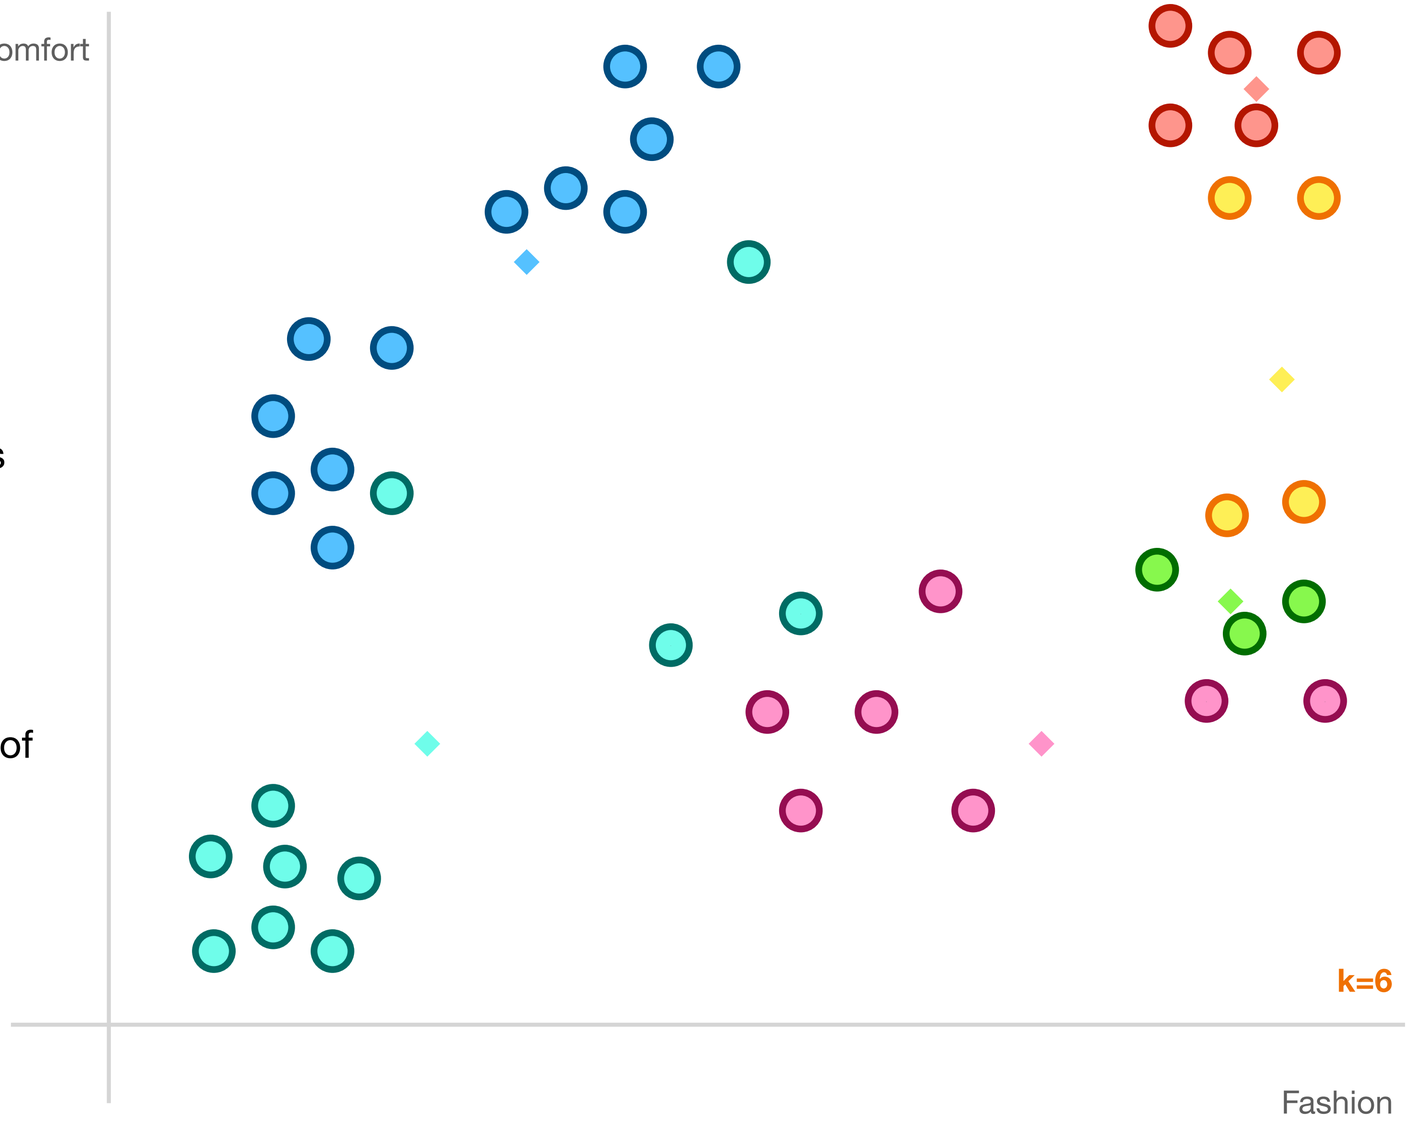

k-means clustering

1. pick a K-number of clusters
2. randomly pick a series of "**centroids**"
3. assign each particle to the **centroid** closest to it
4. move the **centroid** to the weighted geometric center of samples assigned to it
5. Repeat 3-4 until centroids stop moving!

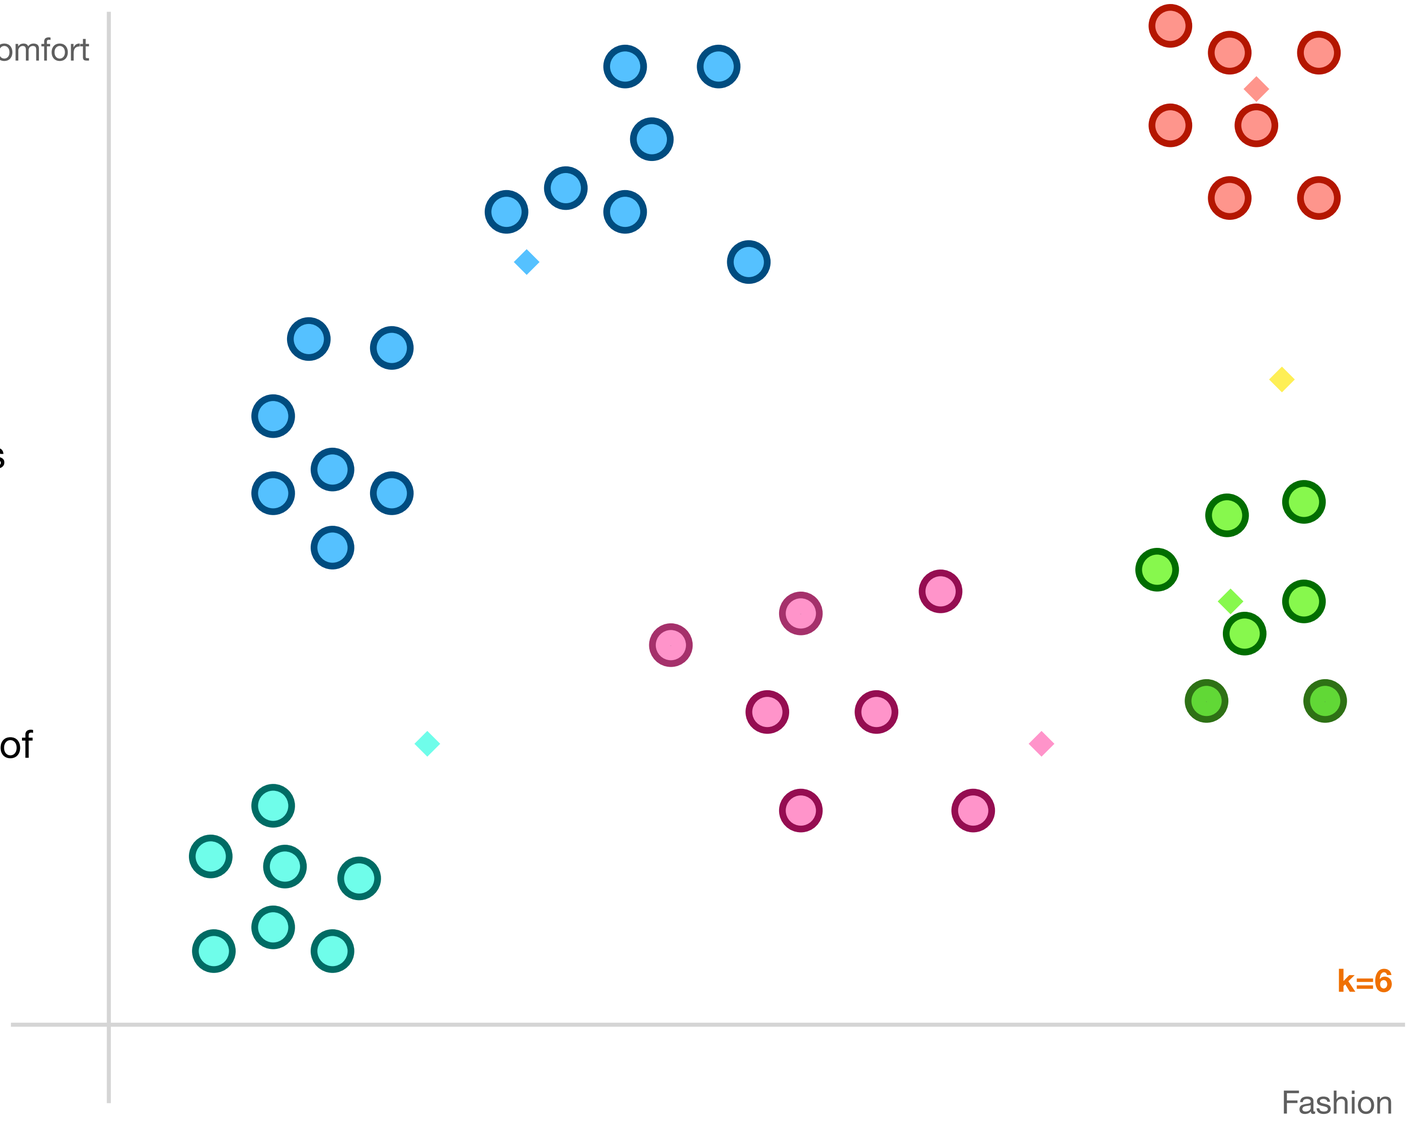

k-means clustering

1. pick a K-number of clusters
2. randomly pick a series of "**centroids**"
3. assign each particle to the **centroid** closest to it
4. move the **centroid** to the weighted geometric center of samples assigned to it
5. Repeat 3-4 until centroids stop moving!

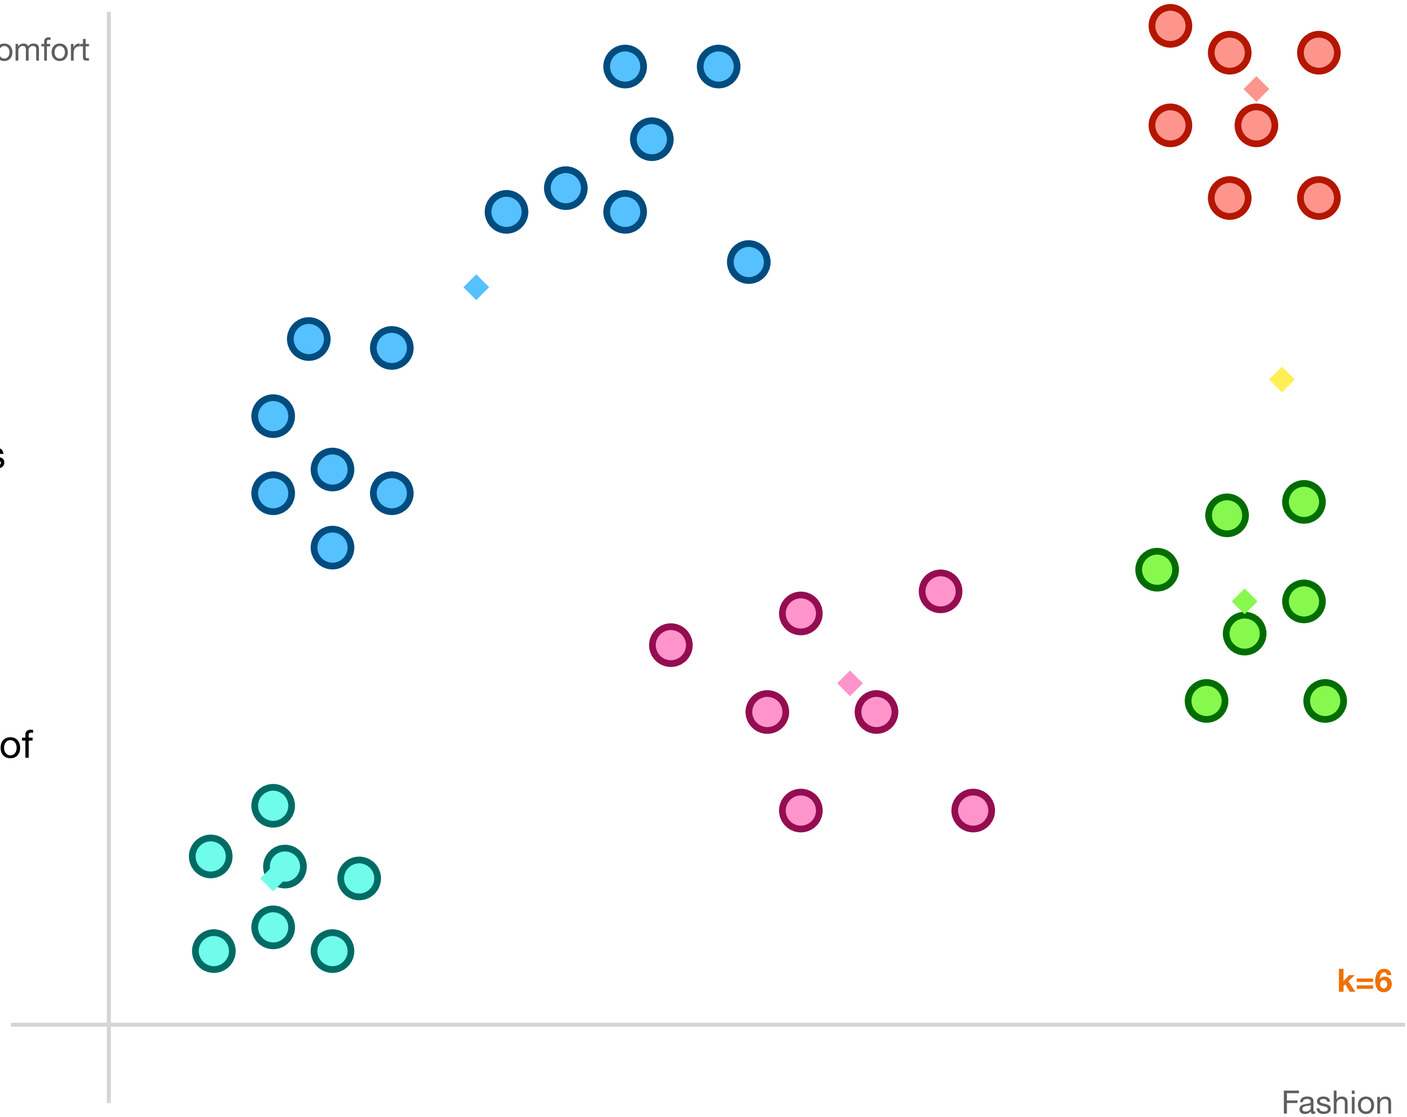

Do we always get back the same clusters?

Nope. And that's OK.

Do we always get back the same clusters?

Nope. And that's OK.

- **K-means** is an *indeterminisitic* algorithm—it has built-in randomness

Unsupervised Learning

Clustering

Dimension reduction

ChatGPT

Principle Component Analysis

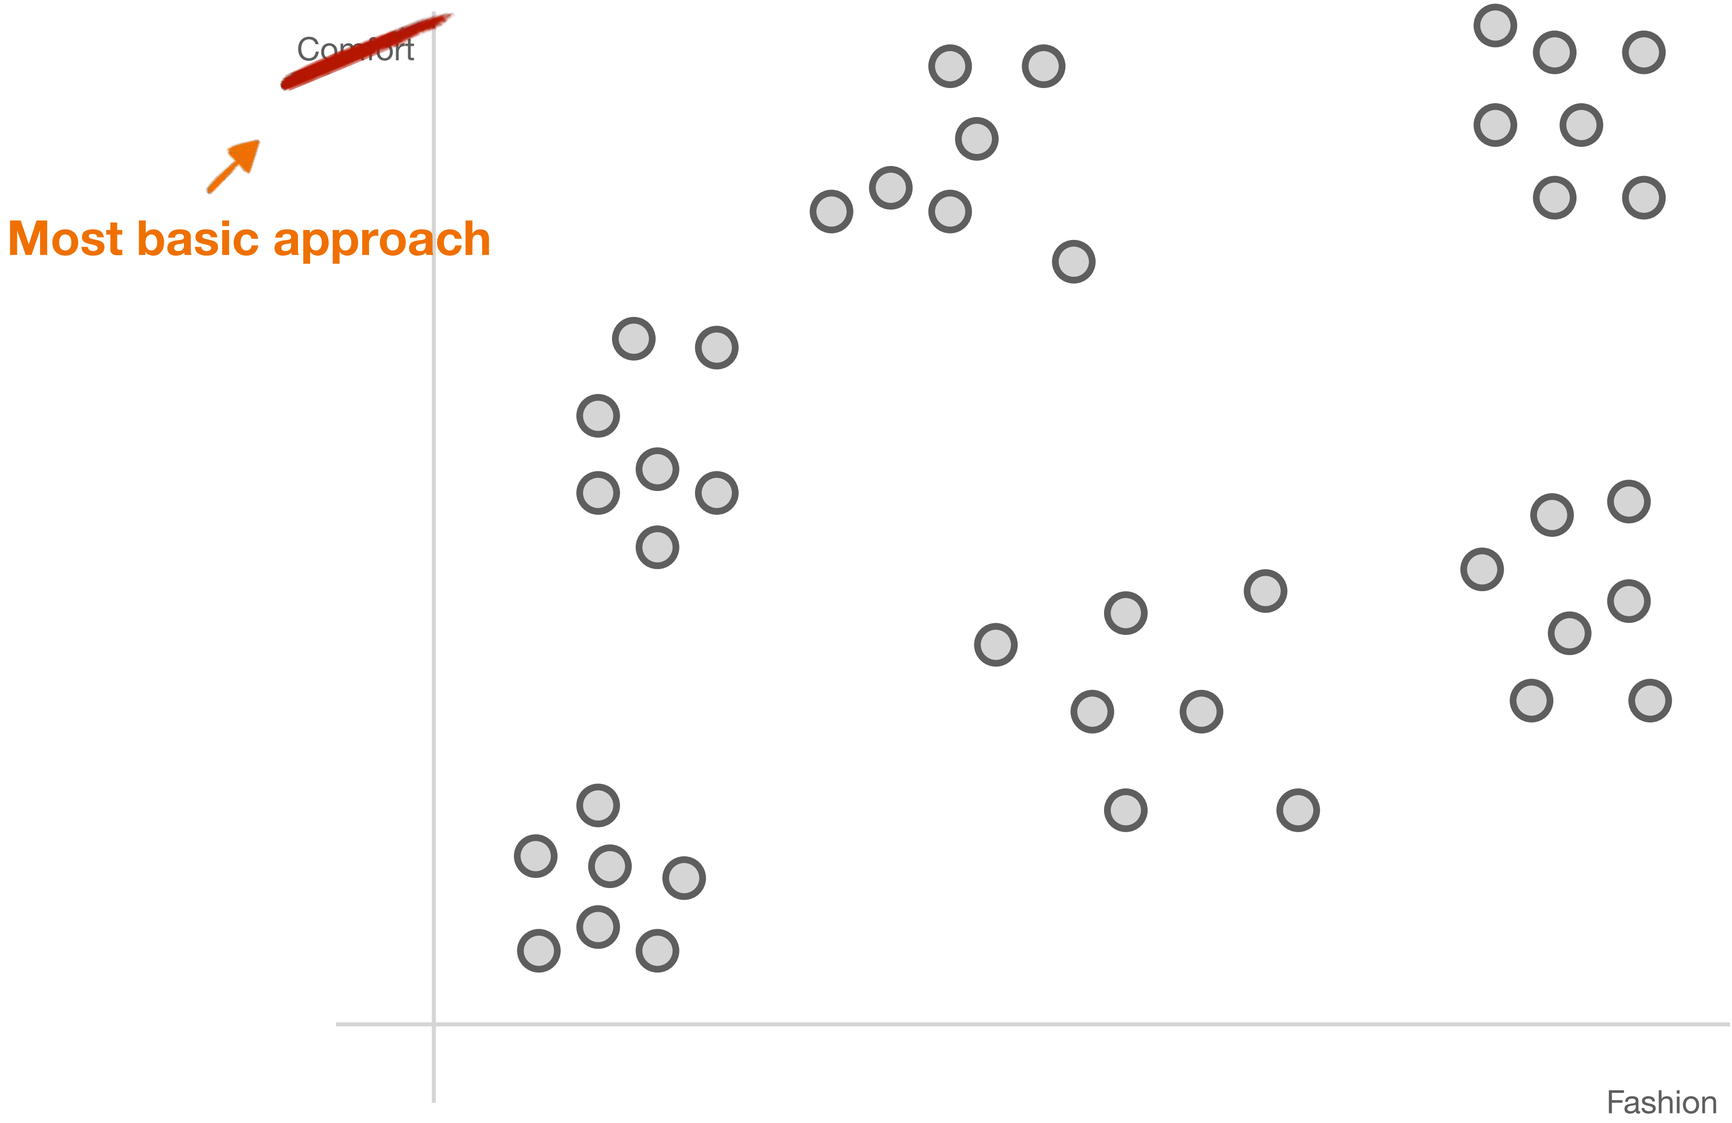

Principle Component Analysis

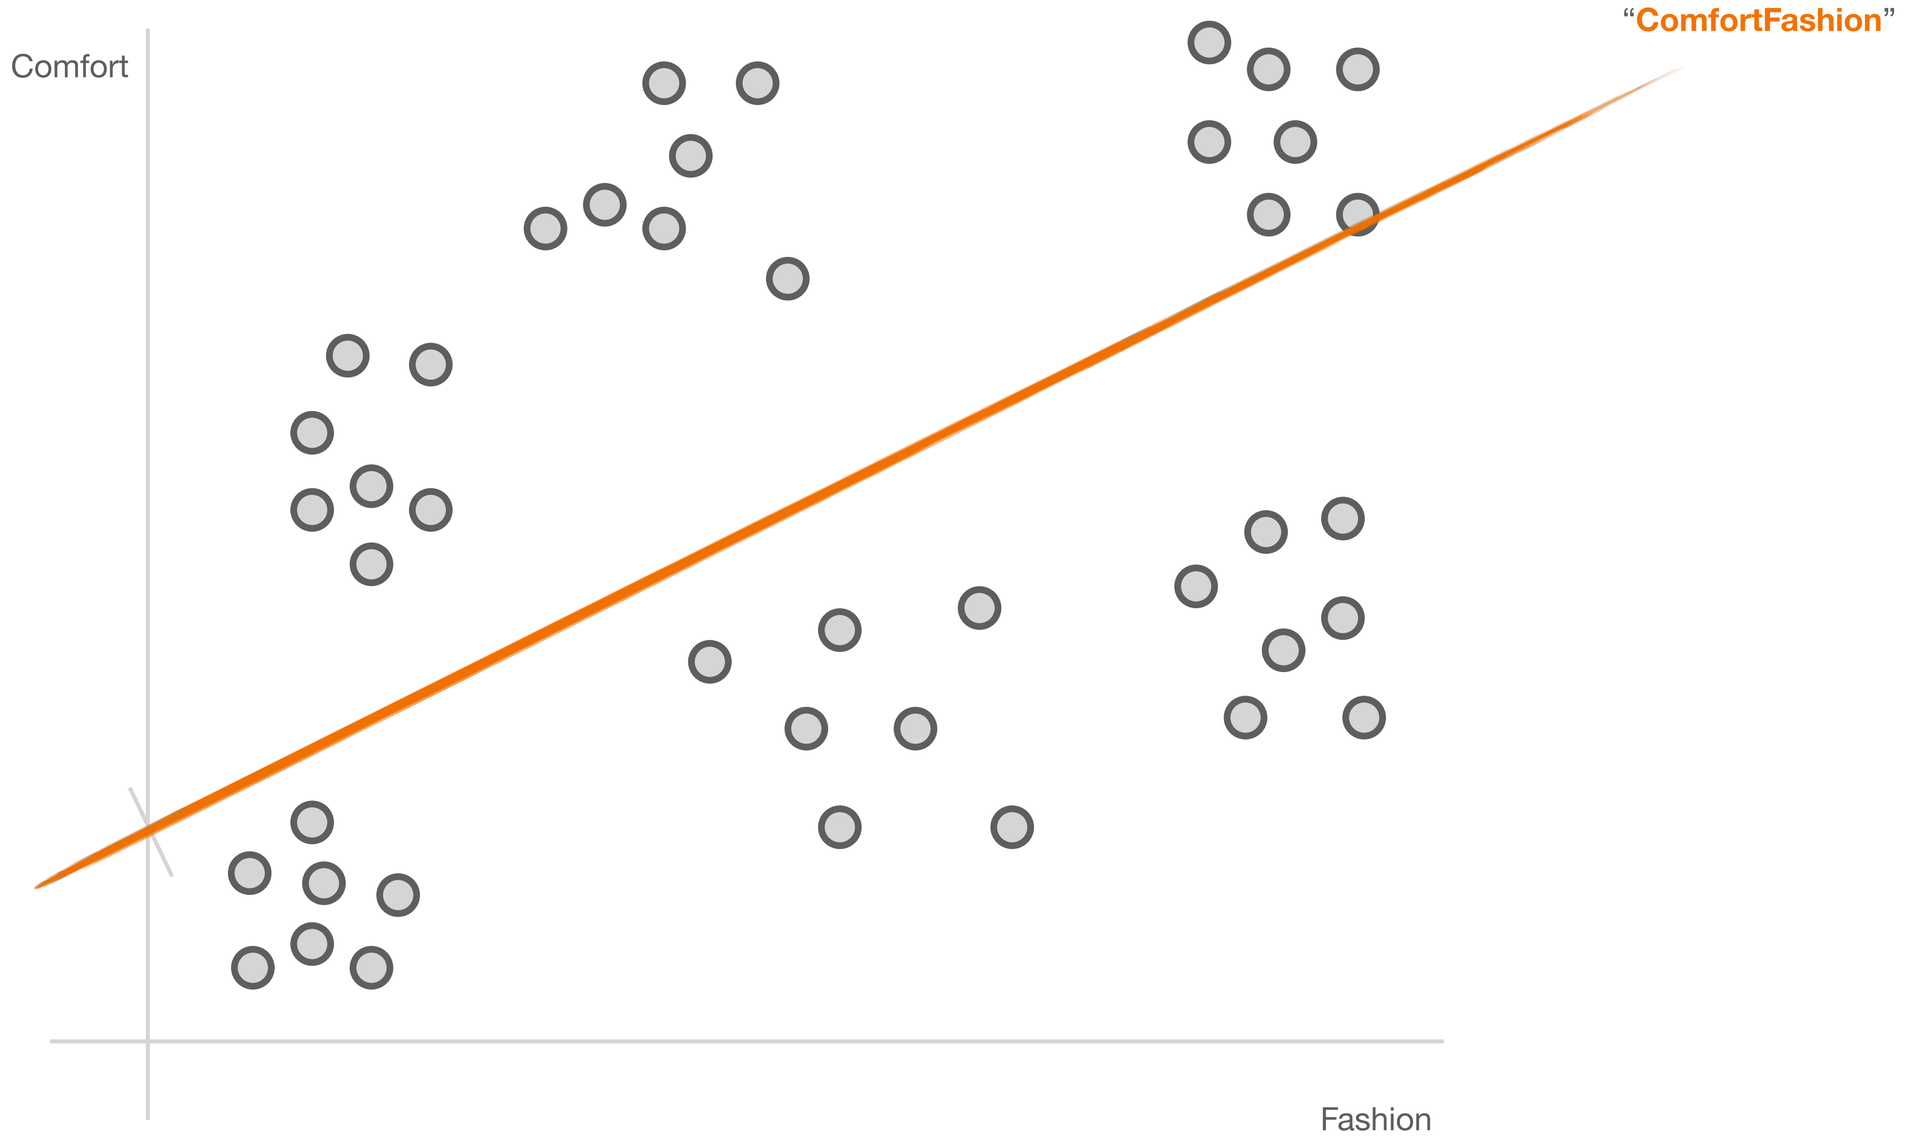

Principle Component Analysis

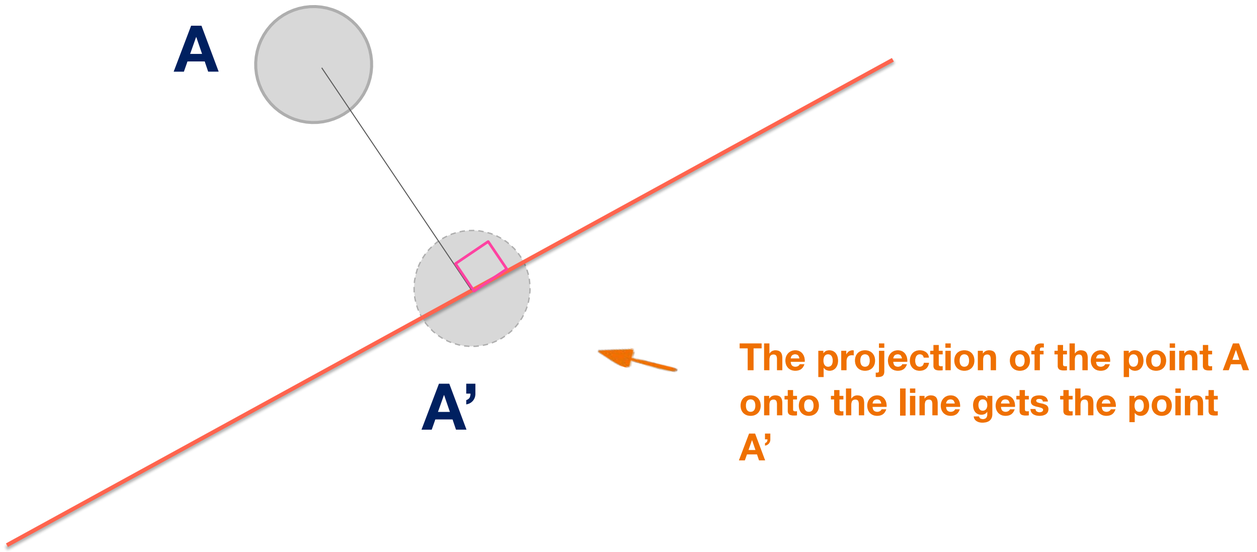

Principle Component Analysis

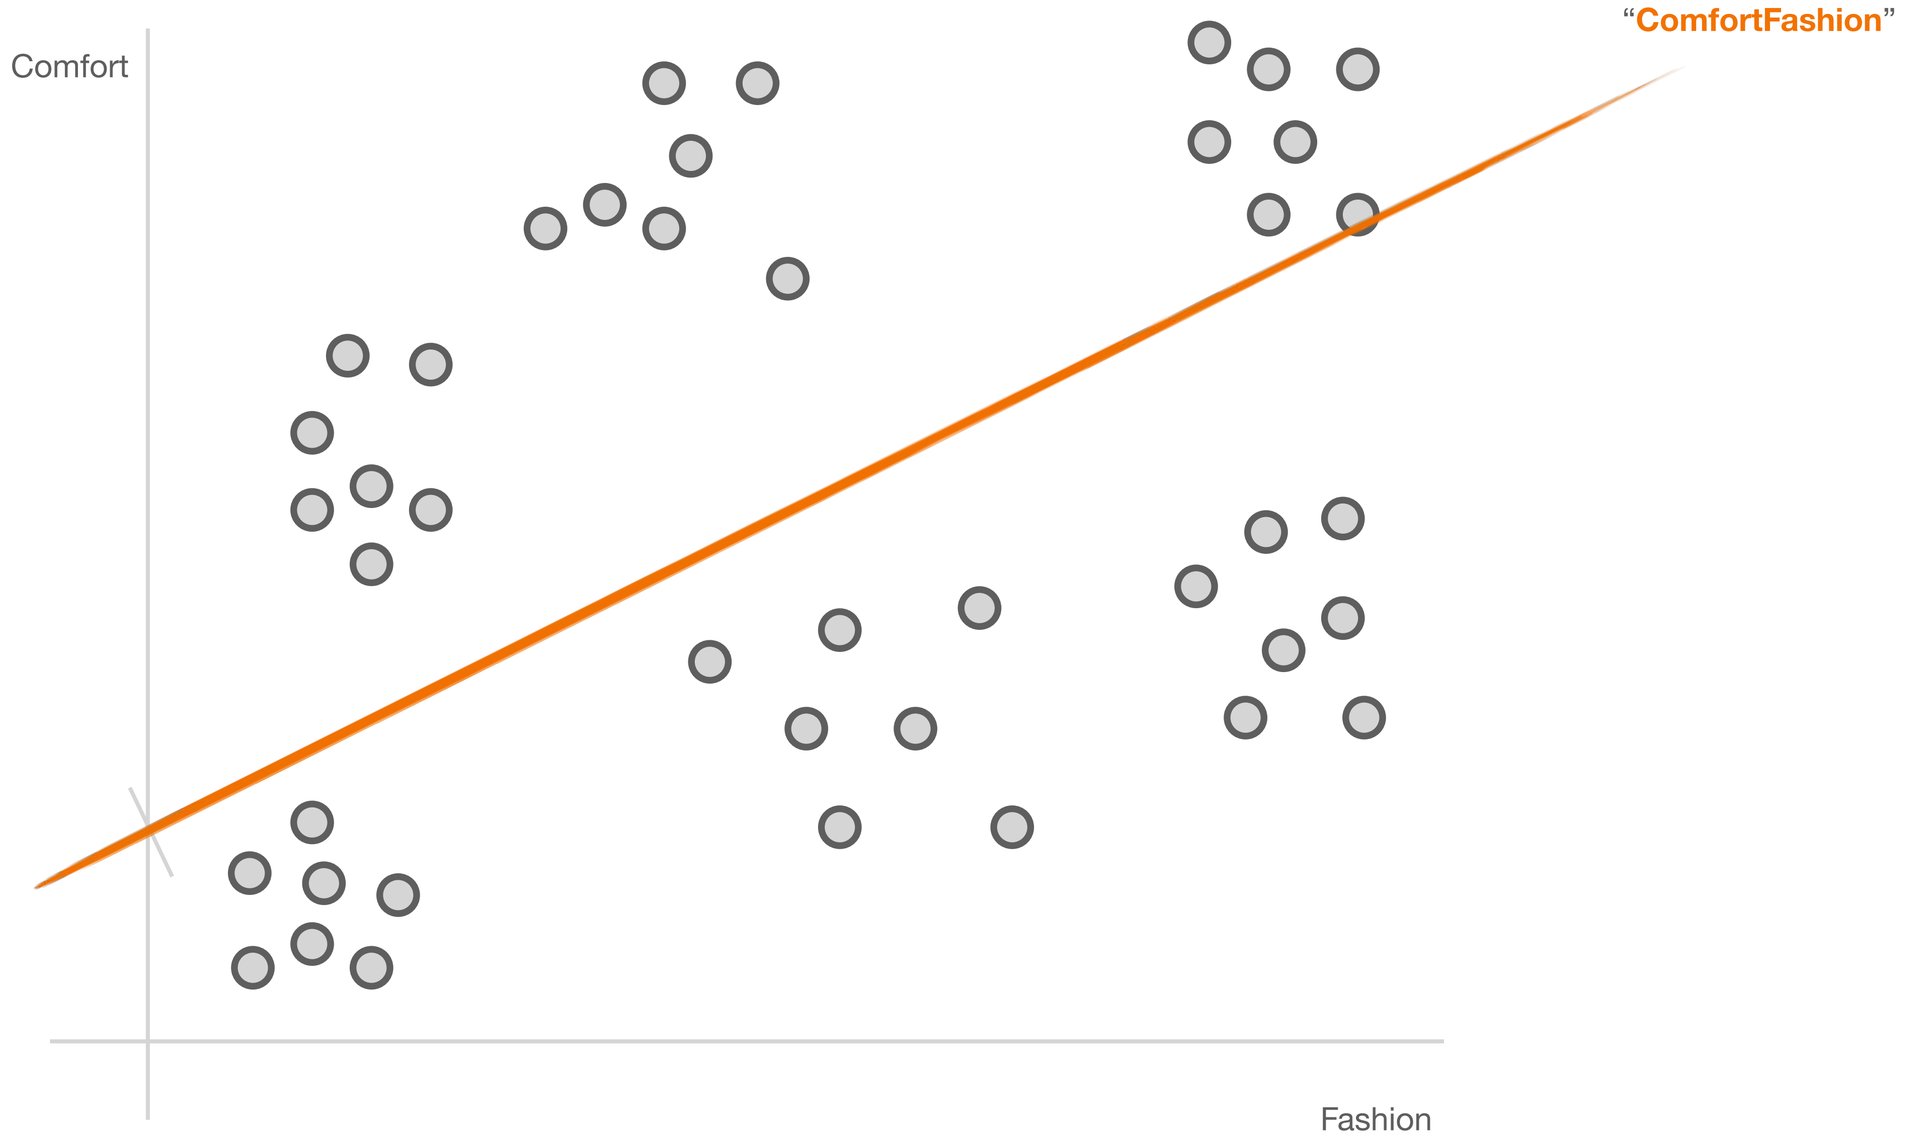

Principle Component Analysis

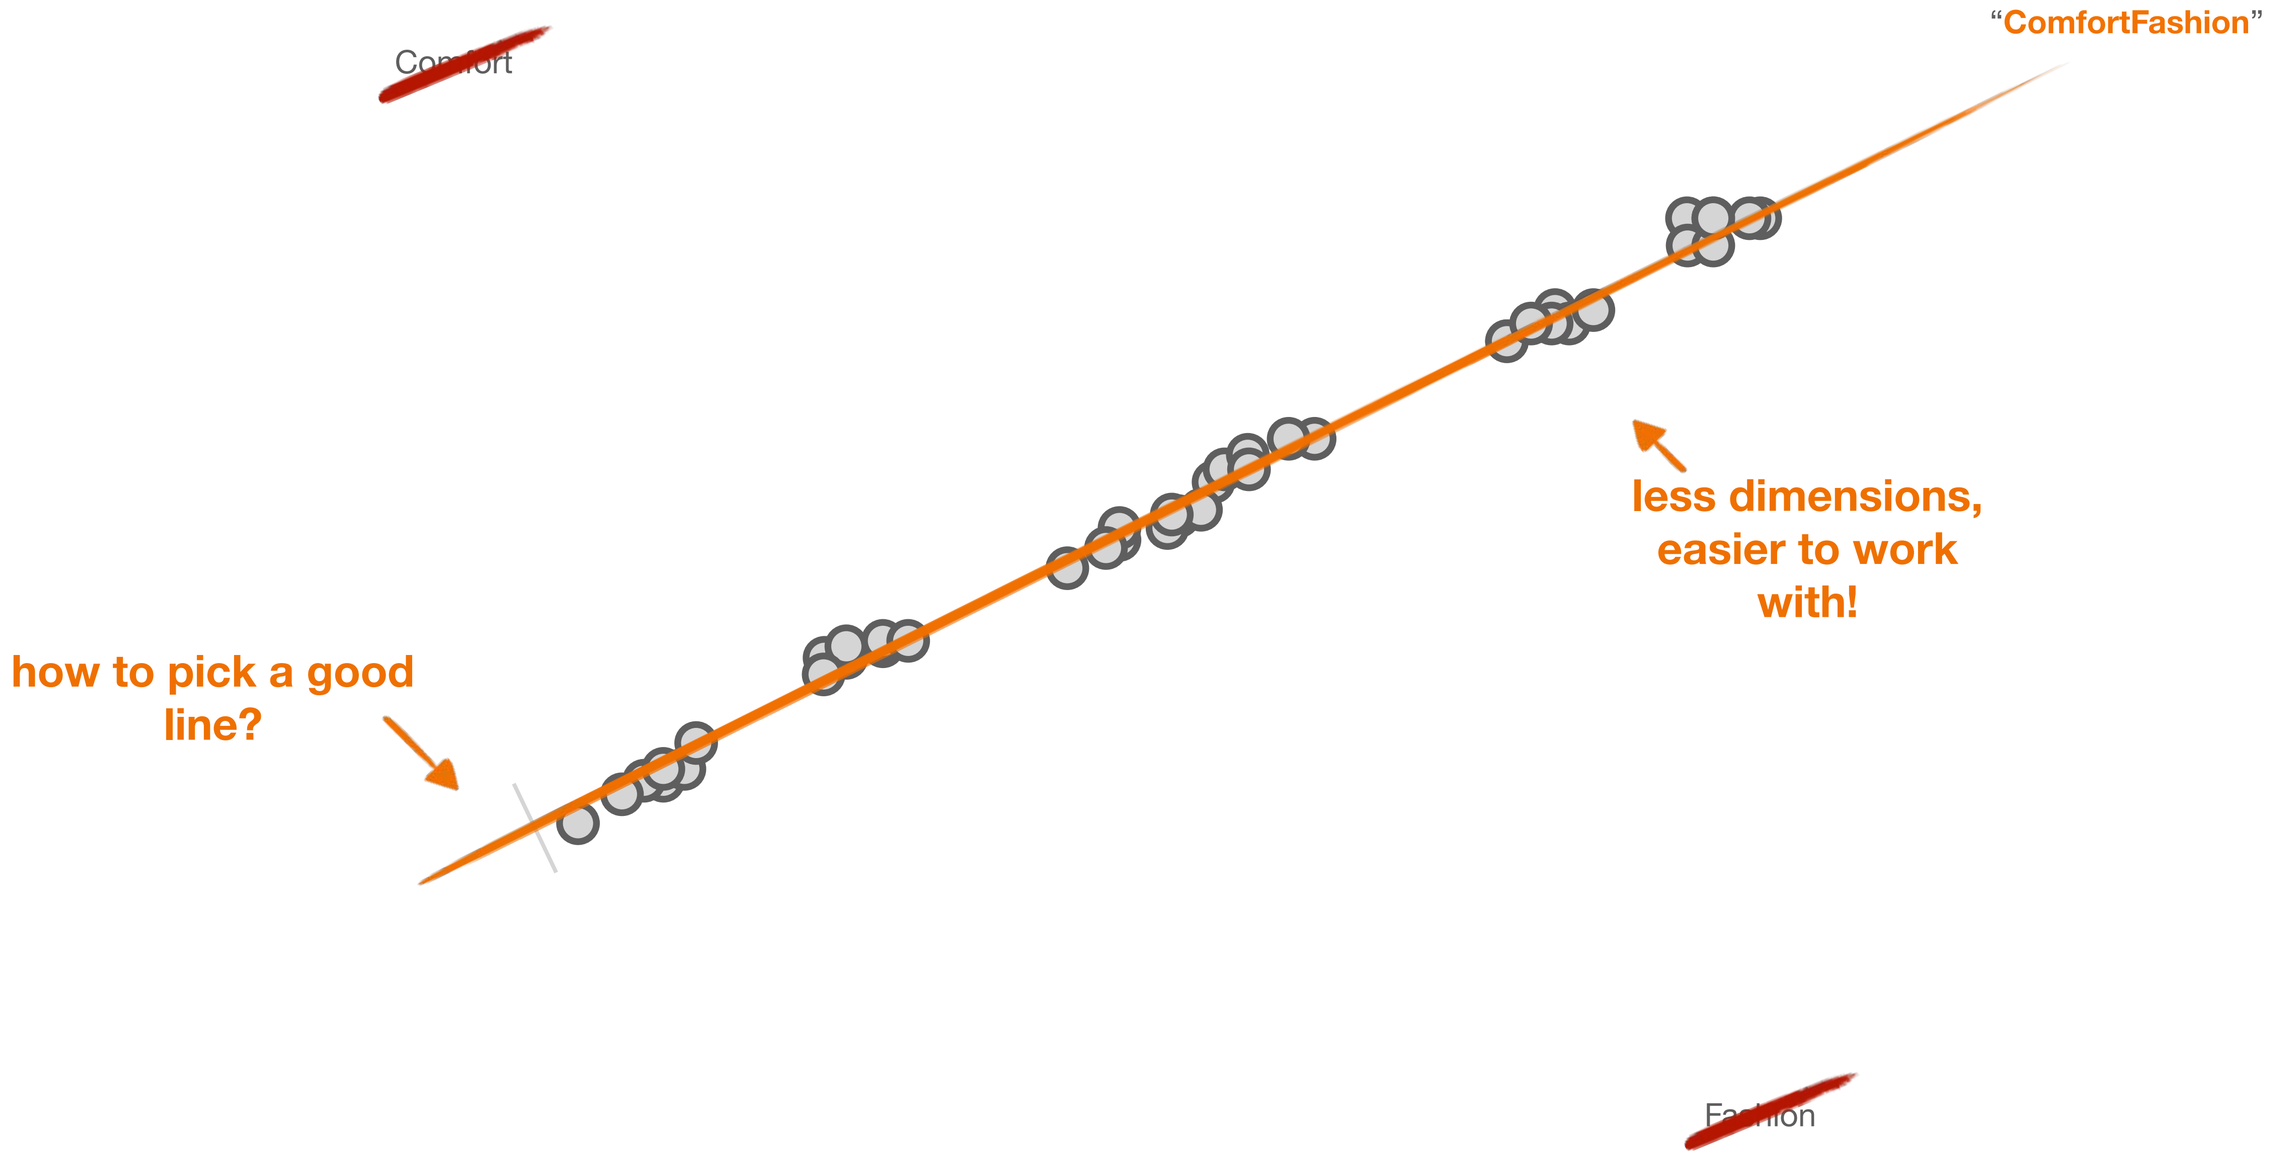

- by **projecting** the samples down to a smaller dimension, they are easier to work with.

Principle Component Analysis

how to pick a good line?

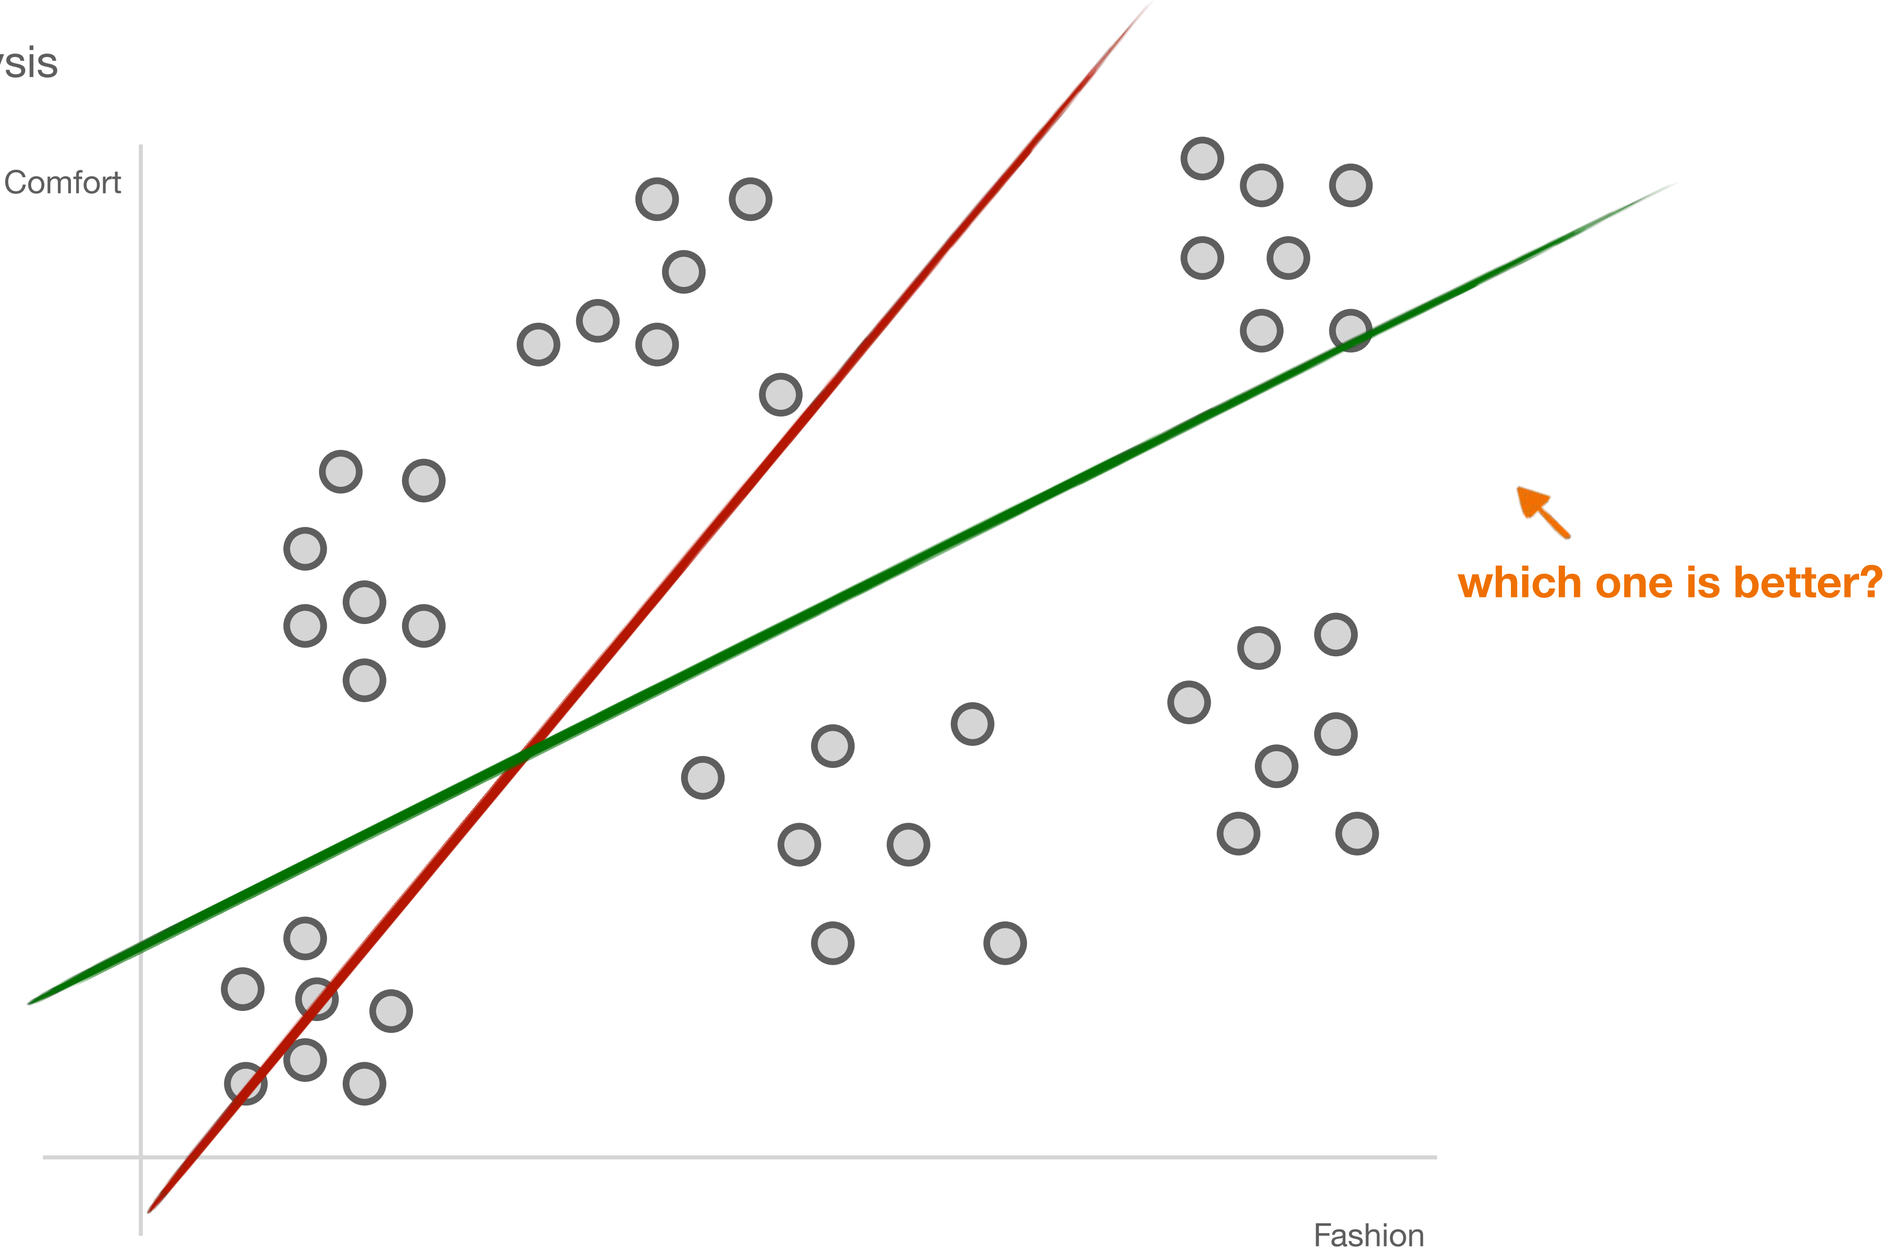

Principle Component Analysis

how to pick a good line?

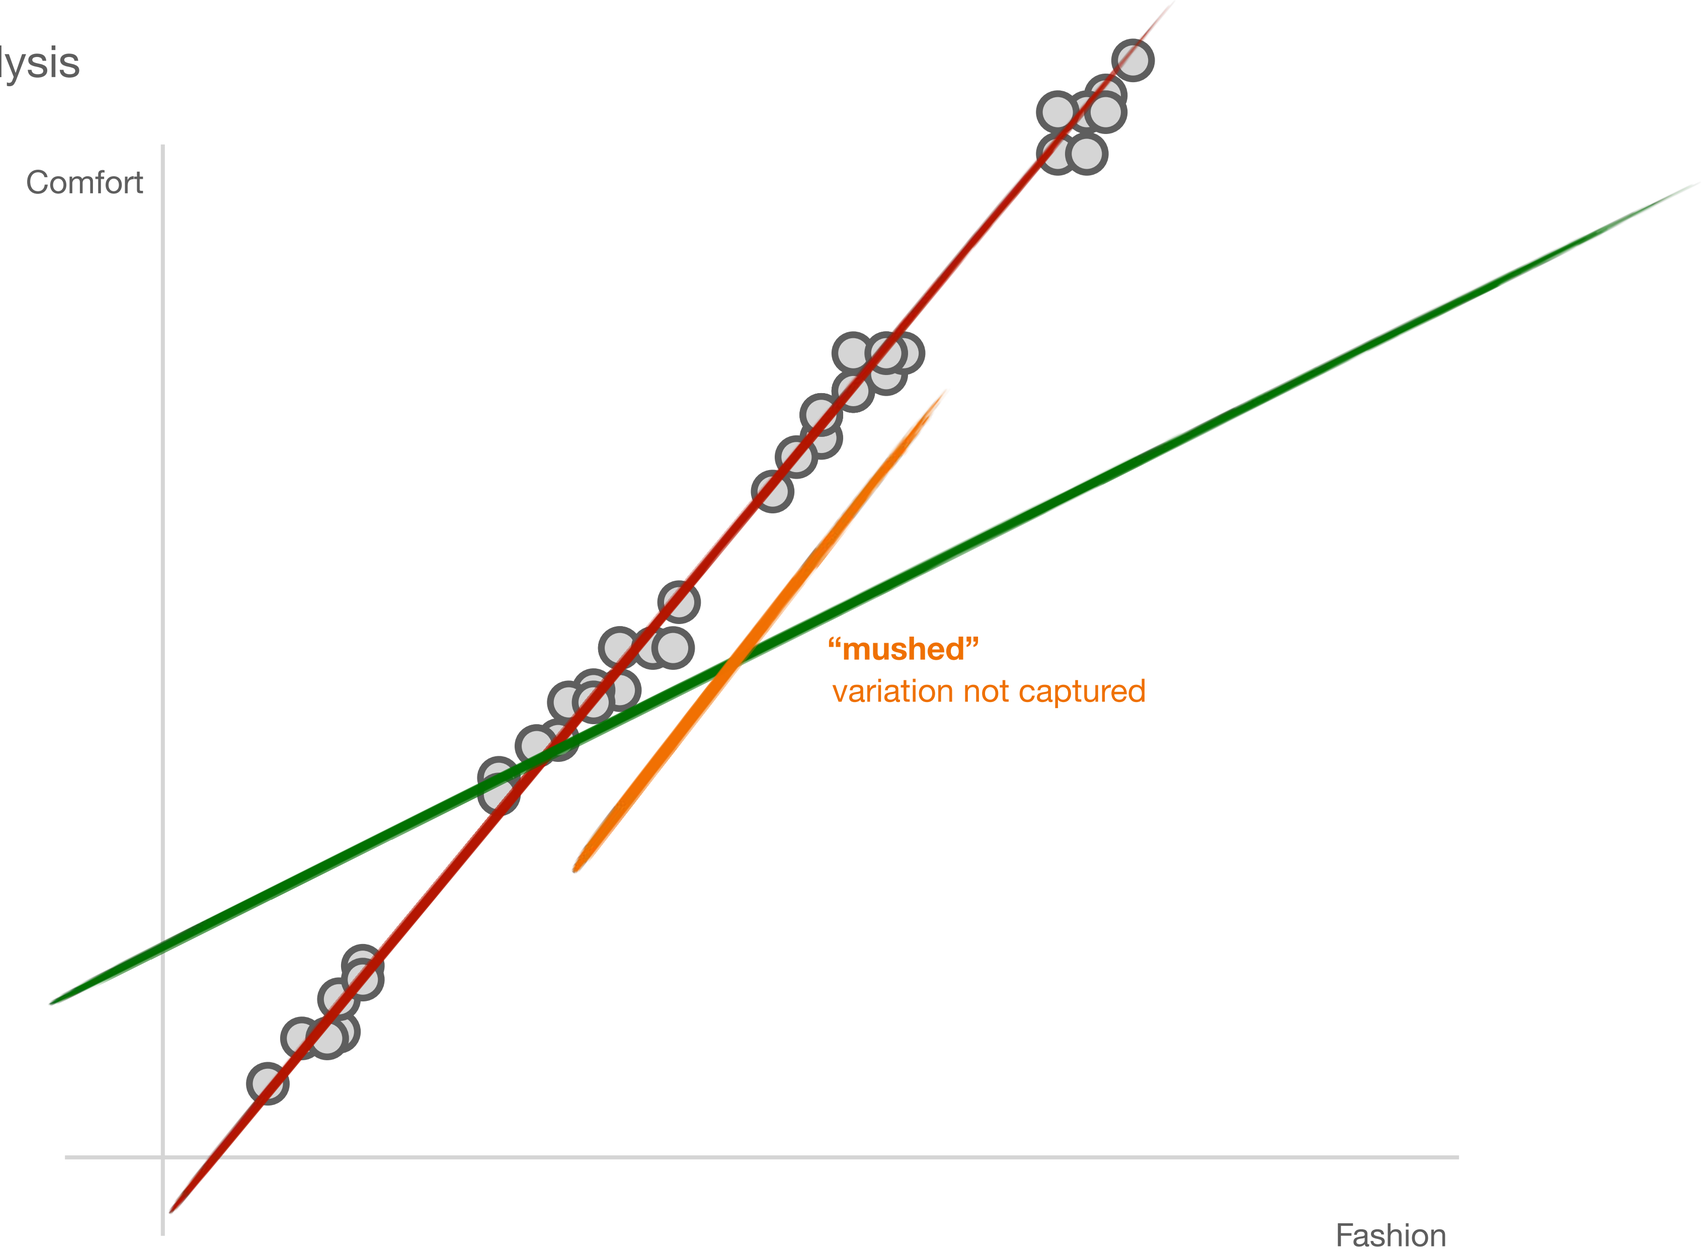

- a good "**projection**" captures the **variation** in the data

Principle Component Analysis

how to pick a good line?

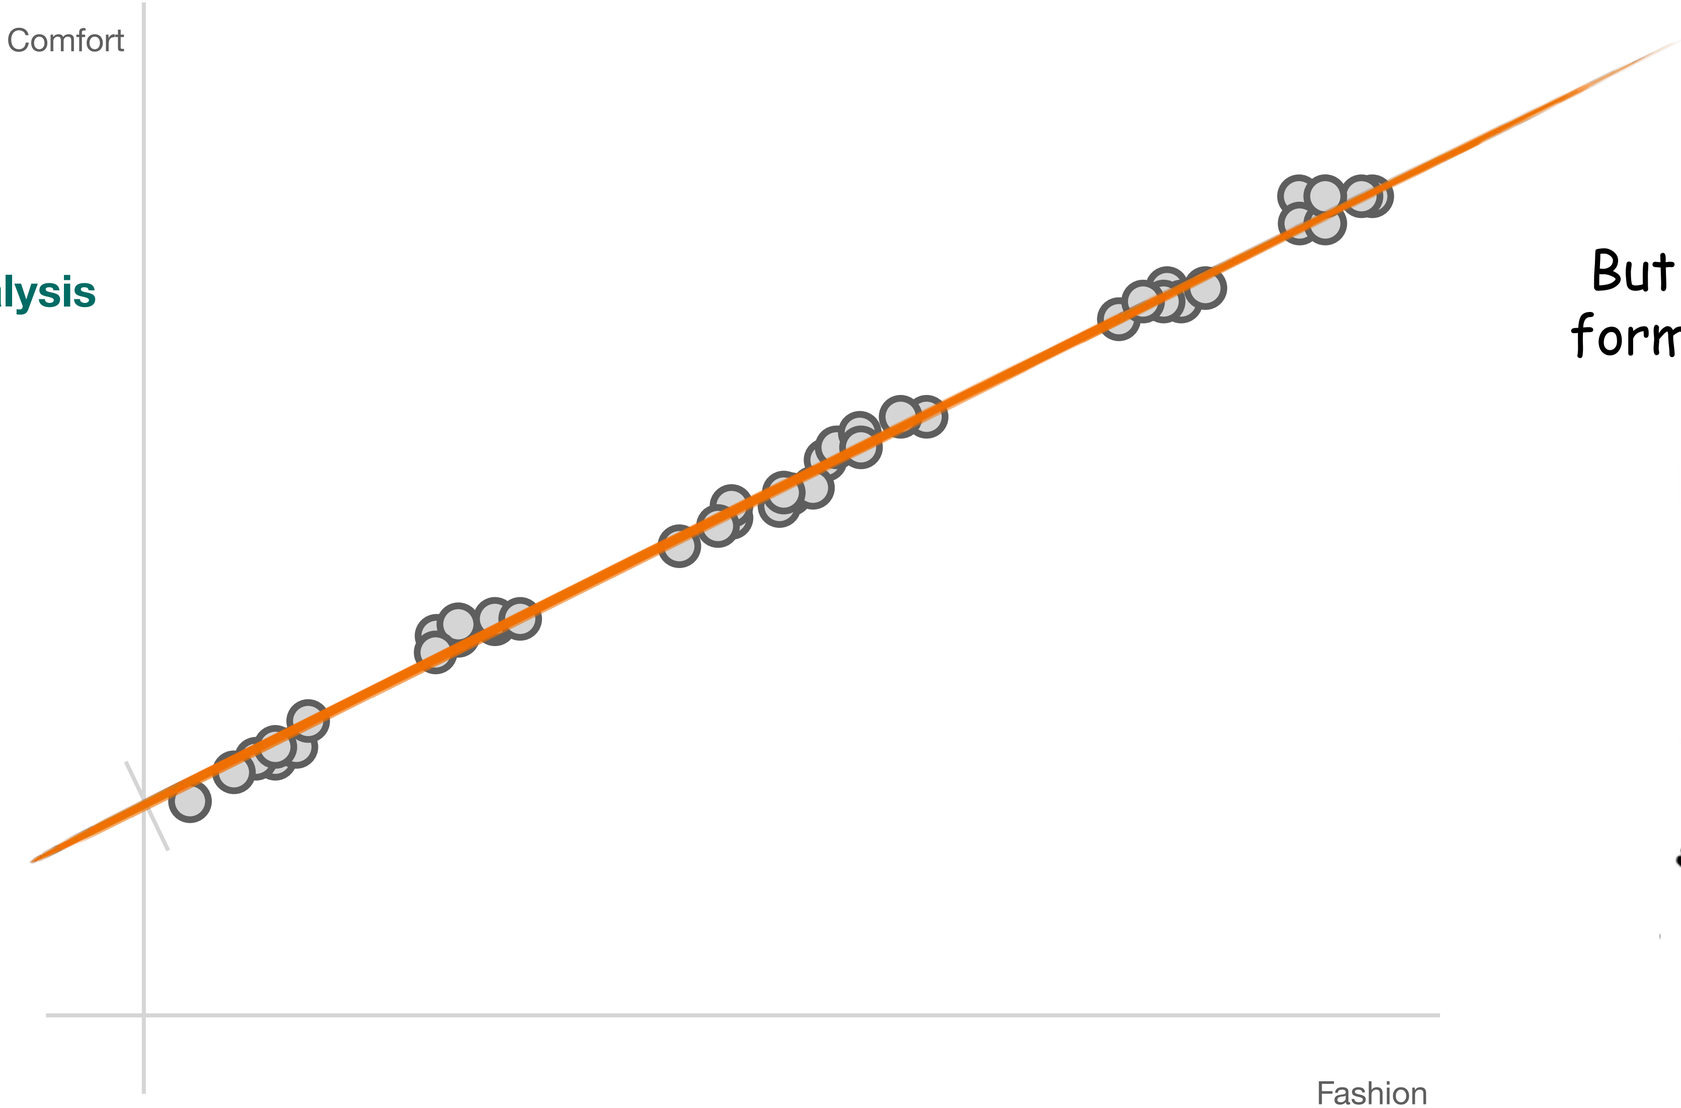

principle component analysis

- Find a good line (**basis**) that **maximizes variation**
- **Project** samples down

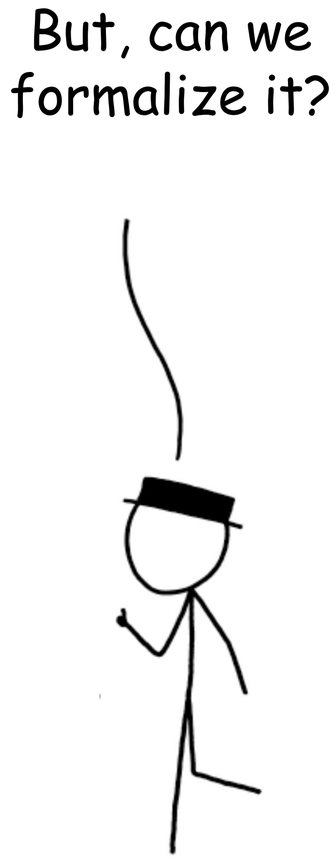

■

So you’ve undoubtedly heard about ChatGPT…

How ChatGPT Works

Training ChatGPT

Using ChatGPT

Training ChatGPT

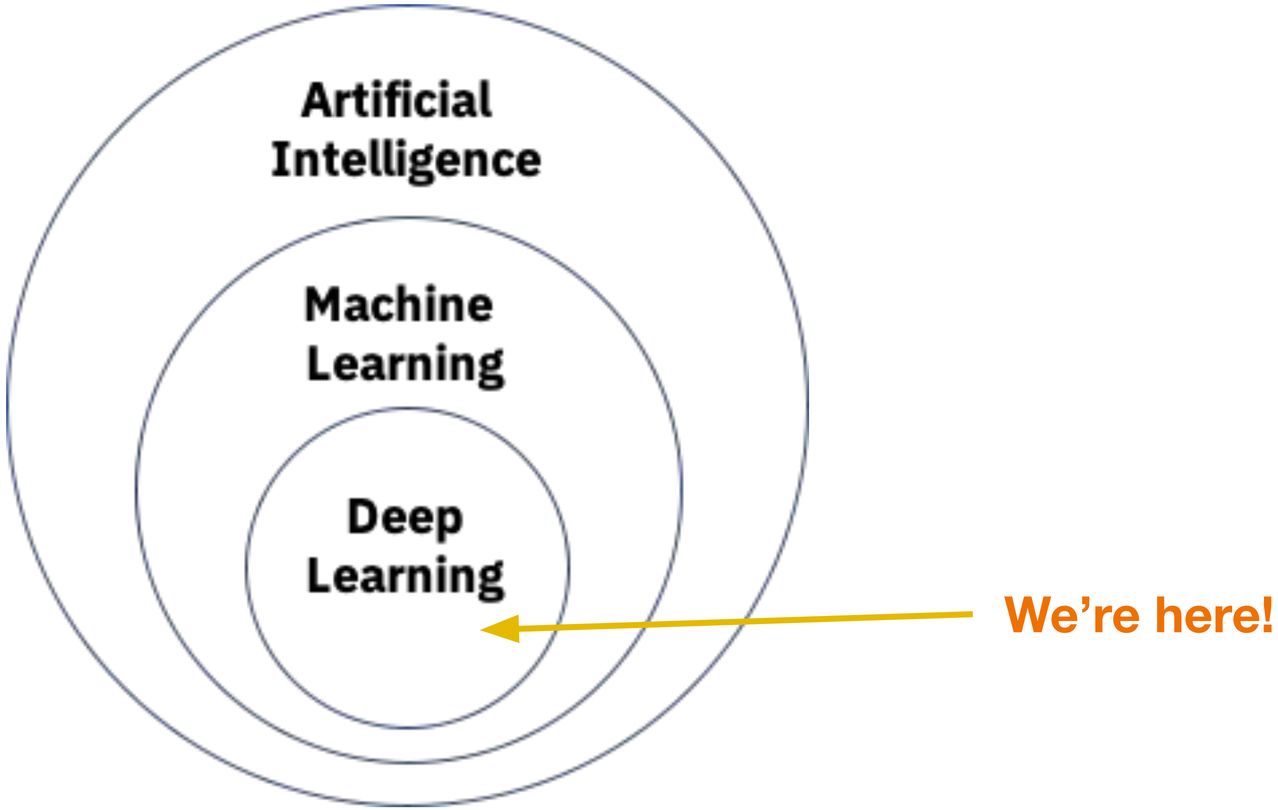

- ChatGPT's goal: **Generating** new, human-like text for conversations

Training ChatGPT: The Beginning

- ChatGPT's goal: **Generating** new, human-like text for conversations
- 

"Generative Pretrained Transformers"

- Large Language Model (LLM)
- *Next-Token Prediction*

Large language models

- **Large**
  - Large training dataset
  - Large number of parameters
- **General purpose**
  - Commonality of human languages
  - Resource restriction
- **Pre-trained and fine-tuned**

Training ChatGPT: The Beginning

Training ChatGPT: The Beginning

GPT

Text input

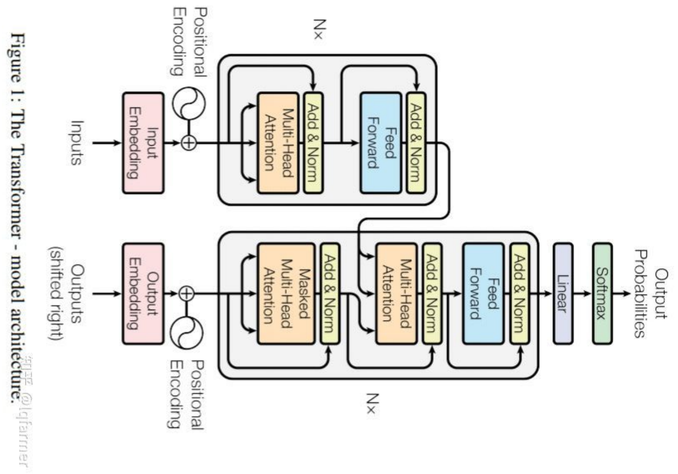

Figure 1: The Transformer - model architecture.

Next word

Training ChatGPT: The Beginning

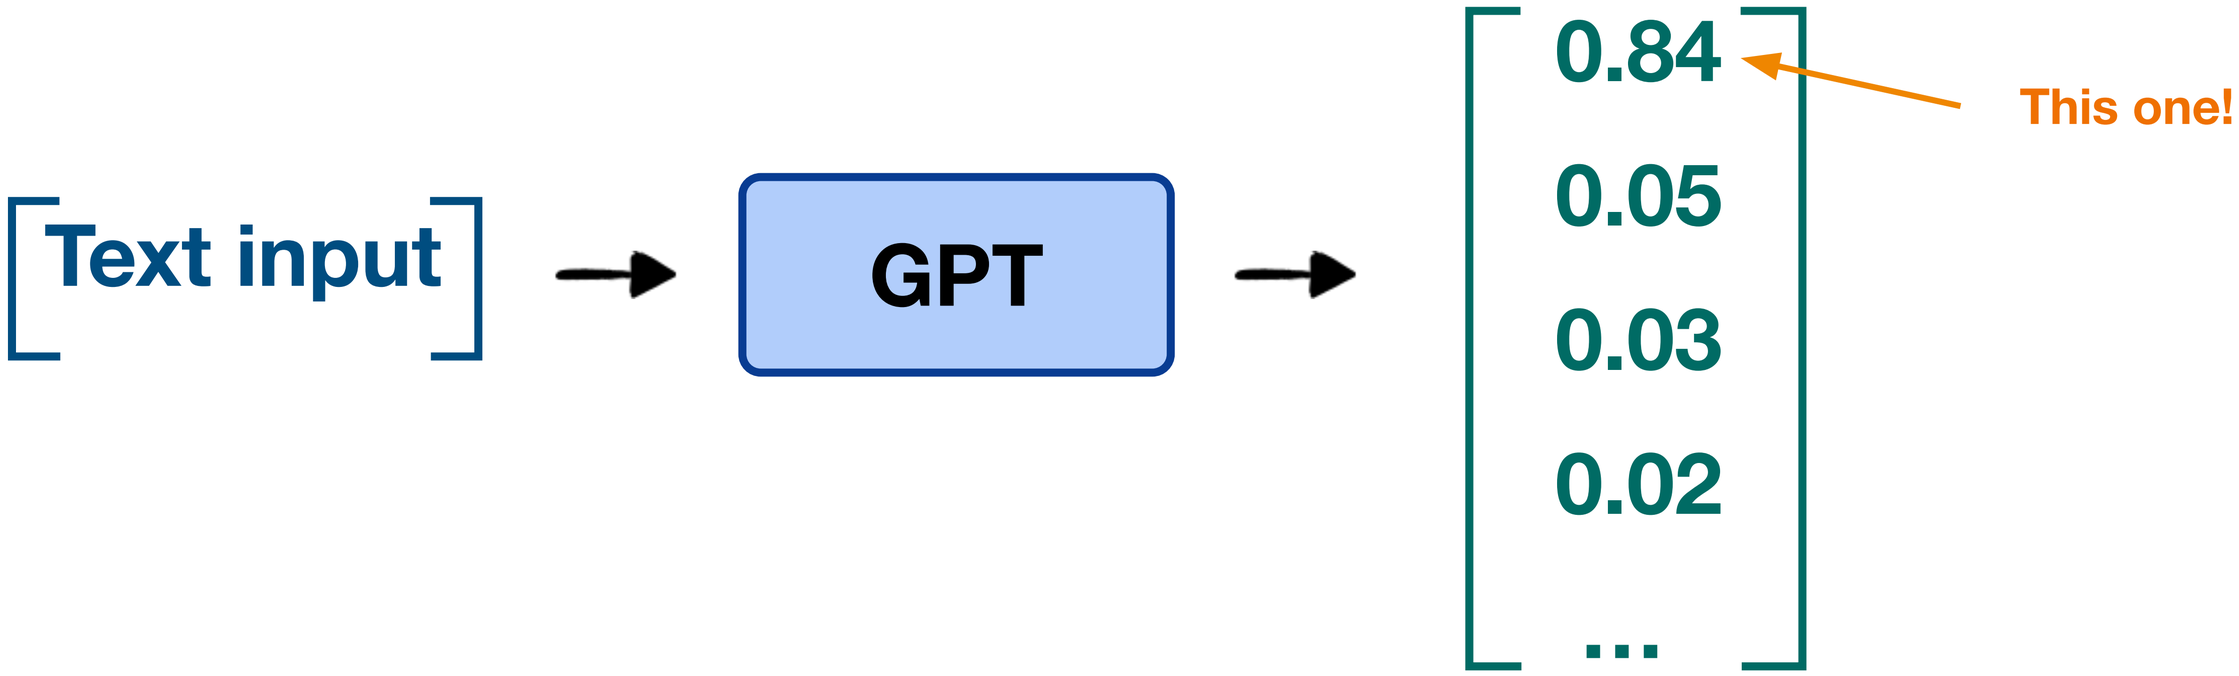

Predicting Next Word

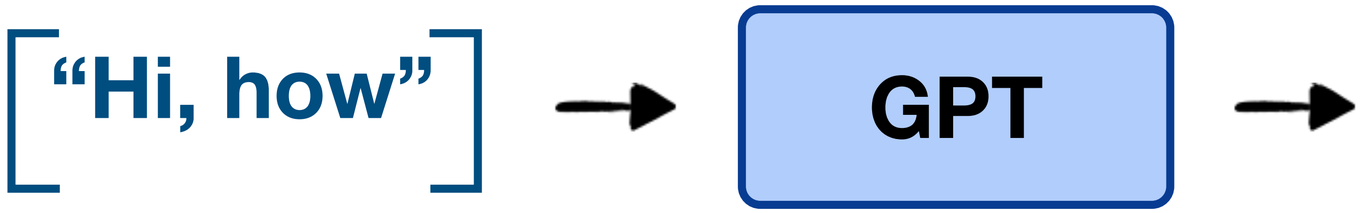

Predicting Next Word

- Now we **repeat** and **generate more text**!

Predicting Next Word

- Now we **repeat** and **generate more text**!

Predicting Next Word

- Now we **repeat** and **generate more text**!

Predicting Next Word

- Now we **repeat** and **generate more text**!

Already a powerful tool. But…

“What is the highest mountain range?”

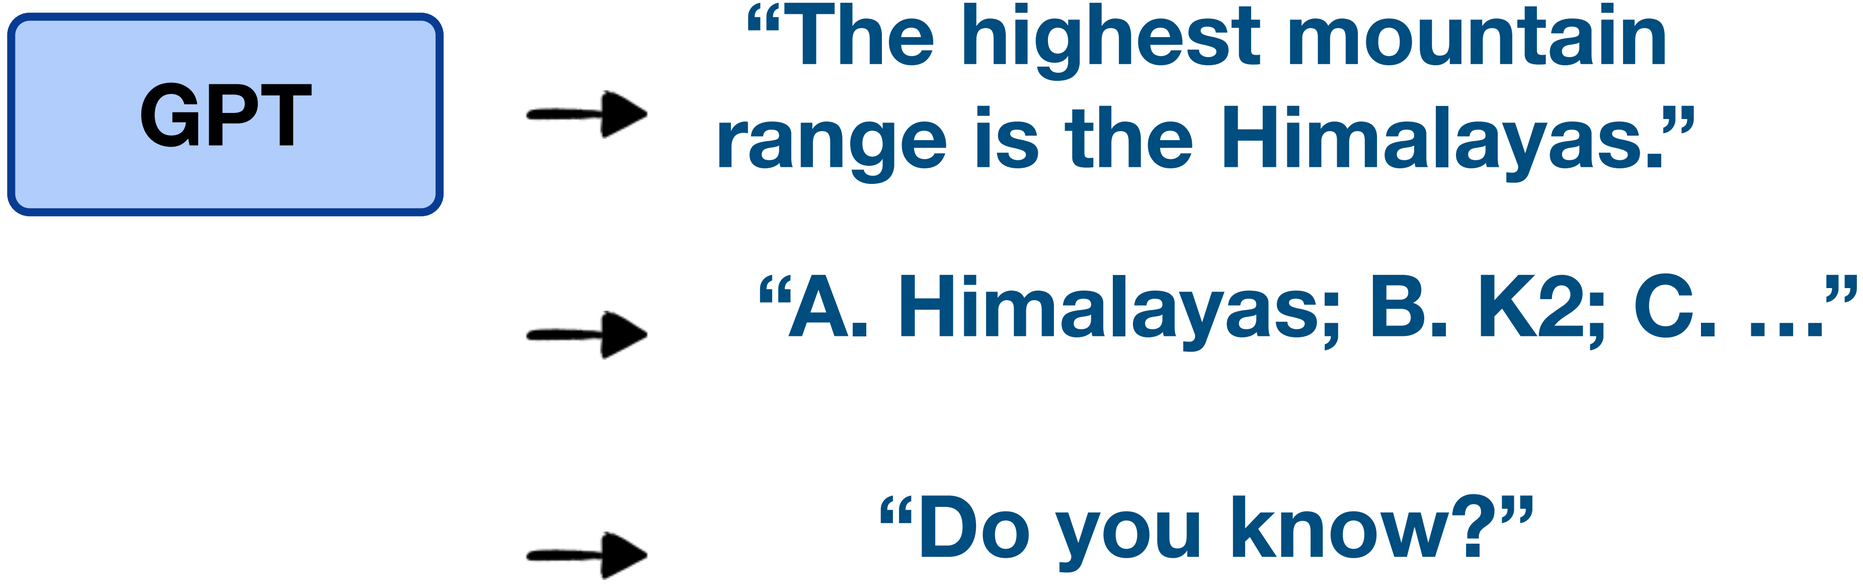

Training ChatGPT from GPT: Fine-tuning

- Reinforcement learning from human feedback (RLHF)

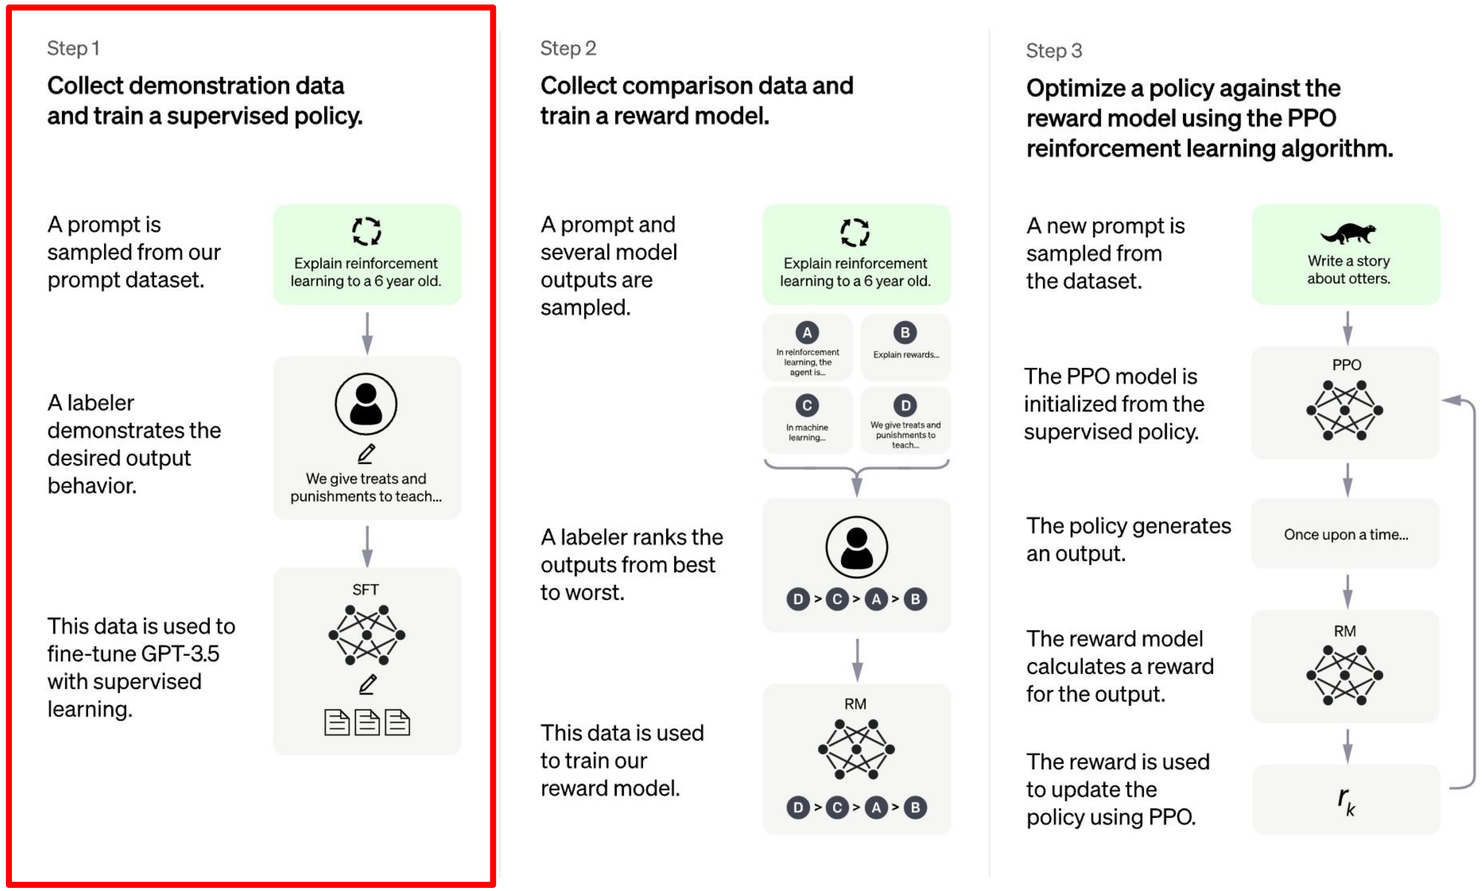

CREDIT: https://openai.com/blog/chatgpt

“What is the highest mountain range?”

human

“The highest mountain range is the Himalayas.”

Fine-tuning

ChatGPT

Training ChatGPT: Specialisation

- Reinforcement learning from human feedback (RLHF)

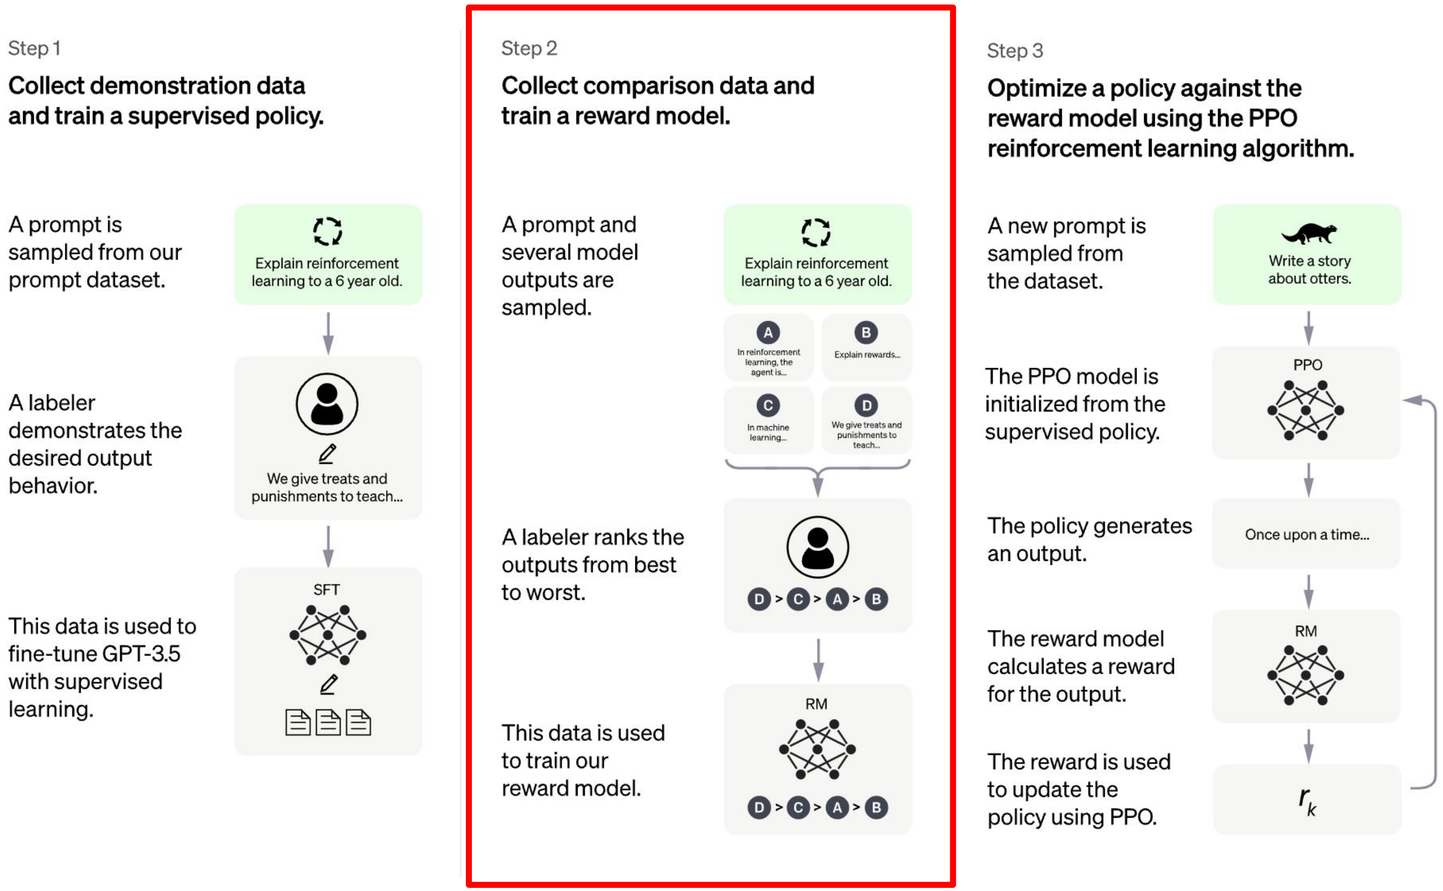

CREDIT: https://openai.com/blog/chatgpt

“What is the highest mountain range?”

GPT

→ “The highest mountain range is the Himalayas.” ✓

Human labeling

→ “A. Himalayas; B. K2; C. .” ✗

“What is the highest mountain range?”

“The highest mountain range is the Himalayas.”

“A. Himalayas; B. K2; C. .”

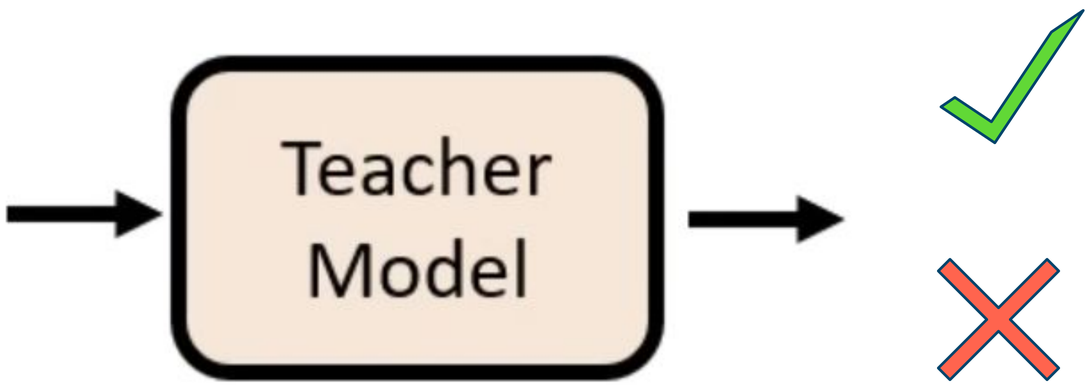

Training ChatGPT: Specialisation

- Reinforcement learning from human feedback (RLHF)

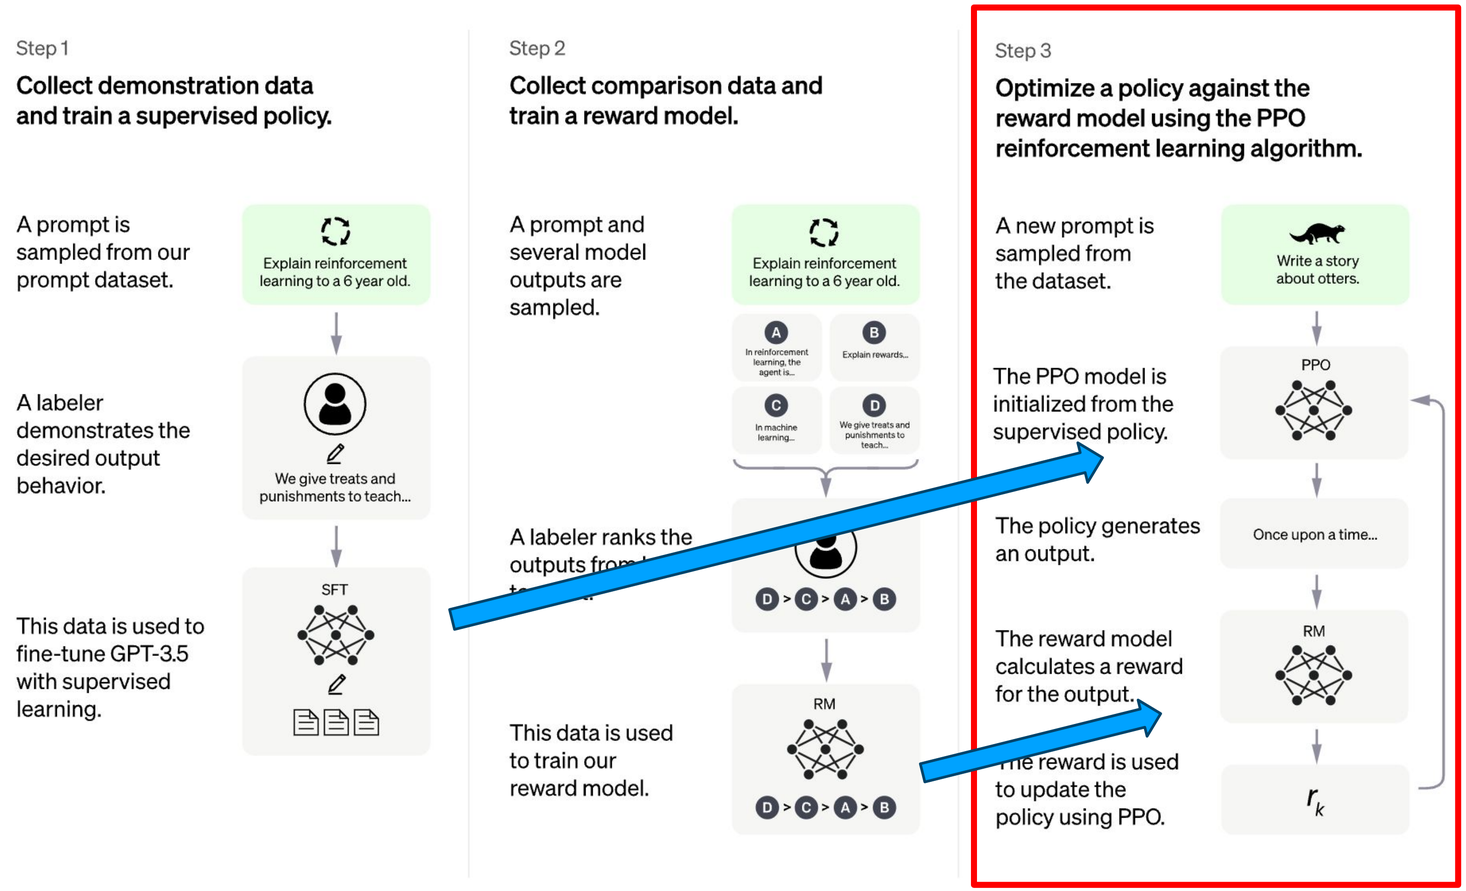

CREDIT: https://openai.com/blog/chatgpt

“What is the highest mountain range?”

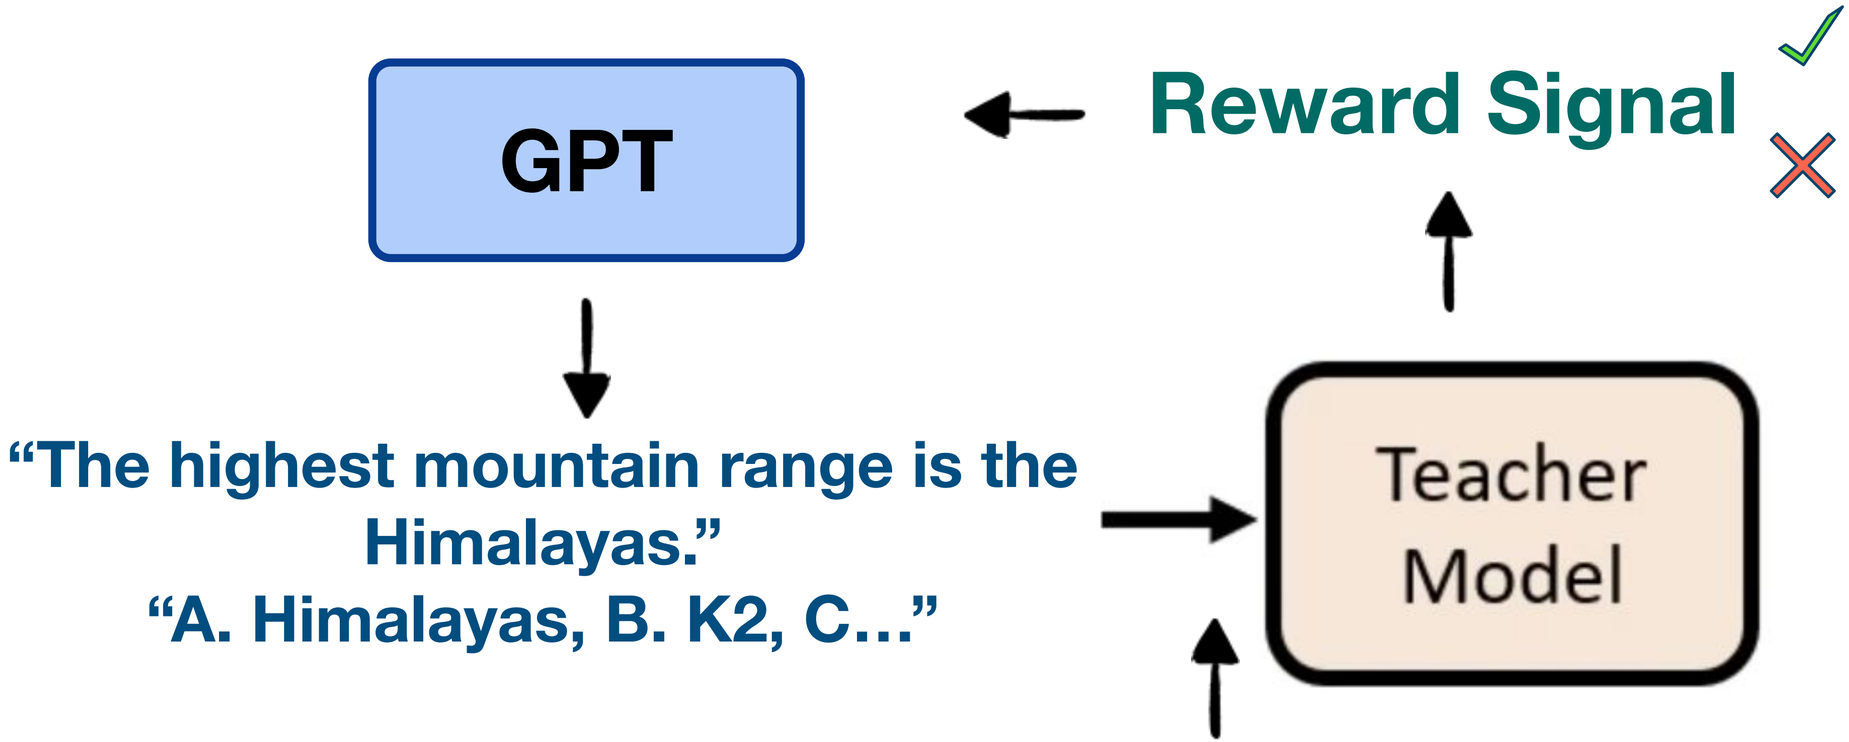

Training ChatGPT

- Pretrained large language models + fine-tuning = **ChatGPT**.

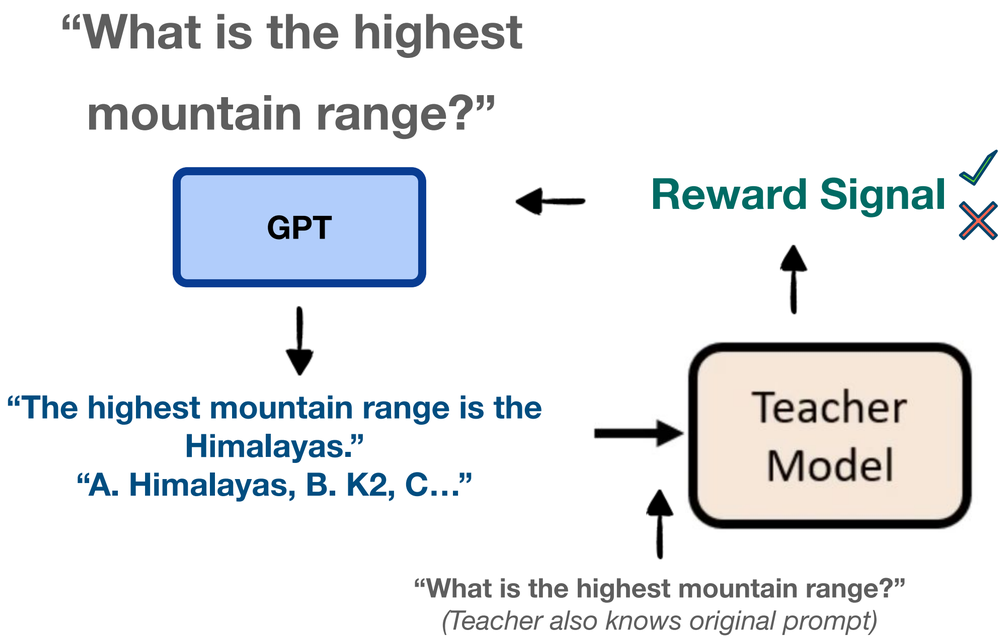

Training ChatGPT

Using ChatGPT

ChatGPT

- Importantly, ChatGPT's **primary goal** is to **generate as human-like text as possible**.
- The model is **probabilistic**.
- It is **by no means** guaranteed to be correct.
- It provides good starting points, but **verify**!

Some Warnings

1. **Don’t take what it says for granted!** Always fact check its information.

Steven A. Schwartz told a judge considering sanctions that the episode had been “deeply embarrassing.” Jefferson Siegel for The New York Times

Some Warnings

1. **Don’t take what it says for granted!** Always fact check its information.

2. **Your conversations are part of its learning process**. Data privacy rules apply; be careful!

https://openai.com/policies/privacy-policy

1. Personal information we collect

We collect personal information relating to you (“Personal Information”) as follows:

Personal Information You Provide: We collect Personal Information if you create an account to use our Services or communicate with us as follows:

- *Account Information:* When you create an account with us, we will collect information associated with your account, including your name, contact information, account credentials, payment card information, and transaction history, (collectively, “Account Information”).
- *User Content:* When you use our Services, we collect Personal Information that is included in the input, file uploads, or feedback that you provide to our Services (“Content”).

Some Warnings

1. **Don’t take what it says for granted!** Always fact check its information.

2. **Your conversations are part of its learning process**. Data privacy rules apply; be careful!

3. **ChatGPT is a neutral tool.** It falls on you to ensure your uses are ethical and sanctioned.

Some Warnings

1. **Don’t take what it says for granted!** Always fact check its information.
2. **Your conversations are part of its learning process**. Data privacy rules apply; be careful!
3. **ChatGPT is a neutral tool.** It falls on you to ensure your uses are ethical and sanctioned.

- Give specific instructions to the bot
  - Make sure you leave little to no room for error!
- Treat the responses like suggestions, not facts
  - It’s good at a lot of things, but not great at one particular thing

Xin Liu, Professor, Computer Science. UC Davis

AIBridge

With acknowledgments to

AI Institute for Food Systems

Cupertino Library