Women of Search

Erika Cardenas

Women of Search

For women (and allies!), started by women

~150 members on slack (since March 2021)

Non-judgmental place to bring cool ideas, questions or *brag!*

Practice public speaking and sessions to help your soft skills

Women of Search

omen of Search session at Haystack US 2022
How to create inclusive environments
What is desired at employers when job seeking

Women of Search session at Haystack EU 2022
• Presented the survey results identifying how women feel in the Search community
• Invited speakers to celebrate their achievements in Search

Women of Search at Haystack-On-Tour Krakow 2023
• Presented 'Vectors in Ecommerce Search'
• How to grow Women-of-Search in Poland

Women of Search

“Happy Hour” continu🥳s ...
every **1st Wednesday of the month at 12 pm EST**.

Lots of workshops and cool sessions planned!

Audrey Sage Lorberfeld, MLIS (She/Her) • 1st
Data / AI / Search lady
4mo

Hello, friends! On January 6, at 6:30 pm CET, we are hosting Sebastian Witalec of SeMI technologies (https://lnkd.in/d3FuWrq7). He is speaking to us about, ironically, public speaking!

All members & allies are welcome to join this session. Make sure to bring your questions about public speaking!

To get our calendar of events, join our Google Group: https://lnkd.in/d4jQAuHE.

See you there! #events #publicspeaking

Sebastian Witalec - Senior Developer Advocate - Progress | LinkedIn

Join us on Slack:

https://relevancy.slack.com/archives/C022UJ7SW2W

Building Recommendation Systems with Vector Search

Erika Cardenas

Erika Cardenas

Developer Advocate at Weaviate

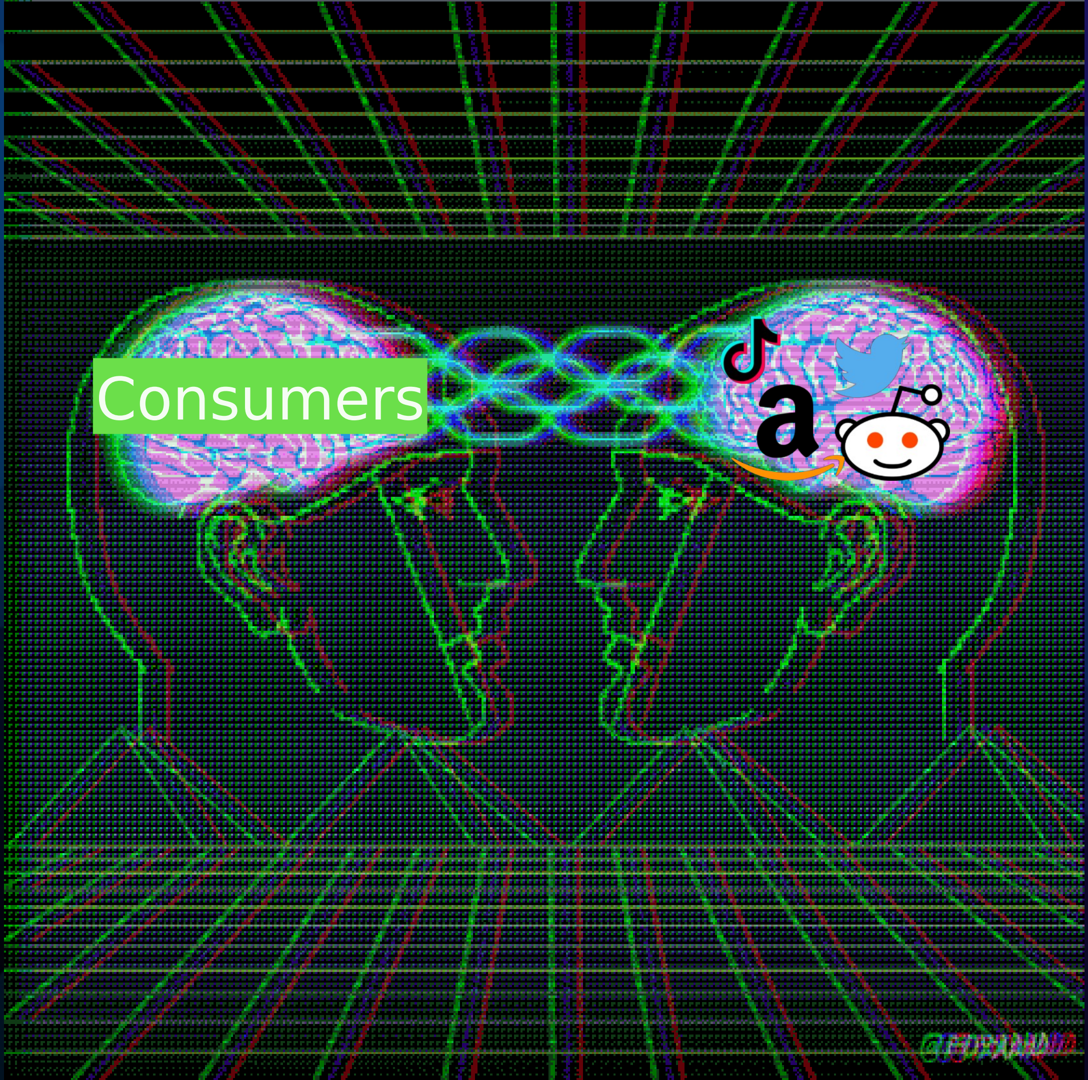
Consumers

Takeaways

1. What is the impact of recommendation?
2. Traditional algorithms used in recommender systems
3. New opportunities with vector embeddings
4. Demo and system design
5. Benefits of Ref2Vec
6. Discussion topics

What is the impact of search and recommendation

“Nearly **$300 billion** is lost **each year** from bad online search experiences.”

Dmitry, Kan

35% of Amazon's revenue comes from their recommendation

tps://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers

Search and Recommendation

"A recommendation system (or recommender system) is a class of machine learning that uses **data to help predict, narrow down**, and find **what people are looking for** among an exponentially growing number of options.”

Nvidia

https://www.nvidia.com/en-us/glossary/data-science/recommendation-system/

Search

Query driven

Recommendation

User preferences

Product features

How do we represent user preferences?

Traditional Recommendation Algorithms

- Content-Based Filtering
- Collaborative Filtering

Content-Based Filtering

Content-Based Filtering (Traditional Approach)

6

Content-Based Filtering (Traditional Approach)

4

6

Content-Based Filtering (Traditional Approach)

9

4

6

Content-Based Filtering (Traditional Approach)

	Comedy	Super Hero	Rom-Com
Elf	1	0	0
Spider Man	0	1	0
The Proposal	1	0	1

$\times$

Rating
6
4
9

Content-Based Filtering (Traditional Approach)

	Comedy	Super Hero	Rom-Com
Elf	6	0	0
Spider Man	0	4	0
The Proposal	9	0	9

Content-Based Filtering (Traditional Approach)

	Comedy	Super Hero	Rom-Com	Total
Rating per Genre	15	4	9	28

Content-Based Filtering (Traditional Approach)

	Comedy	Super Hero	Rom-Com	Total
Rating per Genre	15	4	9	28

	Comedy	Super Hero	Rom-Com
Rating per Genre	0.54	0.14	0.32

Content-Based Filtering (Traditional Approach)

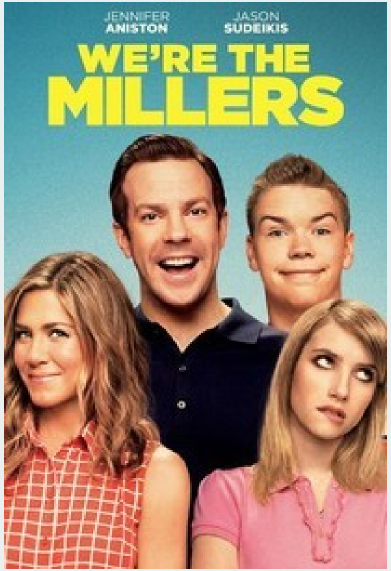

	Comedy	Super Hero	Rom-Com
We’re the Millers	1	0	1

Content-Based Filtering (Traditional Approach)

	Comedy	Super Hero	Rom-Com
We’re the Millers	1	0	1
Venom	0	1	0

Content-Based Filtering (Traditional Approach)

	Comedy	Super Hero	Rom-Com
We're the Millers	1	0	1
Venom	0	1	0
Napoleon Dynamite	1	0	0

Content-Based Filtering (Traditional Approach)

	Comedy	Super Hero	Rom-Com
Rating per Genre	0.54	0.14	0.32

×

	Comedy	Super Hero	Rom-Com
We’re the Millers	1	0	1
Venom	0	1	0
Napoleon Dynamite	1	0	0

	Comedy	Super Hero	Rom-Com
We’re the Millers	0.54	0	0.32
Venom	0	0.14	0
Napoleon Dynamite	0.54	0	0

Content-Based Filtering (Traditional Approach)

	Comedy	Super Hero	Rom-Com	Total	**Rank**
We’re the Millers	0.54	0	0.32	0.86	1
Venom	0	0.14	0	0.14	3
Napoleon Dynamite	0.54	0	0	0.54	2

Collaborative Filtering

Collaborative Filtering (Traditional Approach)

	p_1	p_2	p_3	p_4
u_1	4		5	
u_2	5	1	3	3
u_3	3	4	1	5

Collaborative Filtering (Traditional Approach)

	p_1	p_2	p_3	p_4
u_1	4	?	5	?
u_2	5	1	3	3
u_3	3	4	1	5

Collaborative Filtering (Traditional Approach)

	p_1	p_3
u_1	4	5
u_2	5	3

	p_1	p_3
u_1	4	5
u_3	3	1

Sim(u1, u2) = 0.94

Sim(u1, u3) = 0.84

Collaborative Filtering (Traditional Approach)

Sim(u1, u2) = 0.94

	p_1	p_2	p_3	p_4
u_1	4		5	
u_2	5	1	3	3
u_3	3	4	1	5

Sim(u1, u3) = 0.84

$$\text{Rating}(p_2) = \frac{(0.94*1) + (0.84*4)}{0.94 + 0.84}$$

Collaborative Filtering (Traditional Approach)

Sim(u1, u2) = 0.94

	p_1	p_2	p_3	p_4
u_1	4	2.4	5	
u_2	5	1	3	3
u_3	3	4	1	5

Sim(u1, u3) = 0.84

$$\text{Rating}(p_2) = \frac{(0.94*1) + (0.84*4)}{0.94 + 0.84}$$

Collaborative Filtering (Traditional Approach)

$Sim(u1, u2) = 0.94$

	p_1	p_2	p_3	p_4
u_1	4	2.4	5	
u_2	5	1	3	3
u_3	3	4	1	5

$Sim(u1, u3) = 0.84$

$$Rating(p_4) = \frac{(0.94*3) + (0.84*5)}{0.94 + 0.84}$$

Collaborative Filtering (Traditional Approach)

	p_1	p_2	p_3	p_4
u_1	4	2.4	5	3.9
u_2	5	1	3	3
u_3	3	4	1	5

Sim(u1, u2) = 0.94

Sim(u1, u3) = 0.84

$$\text{Rating}(p_4) = \frac{(0.94*3) + (0.84*5)}{0.94 + 0.84}$$

Limitations of Traditional Recommender Systems

- Scalability
- User-item interaction data is sparse
- Cold-start problem for users and products
- Lack of diversity in results

Vector Representations

Embeddings

Maxi summer dress with no sleeves → Text Encoder → [-0.10, 0.02, 0.14, ..., 0.48]

 → Image Encoder → [-0.12, 0.03, 0.14, ..., 0.52]

Multimodal Embeddings

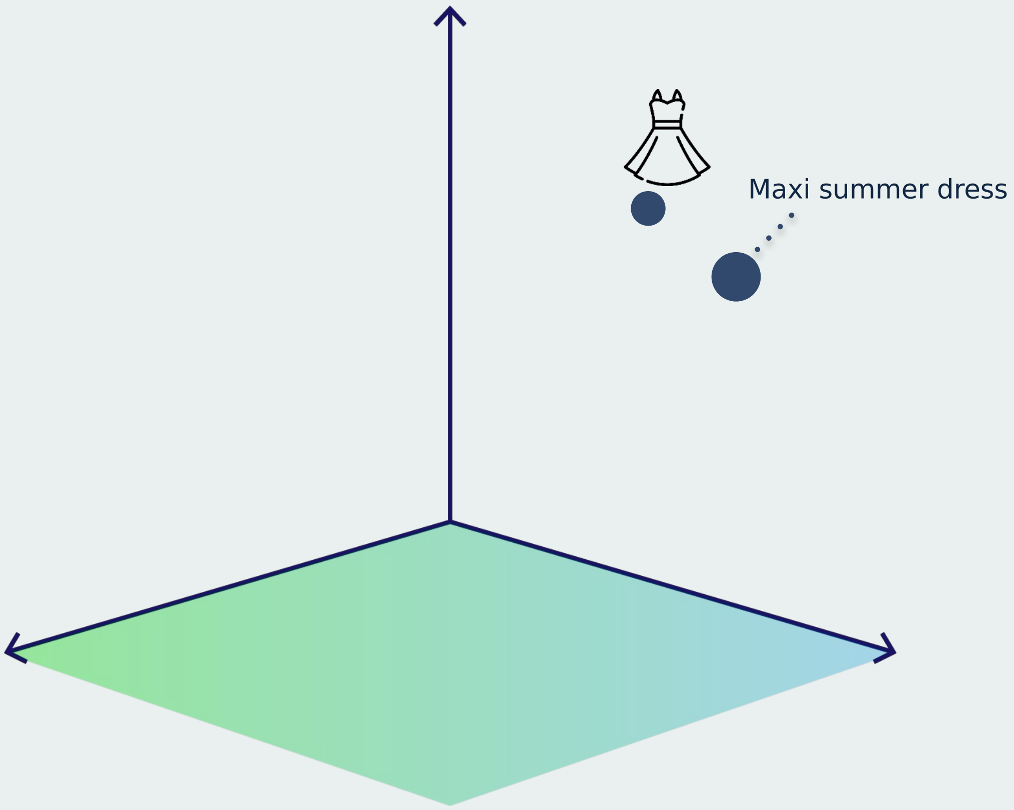

Multimodal Embeddings

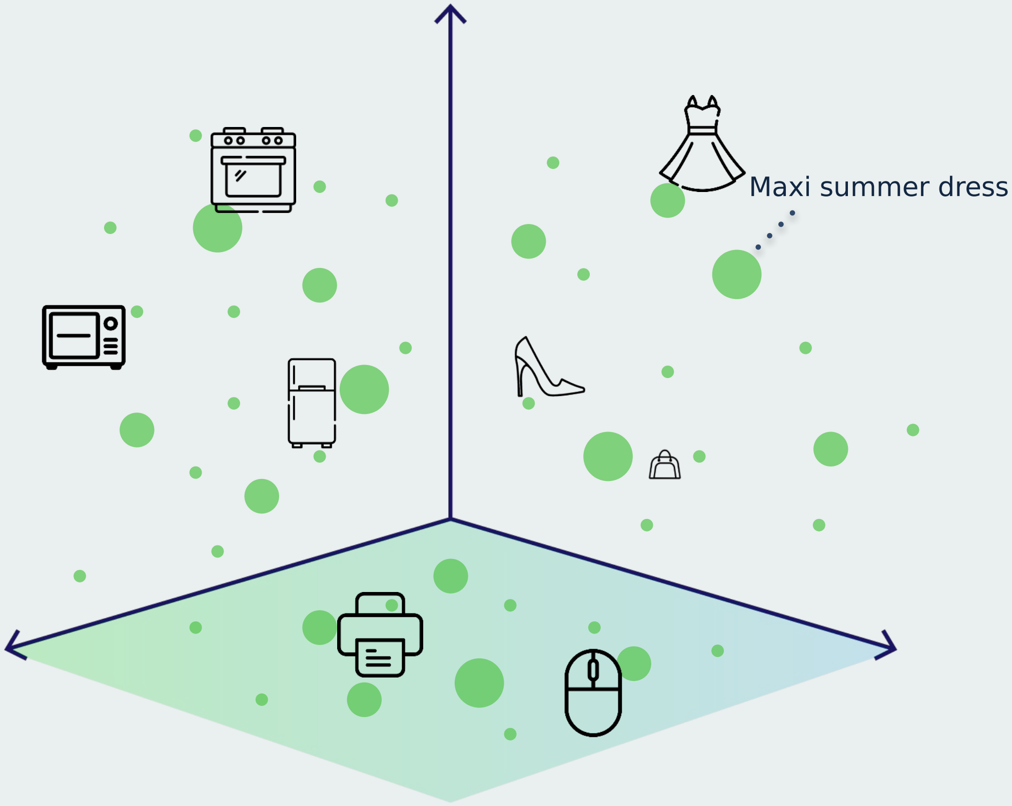

Ref2Vec-Centroid

[-0.12, 0.03, 0.14, ..., 0.52]

[0.04, 0.23, -0.42, ..., 0.01]

[0.29, -0.02, 0.38, ..., 0.54]

[-0.12, 0.03, 0.14, ..., 0.52]

[0.04, 0.23, -0.42, ..., 0.01]

[0.29, -0.02, 0.38, ..., 0.54]

[0.07, 0.08, 0.03, ..., 0.36]

[-0.12, 0.03, 0.14, ..., 0.52]

[0.04, 0.23, -0.42, ..., 0.01]

[0.29, -0.02, 0.38, ..., 0.54]

[0.07, 0.08, 0.03, ..., 0.36]

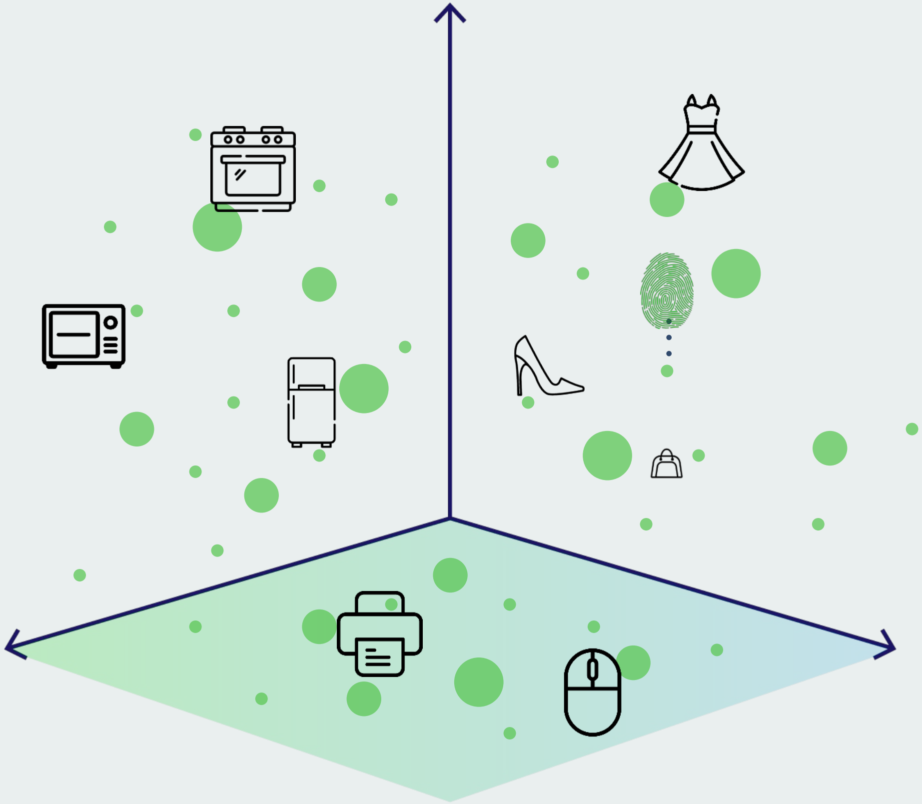

Reference-to-Vector

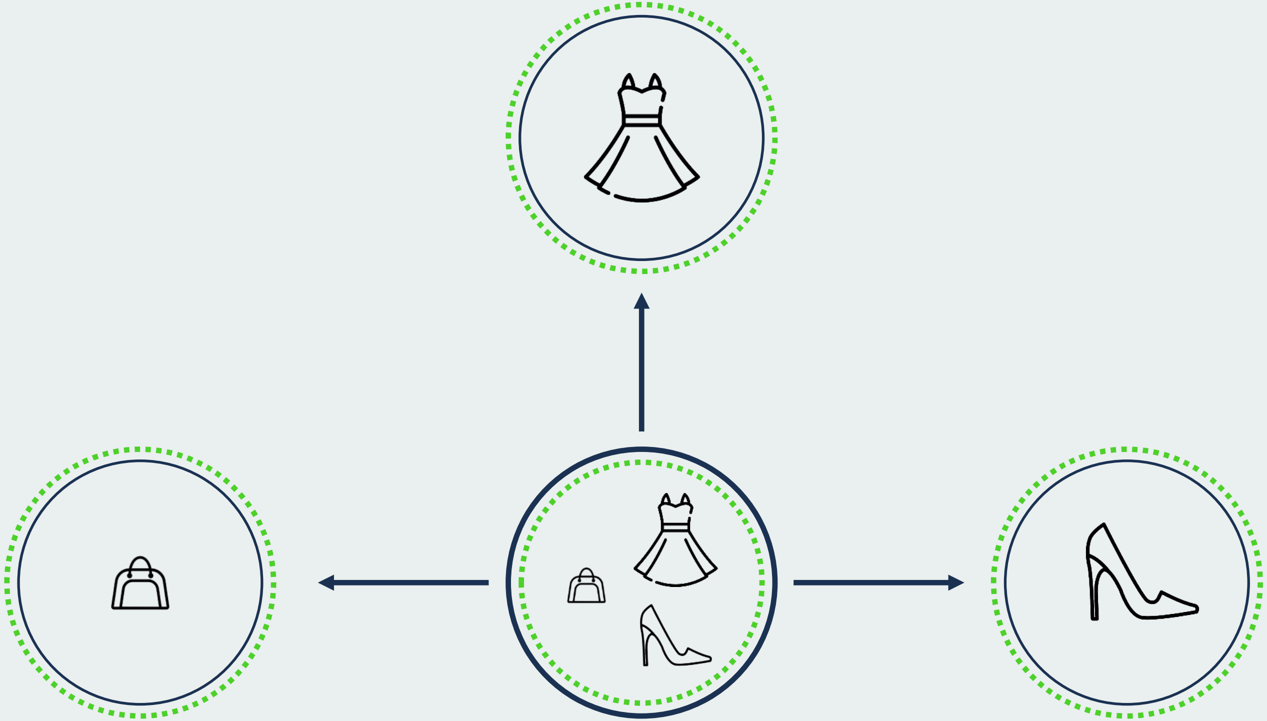

Demo Time

Client-Server Architecture

Client-Server Architecture

Client-Server Architecture

Benefits of Ref2Vec

- Minimal overhead
- Rapid personalization
- Real-time updates

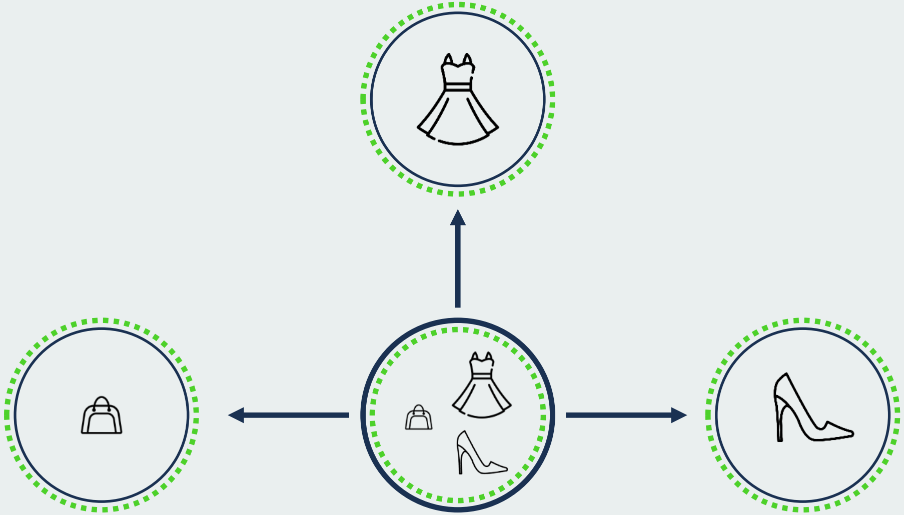

Discussion Topics

Discussion Topics

- Ranking models
- Diversity in results
- Multiple user embeddings

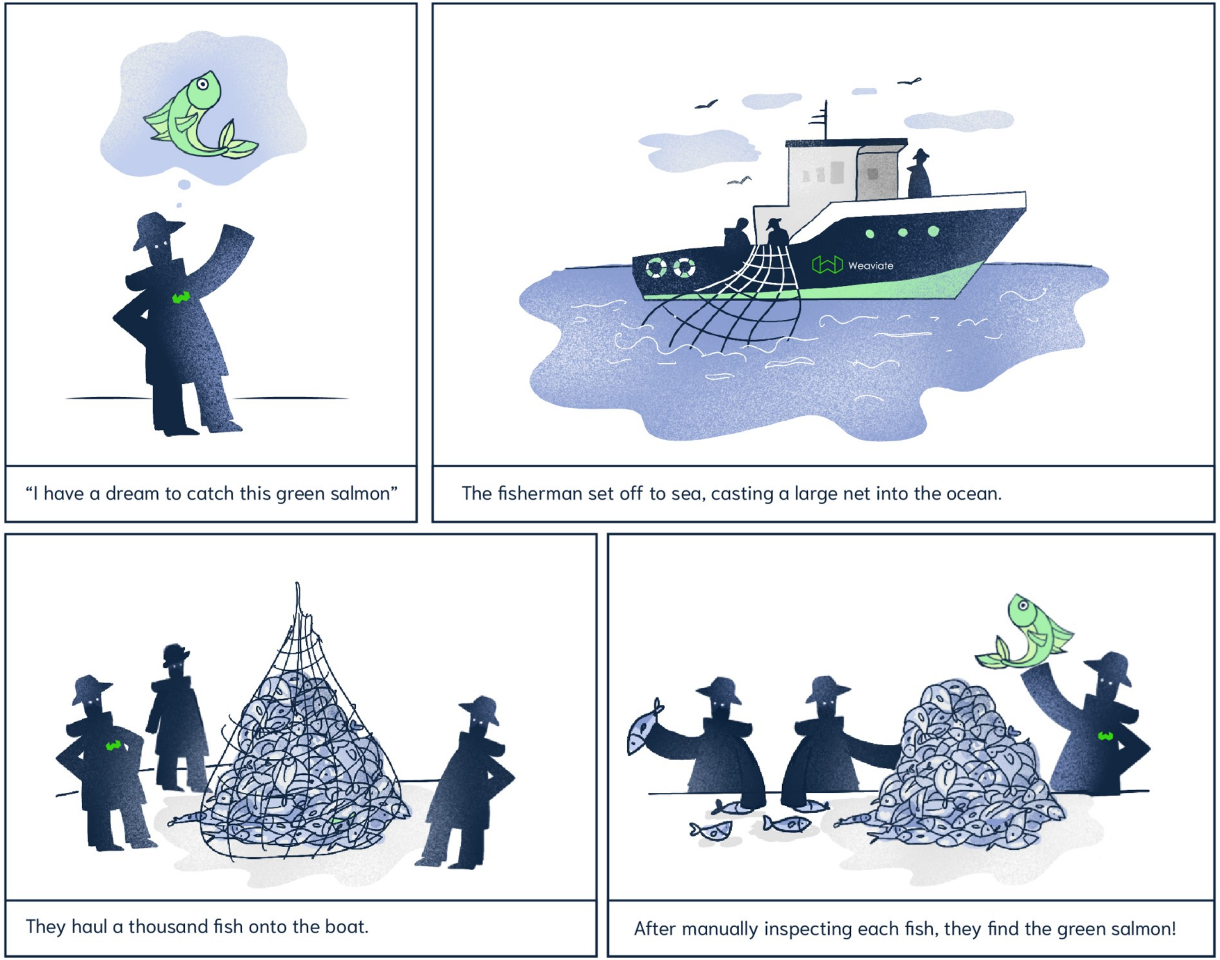
“I have a dream to catch this green salmon”
The fisherman set off to sea, casting a large net into the ocean.
Weaviate
They haul a thousand fish onto the boat.
After manually inspecting each fish, they find the green salmon!

Diversity in Results

Women Shoe icon by Icons8

Diversity in Results

Multiple User Embeddings

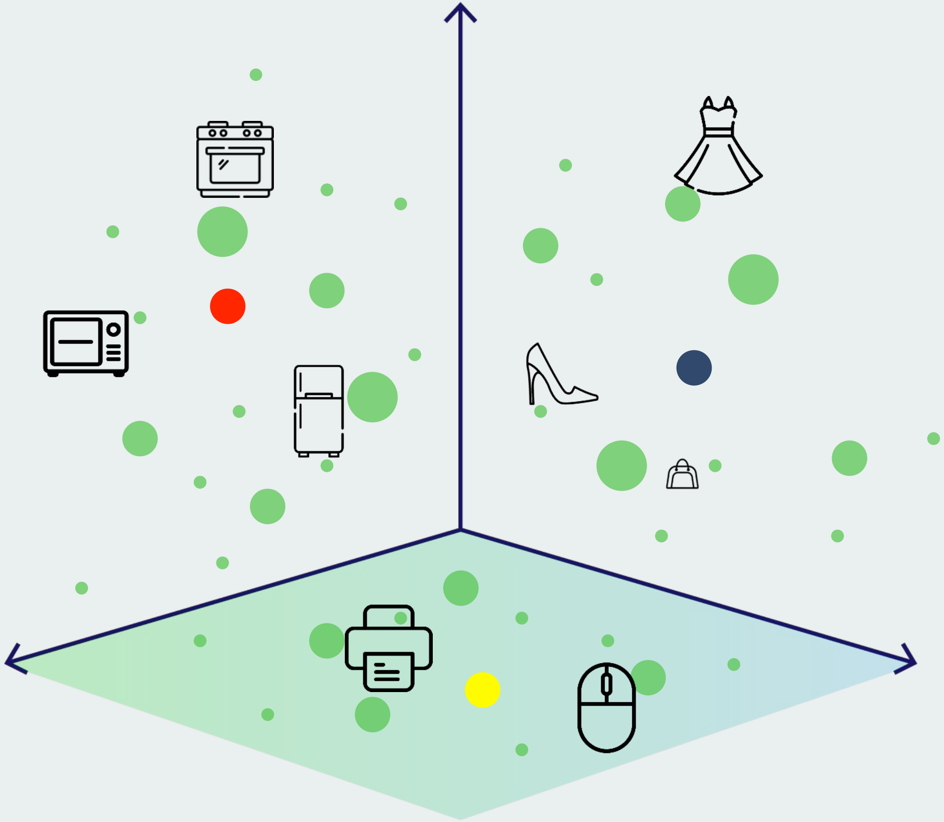

Key Takeaways

- **No longer need** a **treasure trove** of data to build a recommender system
- Vector embeddings allow you to **scale** to **millions** or **billions** of products and users
- Ref2Vec offers **rapid personalization** for your users and has a **low overhead**

Resources

1. https://www.mckinsey.com/industries/retail/our-insights/how-retailers-can-keep-up-with-consumers

2. https://www.nvidia.com/en-us/glossary/data-science/recommendation-system/

Connect with us!

weaviate.io

weaviate/weaviate

weaviate_io

Erika Cardenas

Twitter: @ecardenas300

LinkedIn: https://www.linkedin.com/in/erikacardenas300/

Thank you