

Building a visual similarity search from scratch

SLO AI Lightning Talks #2
December 11, 2025

Robert Sedovšek
<https://galjot.si>

To break the ice...

Career so far

^-_(ツ)_/^- (now)

Turtl (before)

Celtra (way before)

PhD

Faculty of Natural Sciences and
Engineering;
University of Ljubljana

Normal hobbies

Running, XC  and  

Photography

Unusual trivia

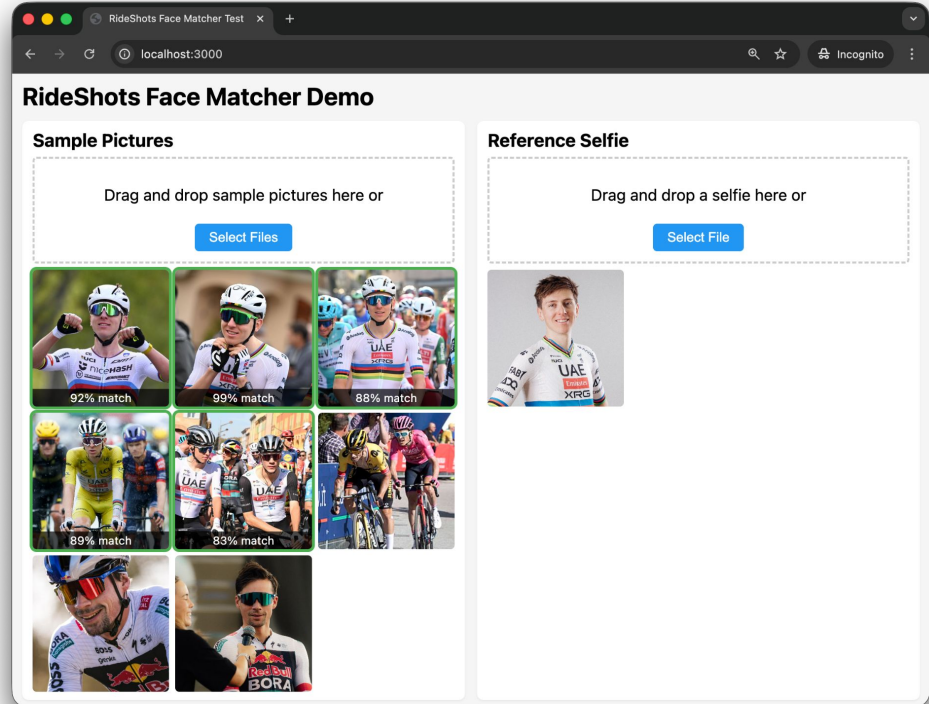
I ran a **100k** (5,500m+ elevation)
race **twice**

Background & motivation

RideShots.cc (*passion project*)

Image search via **facial recognition**:

- ~~face~~api.js
- ~~face~~_recognition
- Insightface
- AWS Rekognition



Let's build it!

— from scratch, with open-source tools —

On a **concrete case**: finding a specific item
(e.g., a dress) within a catalog of product images.

✓ What to expect?

1. **Data collection**
2. **Object detection**
locating the dress in the image
3. **Segmentation & background removal**
isolating the dress itself
4. **Embedding**
turning the image into numerical vectors
5. **Search**
finding similar embeddings

✗ What this is not?

- Scalability/perf. concerns
- Production ready setup
- Alternatives:
numerous other models I tried and tested, less successfully, and not used in the demo

Which two photos are more similar? (1/2)



Which two photos are more similar? (2/2)



1. Data collection

Technologies

- **Puppeteer** — headless browser automation
- **Axios** — HTTP requests for image downloads
- **Node.js** — file system operations

Process

- Reads product URLs from a text file
- Navigates pages with Puppeteer
- Simulates scrolling to trigger lazy-loaded images
- Extracts image URLs via DOM selectors
- Downloads images to local storage



SEARCH

ASYMMETRIC BEADED SATIN EFFECT MIDDRESS
\$ 79.90

ANTHRACITE GREY | 5029/208

Asymmetric neckline and slit. Beads at back. Back button.

PRODUCT MEASUREMENTS

COMPOSITION & CARE

CHECK IN-STORE AVAILABILITY

SHIPPING, EXCHANGES AND RETURNS

Due to the holiday season, orders placed until 30th of December.



CHAT

2. Object detection

Term


Object detection

Task

Find and label objects in an image

Why it matters?

Identify what's in an image



Model: Florence-2 Object Detection (Microsoft)

A lightweight vision-language model open-sourced by Microsoft under the MIT license.

```
# Object Detection: Finding the dress in the image

1 from transformers import AutoProcessor, AutoModelForCausalLM
2
3 # Load Florence-2 (processor for image/text encoding, model for inference)
4 processor = AutoProcessor.from_pretrained("microsoft/Florence-2-base")
5 model = AutoModelForCausalLM.from_pretrained("microsoft/Florence-2-base")
6
7 def detect_clothing(image):
8     od_task = "<OD>" # Object Detection
9
10    inputs = processor(text=od_task, images=image, return_tensors="pt")
11    detections = model.generate(**inputs)
12    parsed = processor.post_process_generation(
13        detections, task=od_task, image_size=(image.width, image.height)
14    )
15
16    bbox = parsed['<OD>']['bbox']
17    label = parsed['<OD>']['label']
18
19    return None if label not in ['dress', 'shirt', 'pants', 'jacket', ...]:
20
21    return {'name': label, 'bbox': bbox}
22
```

3. Segmentation & background removal

Term

Segmentation

Task

Pixel-level object mask (isolating the dress itself)

Why it matters?

Enables background removal and object isolation



Model: SAM 2 – Segment Anything Model v2 (Meta)

Designed specifically to generate pixel-accurate masks for any object in an image.

```
Segmentation & Background Removal: Isolating the dress
1 from transformers import Sam2Model, Sam2Processor
2
3 sam2_processor = Sam2Processor.from_pretrained("facebook/sam2-hiera-large")
4 sam2_model = Sam2Model.from_pretrained("facebook/sam2-hiera-large")
5
6 def segment_clothing(image, bbox):
7     inputs = sam2_processor(
8         images=image,
9         input_boxes=[[bbox[0], bbox[1], bbox[2], bbox[3]]],
10        return_tensors="pt"
11    )
12    outputs = sam2_model(**inputs, multimask_output=False)
13    processed_masks = sam2_processor.post_process_masks(
14        outputs.pred_masks.cpu(),
15        inputs["original_sizes"]
16    )
17    return processed_masks[0][0][0].numpy() > 0 # Pick best mask
18
19 def remove_background(img, mask):
20     original_rgb = np.array(img)
21     bg_removed = np.zeros((img.height, img.width, 4)) # Transparent image
22
23     for y in range(img.height):
24         for x in range(img.width):
25             if mask[y, x]: # If mask is True, keep pixel
26                 bg_removed[y, x] = [*original_rgb[y, x], 255]
27
28     return PILImage.fromarray(bg_removed, 'RGBA')
29
```

4. Embedding

Term

Embedding

Task

Convert a photo into a numerical vector representing its features.

Why it matters?

Enables comparison between vectors
(*similar images*)



Model: FashionCLIP

(patrickjohncyh/fashion-clip)

A vision-language model based on CLIP (OpenAI), fine-tuned on fashion data.

```
Embedding: Turning the image into numerical vectors

1 from transformers import CLIPModel, CLIPProcessor
2
3 model = CLIPModel.from_pretrained("patrickjohncyh/fashion-clip")
4 processor = CLIPProcessor.from_pretrained("patrickjohncyh/fashion-clip")
5
6 def generate_embedding(garment_image):
7     inputs = processor(images=garment_image, return_tensors="pt")
8     embedding = model.get_image_features(**inputs)
9
10    return embedding[0].numpy() # Returns: 512-dimensional vector
```

5. Search

Term

Search

Task

Find similar images by comparing embedding vectors

Why it matters?

Enables fast retrieval of visually similar products




Search

🤔 [sedovsek/garment-similarity-search](https://github.com/sedovsek/garment-similarity-search)

For each query image, compute **cosine similarity** against all product embeddings.

```
Search: Finding similar embeddings
1 import numpy as np
2
3 class SimpleVectorIndex:
4     def __init__(self):
5         self.embeddings, self.product_ids = []
6
7     def insert(self, product_id, embedding):
8         self.embeddings.append(embedding)
9         self.product_ids.append(product_id)
10
11    def search(self, query_embedding, k=5):
12        # CLIP embeddings are already normalized, so we can compute
13        # cosine similarity directly via dot product
14        similarities = np.dot(self.embeddings, query_embedding)
15
16        # Get top-k most similar products
17        top_k_indices = np.argsort(similarities)[::-1][:k]
18
19        results = []
20        for idx in top_k_indices:
21            results.append({
22                'id': self.product_ids[idx],
23                'similarity': float(similarities[idx])
24            })
25
26        return results
27
```



EyeBye

Upload an image to find similar products

Upload

Drag & drop an image here
or click to browse

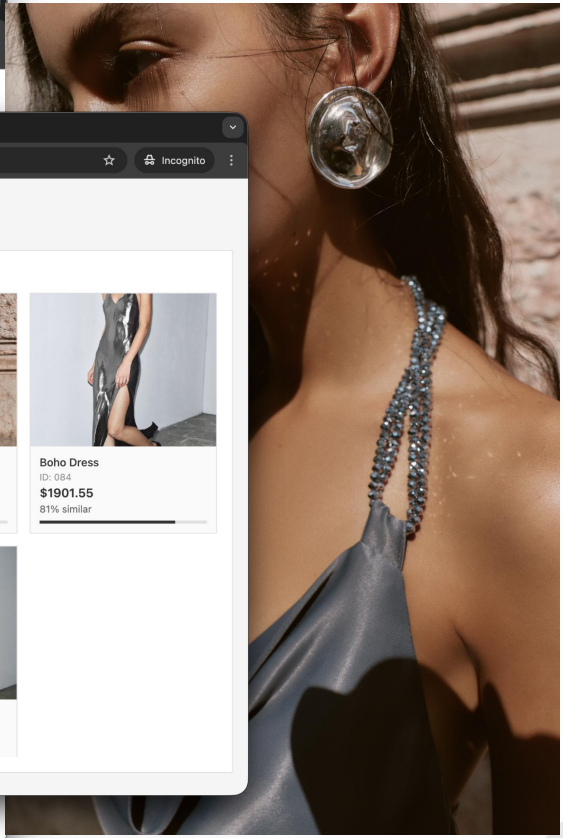
Preview

Number of results:
Search

File: 3. Query image.jpg

Results

<p>Sheath Dress ID: 093 \$1735.49 84% similar</p>	<p>Boho Dress ID: 084 \$1901.55 81% similar</p>
<p>Cocktail Dress ID: 056 \$920.19</p>	



Pipeline

```
# Visual Similarity Search Pipeline

1 vector_index = SimpleVectorIndex()
2
3 # OFFLINE: Build product index
4 for product in load_products("data/products.json"):
5     for image in product["images"]:
6         embedding = process_image(Image.open(image))
7         vector_index.insert(product["id"], embedding)
8
9 # CLIENT: Search for similar products
10 query_embedding = process_image(Image.open("upload.jpg"))
11 similar_products = vector_index.search(query_embedding, k=5)
12
13 def process_image(image):
14     """Process: detect → segment → remove background → embed"""
15     detection = detect_clothing(image)
16     mask = segment_clothing(image, detection["bbox"])
17     garment = remove_background(image, mask)
18     return generate_embedding(garment)
19
```

Q&A

Thank you for your attention!

You could have been scrolling through your phone, but you chose not to – I appreciate it!

@sedovsek on [X](#) · [Instagram](#) · [Threads](#) · [Strava](#)
<https://galjot.si>

