ForgetMeNot: Memory-Aware Forensic Facial Sketch Matching

Authors: Ouyang, Hospedales, Song, Li

Slides by Josh Kelle

Overview

- VIPSL dataset
- experiment goals
- experiment results
- conclusion

VIPSL Dataset

- Photographs of 200 faces with neutral expression
- Each photo was sketched by 5 different artists

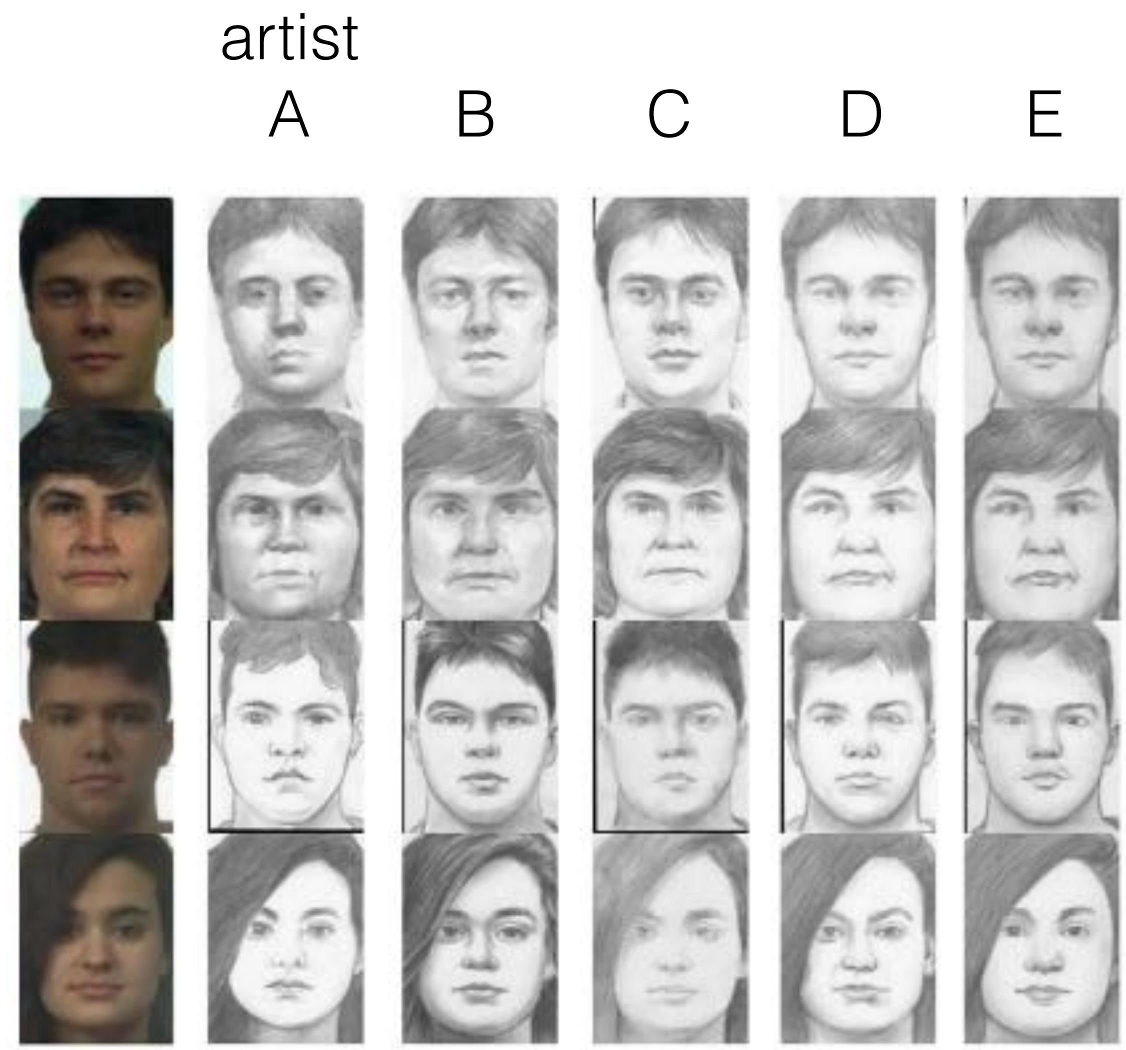

Artist Style

Artist A



Artist B



Goal: re-sketch in a different style

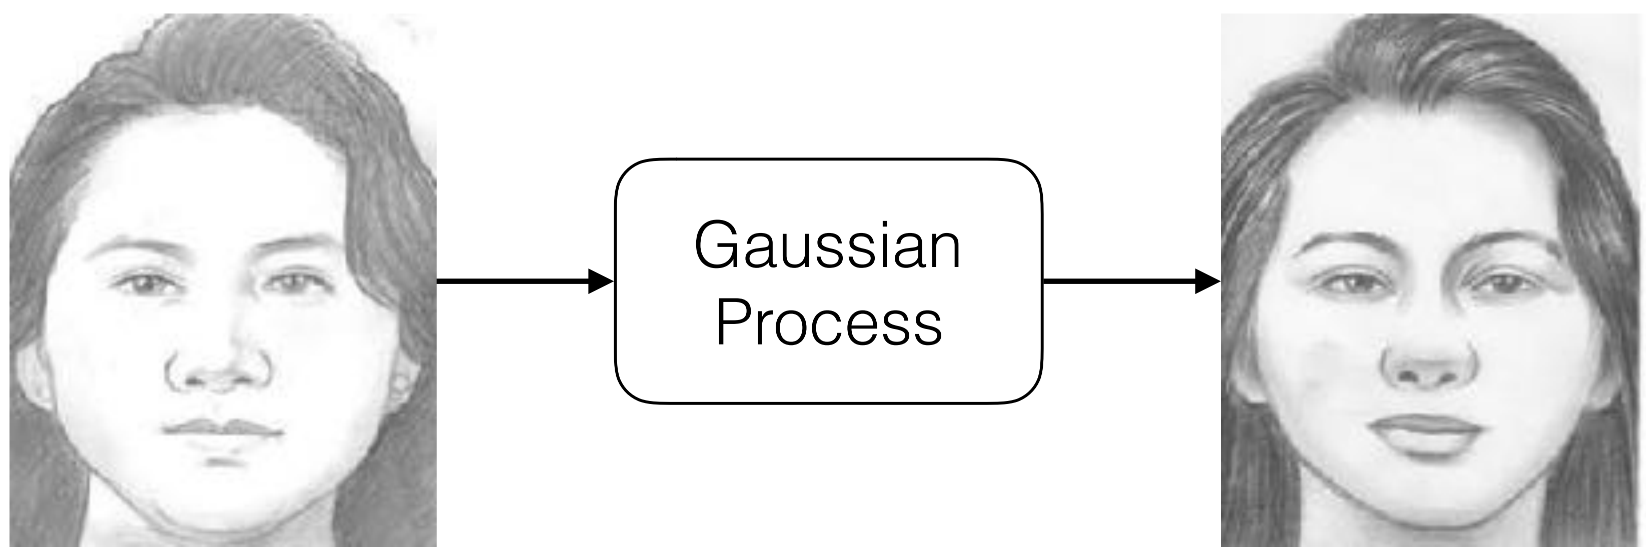

input
sketch from artist A

output
sketch in the style of artist B

HOG representation

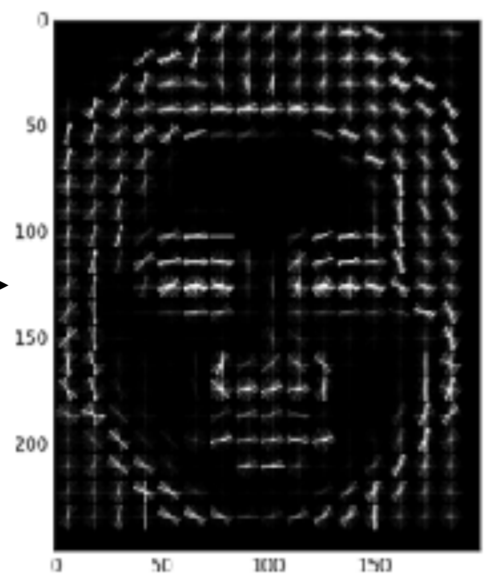

Gaussian Process

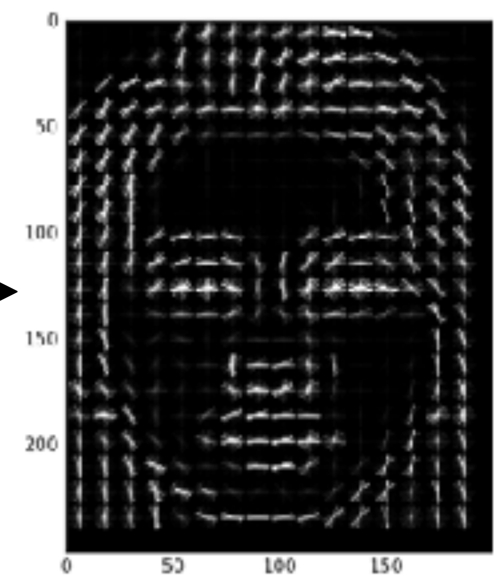

input sketch

input HOG features

output HOG features

invert HOG features

Training the GP

- Treat each HOG image as a vector in $\mathbb{R}^{2560}$.
- Use PCA to reduce this to $\mathbb{R}^{150}$, although this didn't produce a noticeable improvement.
- GP: $\mathbb{R}^{150} \rightarrow \mathbb{R}^{150}$
- Then convert GP output back to $\mathbb{R}^{2560}$ hog space.

X = { 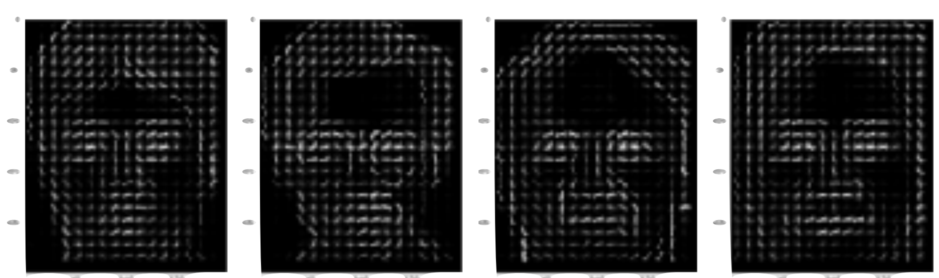 }

Y = { 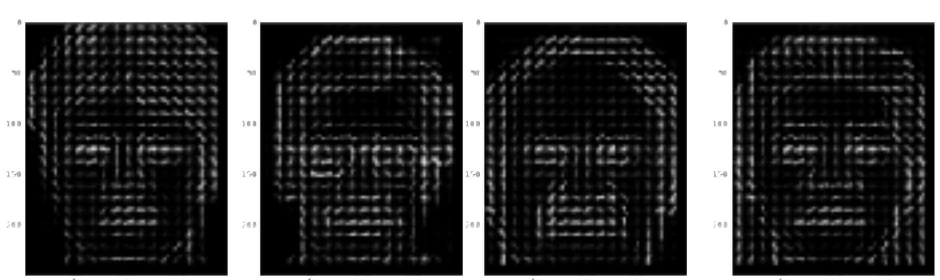 }

Results for A→B model

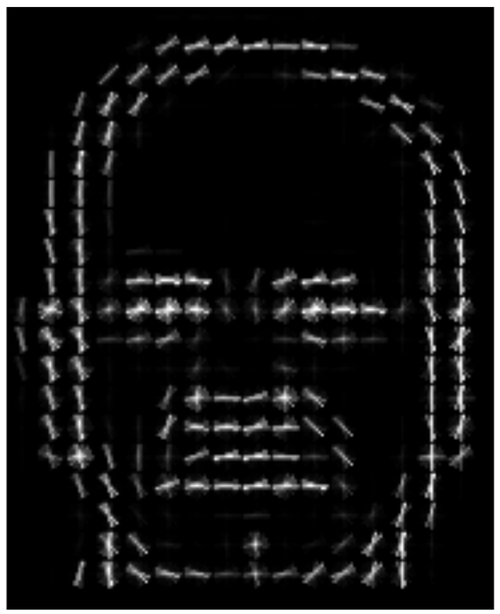

input

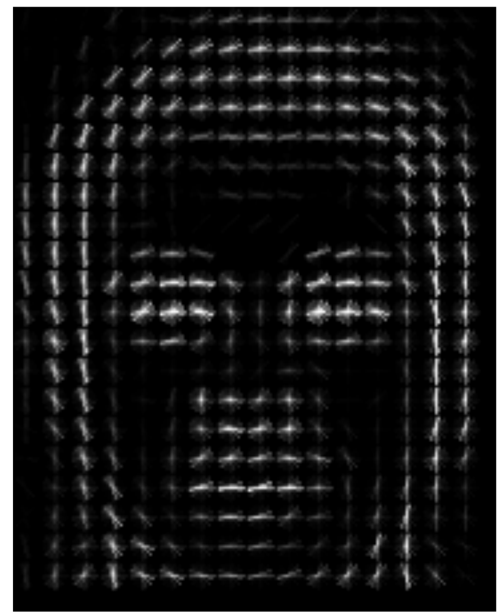

GP prediction

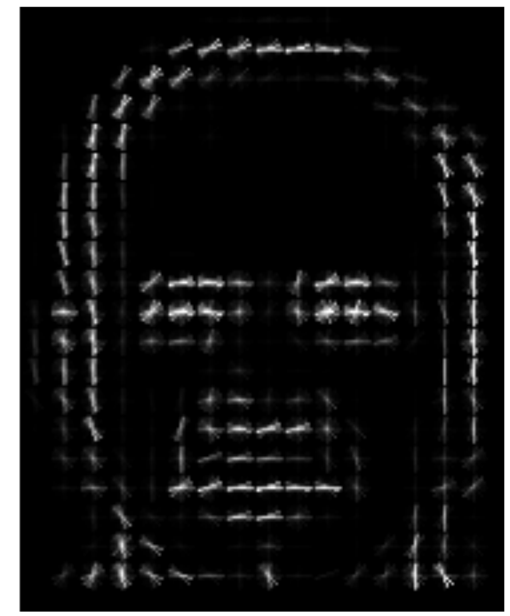

ground truth

- The prediction's gradients look less sharp, which is good.
- I was surprised to see more gradients around the outside of the head.

Results for A→B model

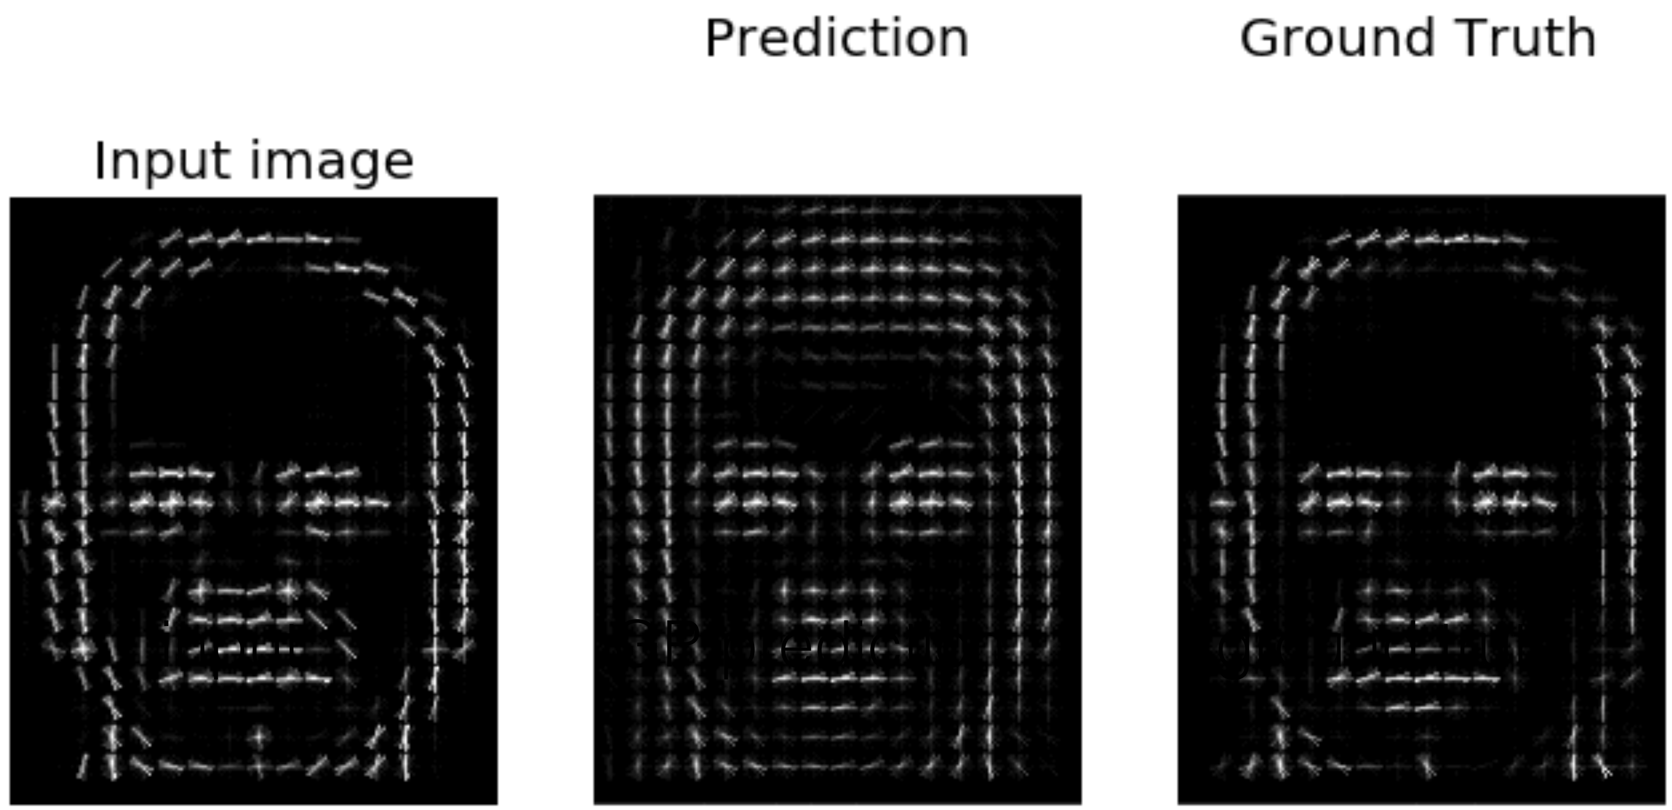

- It looks like the GP is smoothing too much.
- Hypothesis: the GP is putting too much emphasis on the mean face.

Reverse direction: B to A

A

A to B

A has more gradient activity than B

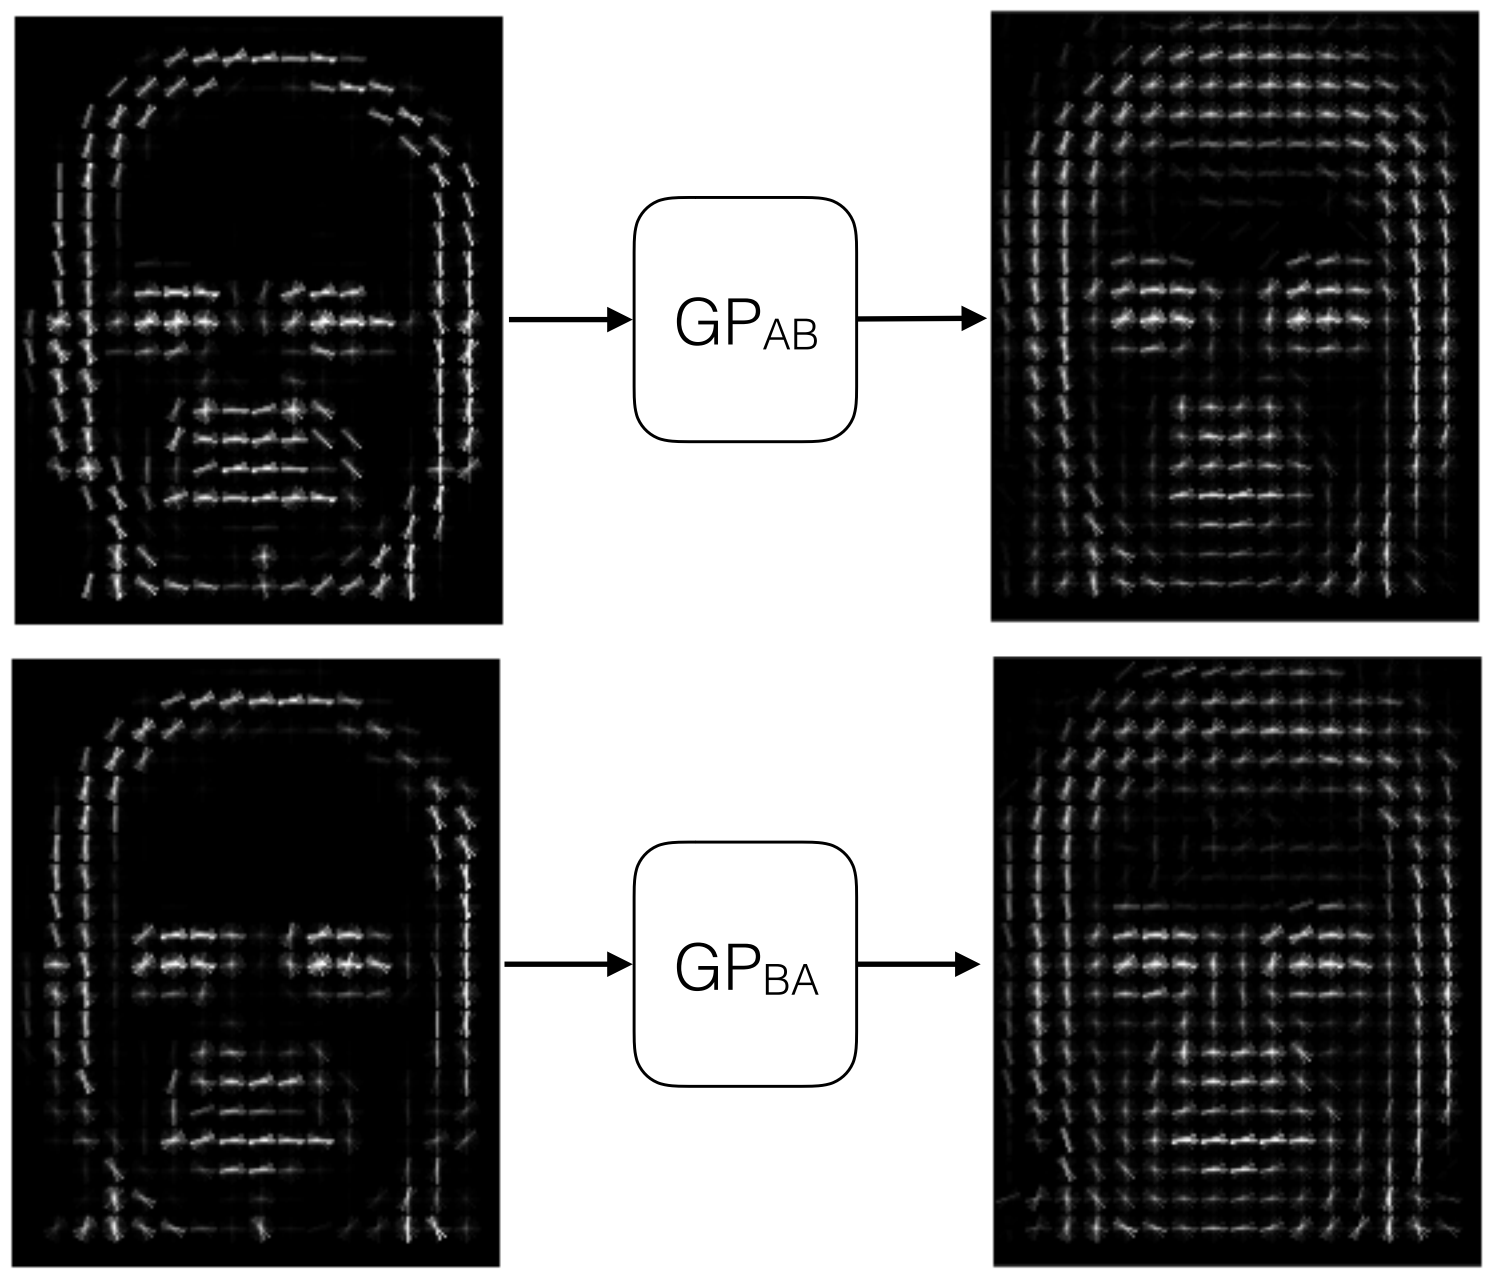

A has more gradient activity than B

B

B to A

Quantifying Style Similarity

- Measure similarity of sketch style by L2 distance in HOG space.

$$sim(A, B) = \frac{1}{n} \sum_{i=0}^{n} \left\| x_i^{(A)} - x_i^{(B)} \right\|_2$$

where $x_i^{(A)}$ is the HOG representation of the i-th sketch from artist A

Quantifying Style Similarity

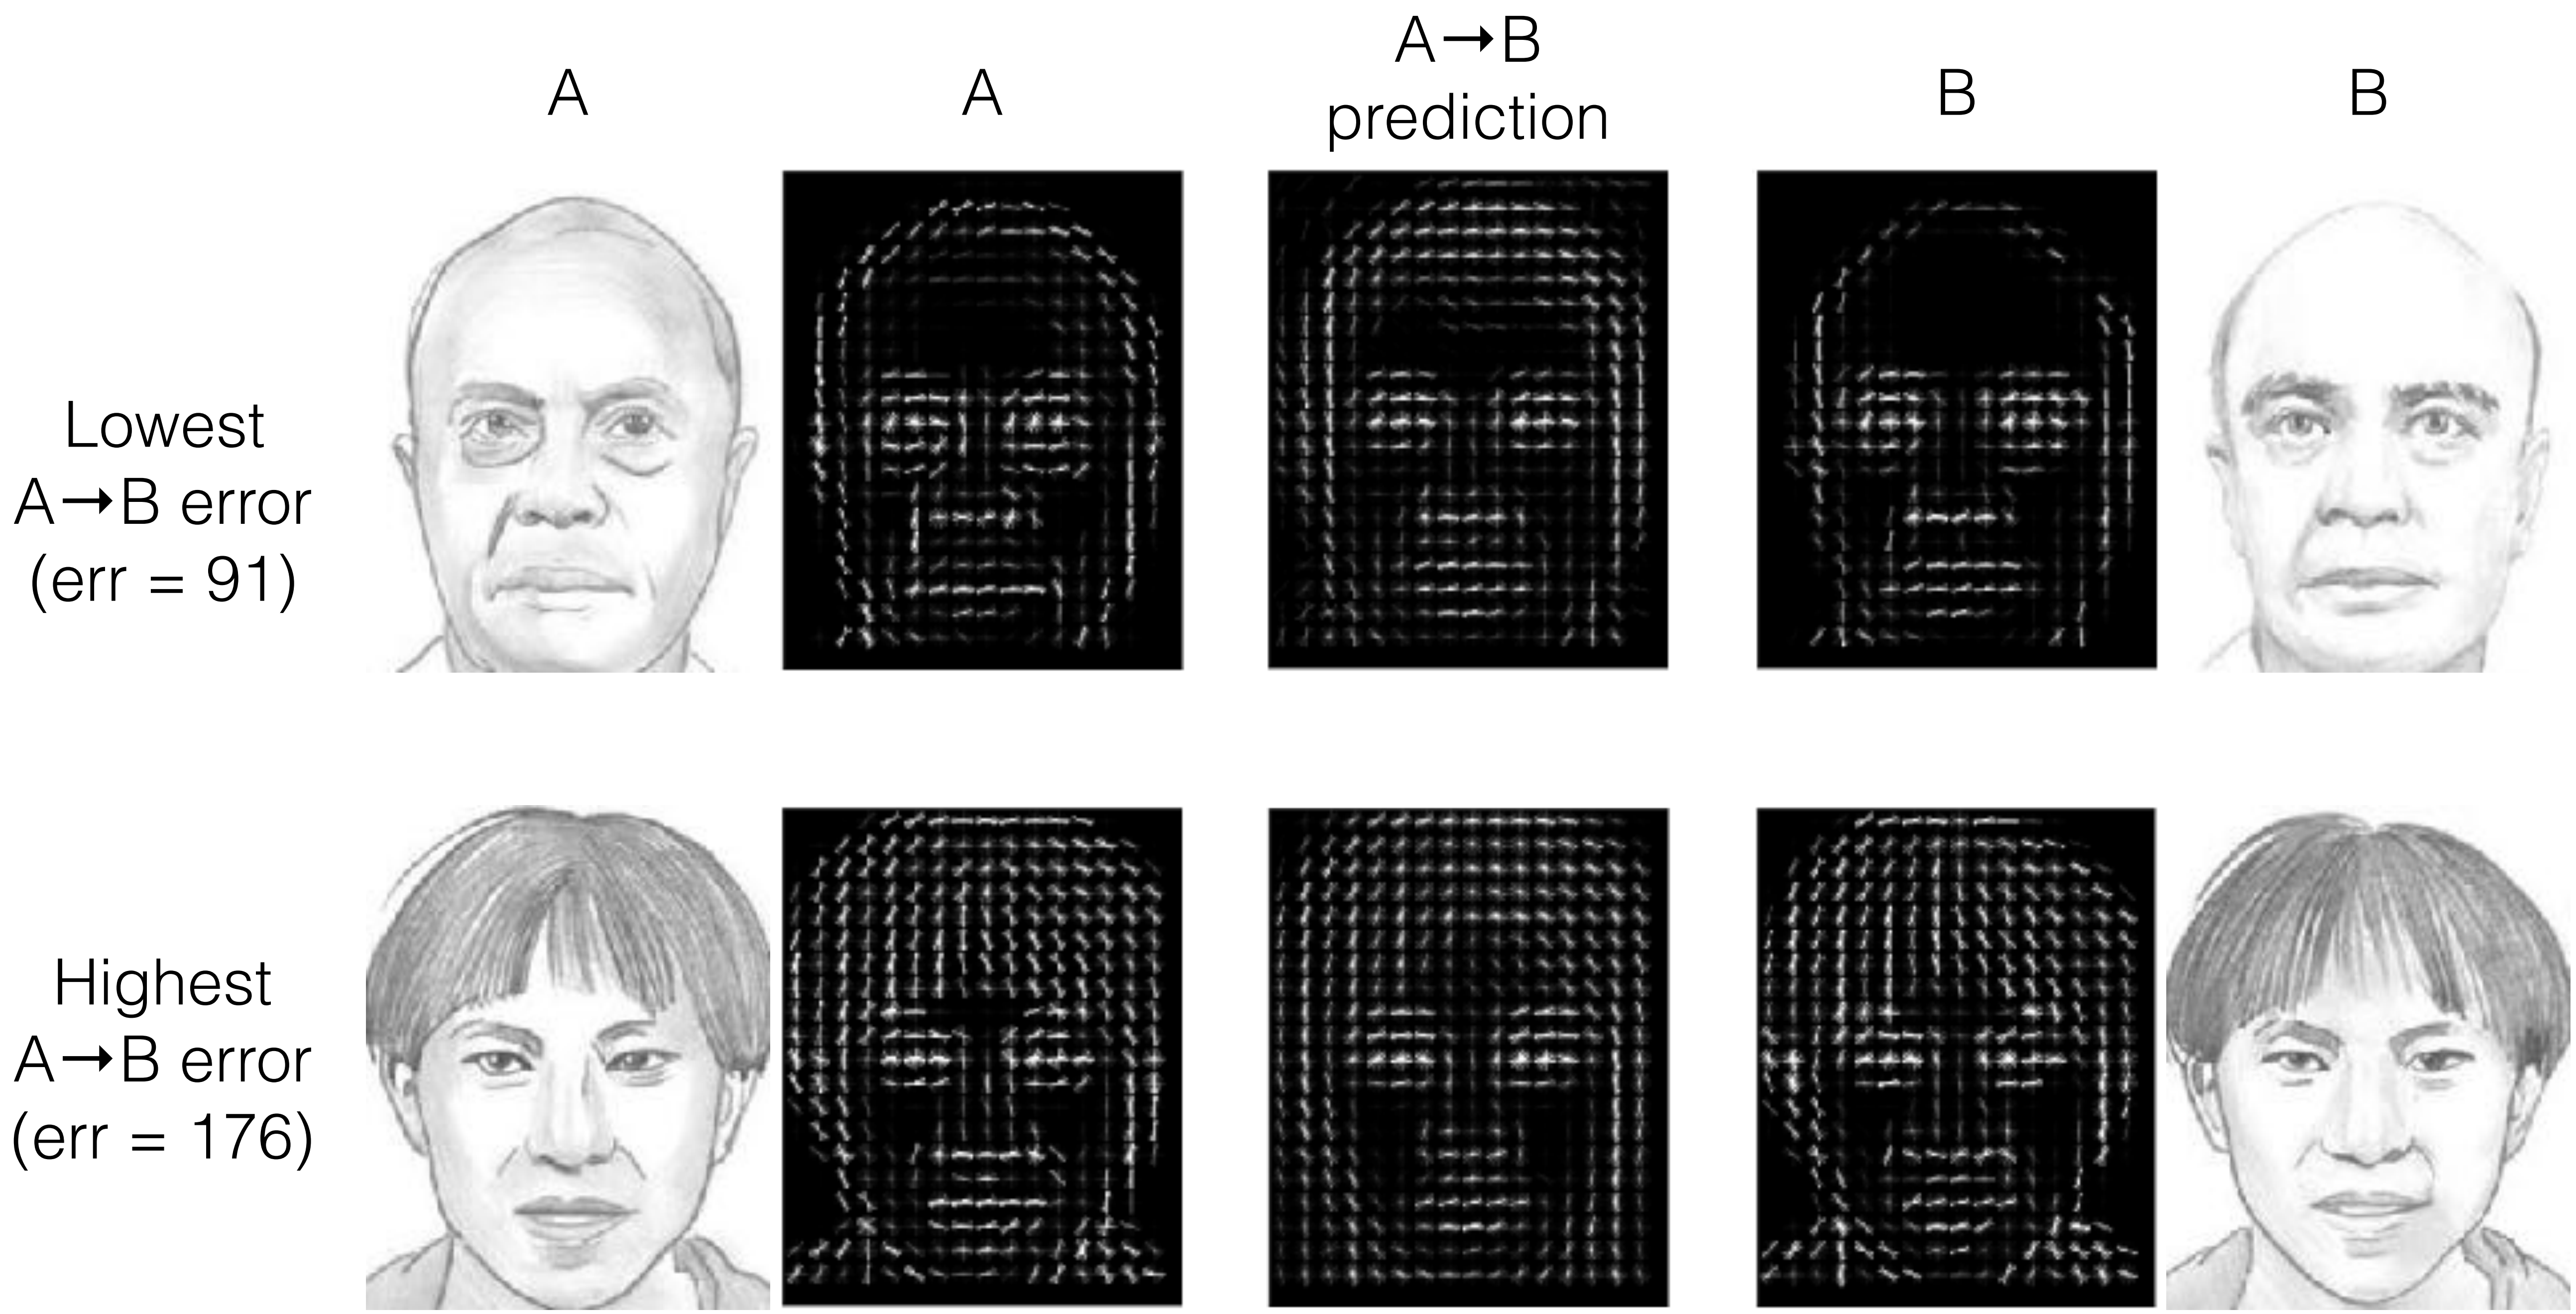

Which artists have similar style?

- For each pair of artists X→Y, measure average prediction error.

A and B are most different

D and E are most similar

	A	B	C	D	E
A	0	129.52	119.99	119.27	125.82
B	129.52	0	120.8	121.32	122.95
C	119.99	120.8	0	114.05	121.02
D	119.27	121.32	114.05	0	104.03
E	125.82	122.95	121.02	104.03	0

Which artists have similar style?

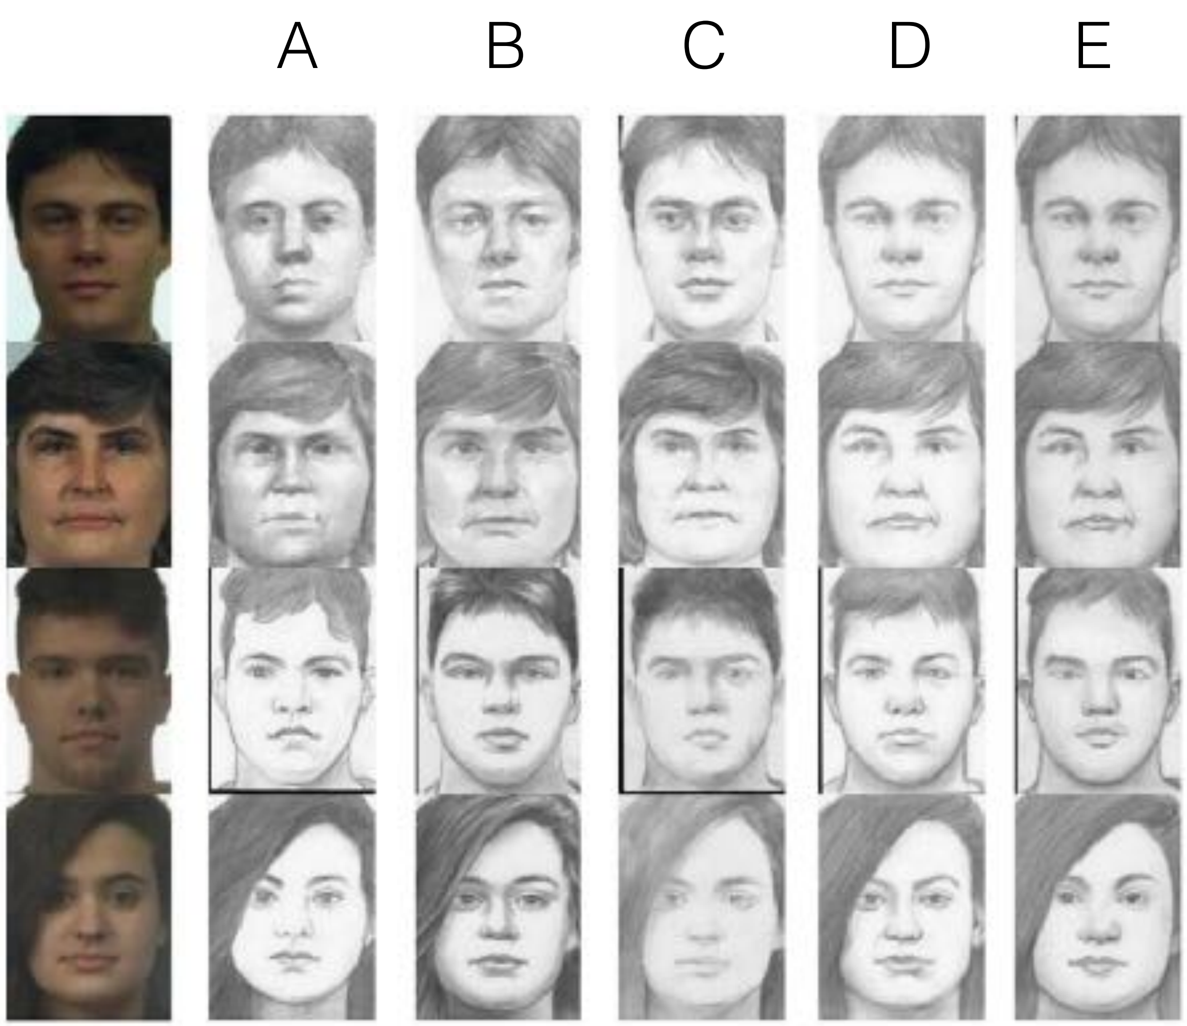

	A	B	C	D	E
A	0	129.52	119.99	119.27	125.82
B	129.52	0	120.8	121.32	122.95
C	119.99	120.8	0	114.05	121.02
D	119.27	121.32	114.05	0	104.03
E	125.82	122.95	121.02	104.03	0

Chaining

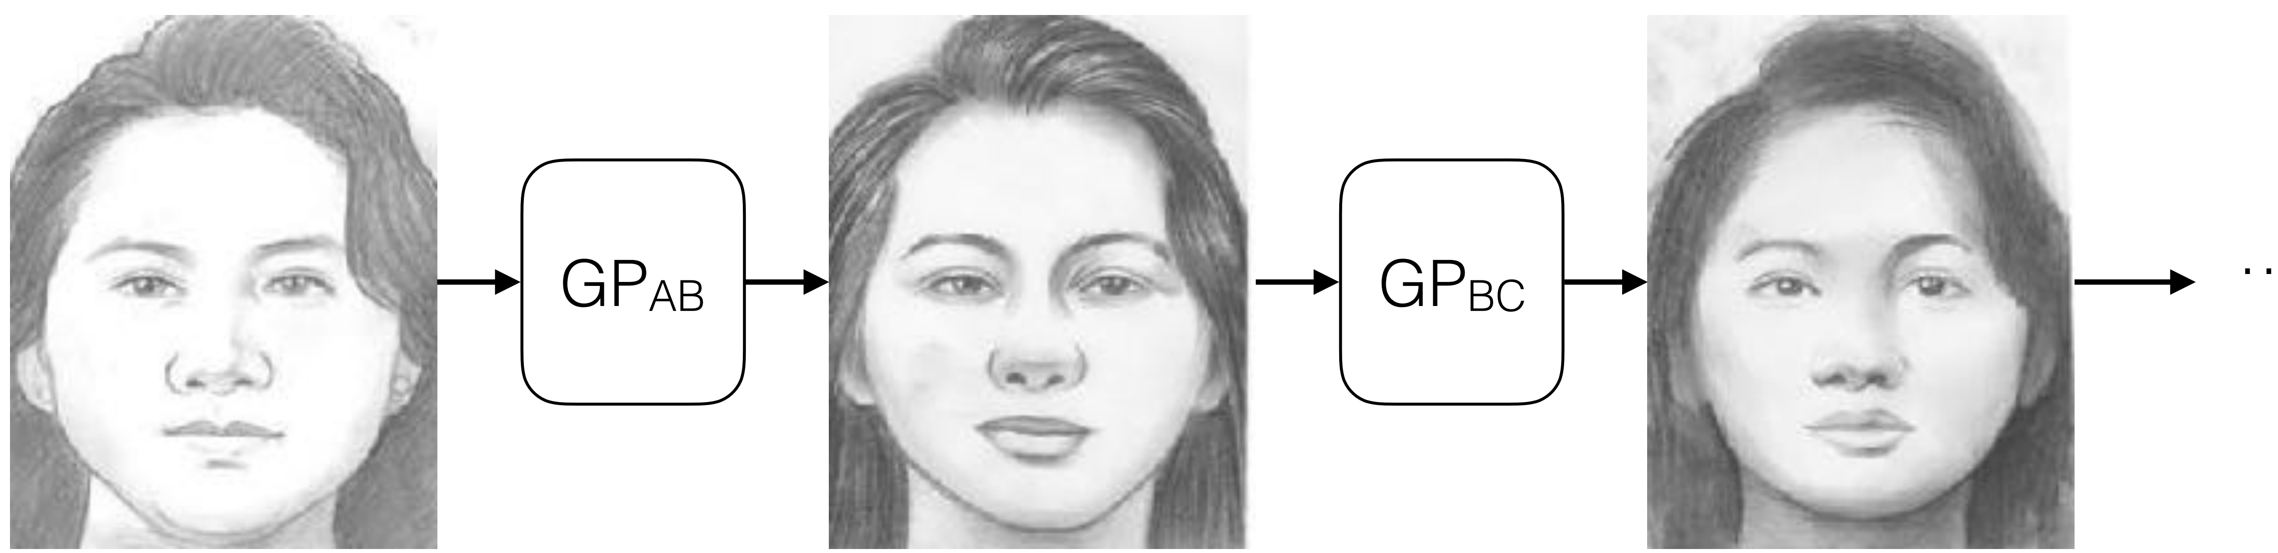

input sketch from artist A

reconstructed sketch in the style of artist B

reconstructed sketch in the style of artist C

Chaining

- Does chaining reduce error?
- Average E→C error is 121.
- avg_err(E→D) = 104
  avg_err(D→C) = 114
- Compare error between E→C vs E→D→C chain.

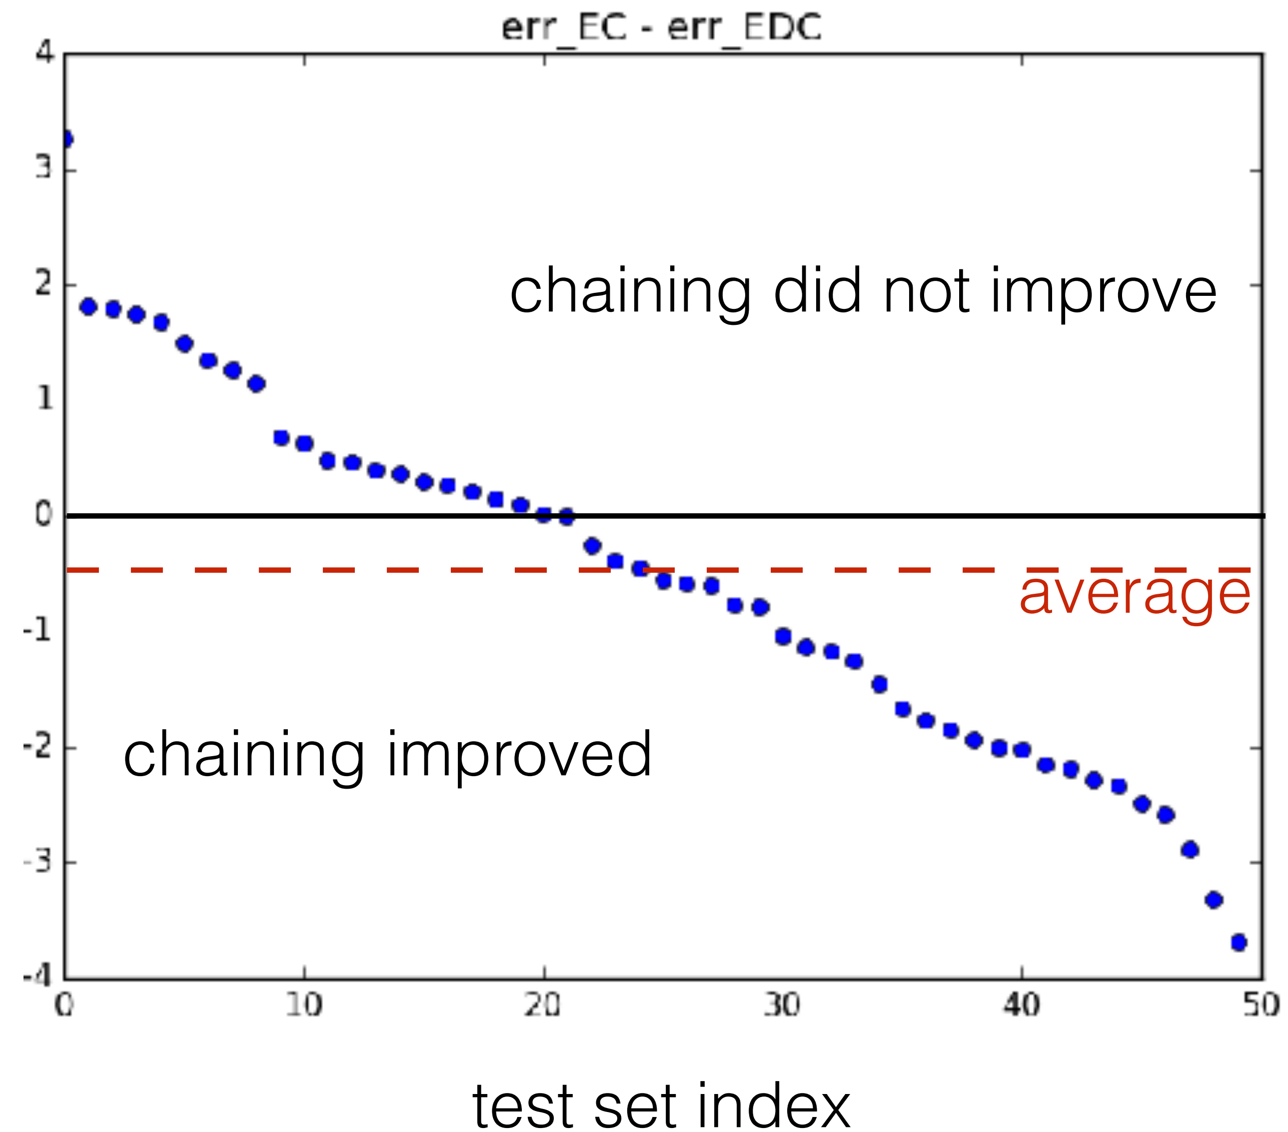

Chaining

- Does chaining reduce error?
- Average E→C error is 121.
- avg_err(E→D) = 104
  avg_err(D→C) = 114
- Compare error between E→C vs E→D→C chain.

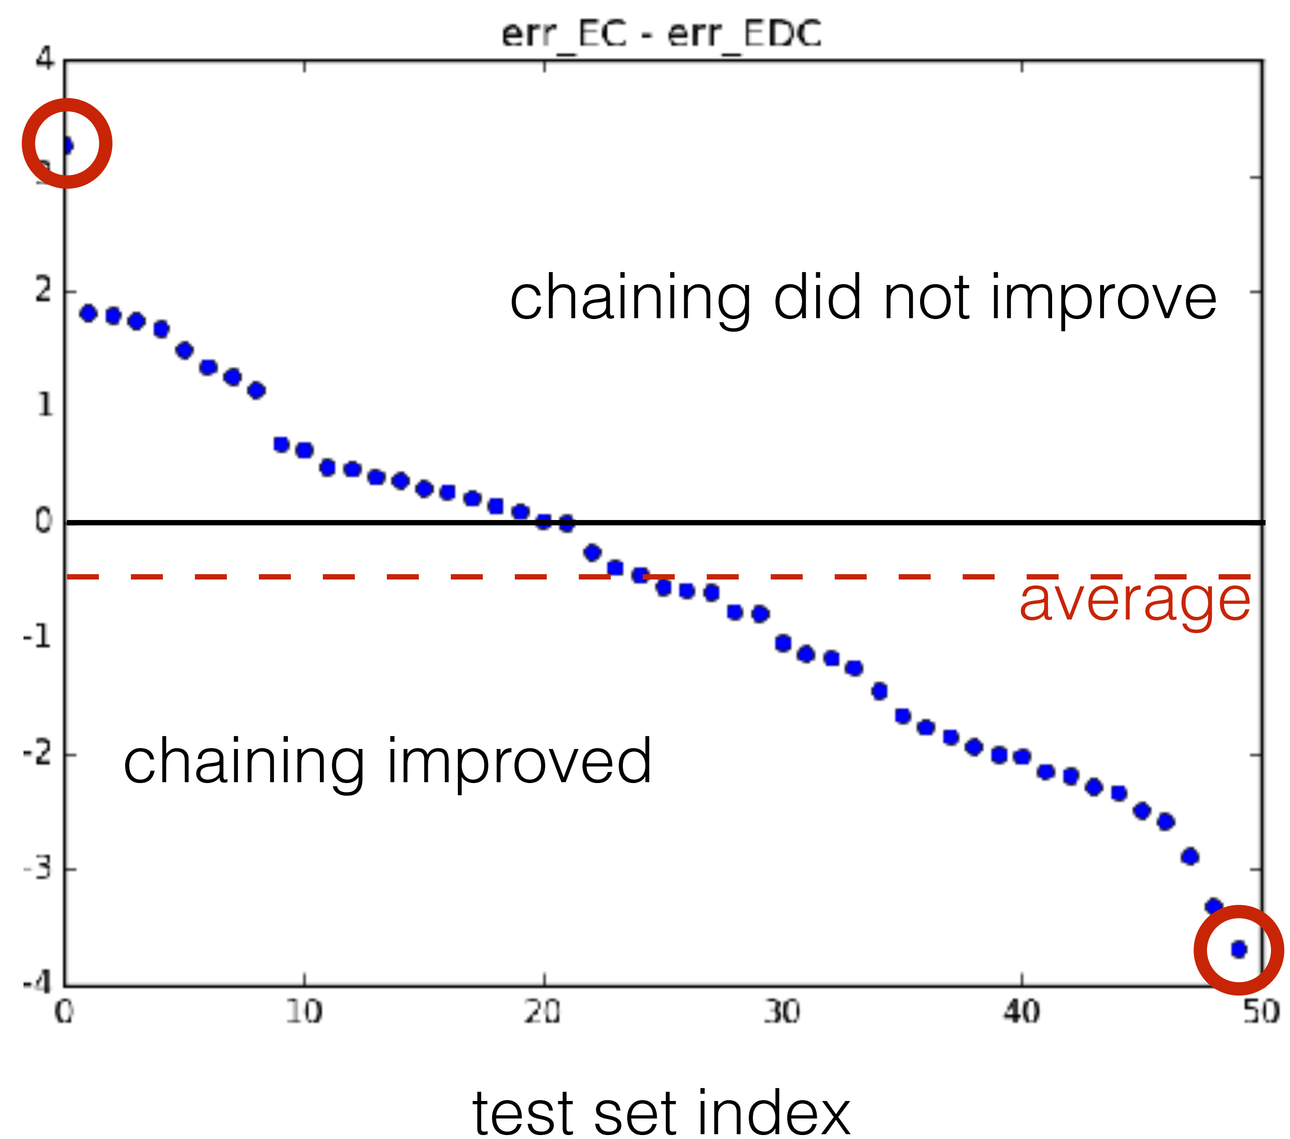

Chaining

(best and worst case example)

E D C

chaining improved the most

chaining improved the least

Chaining

(best test case example)

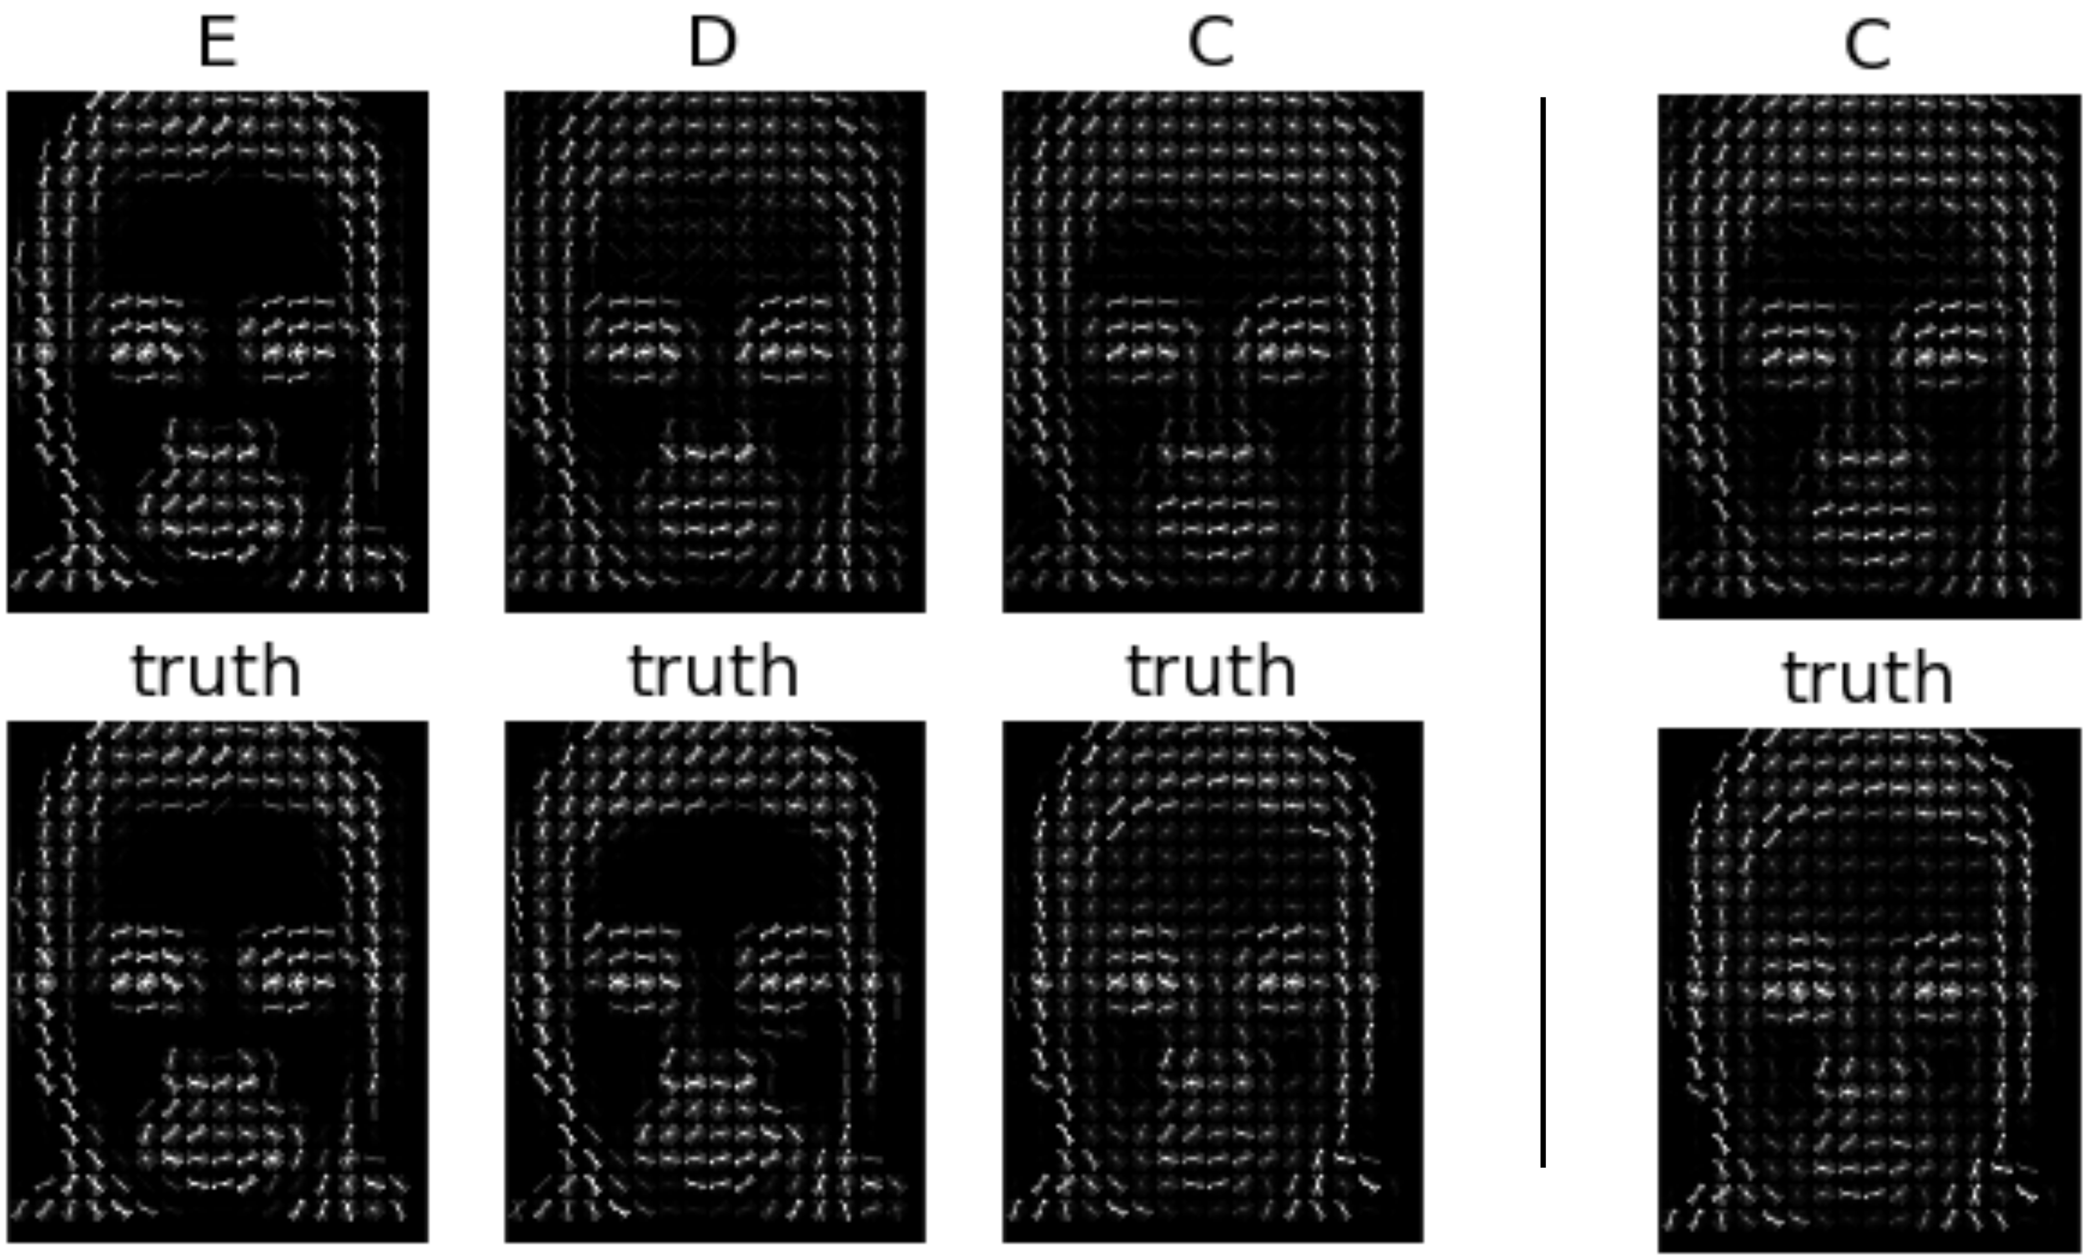

E→D→C

E→C

Chaining

(worst test case example)

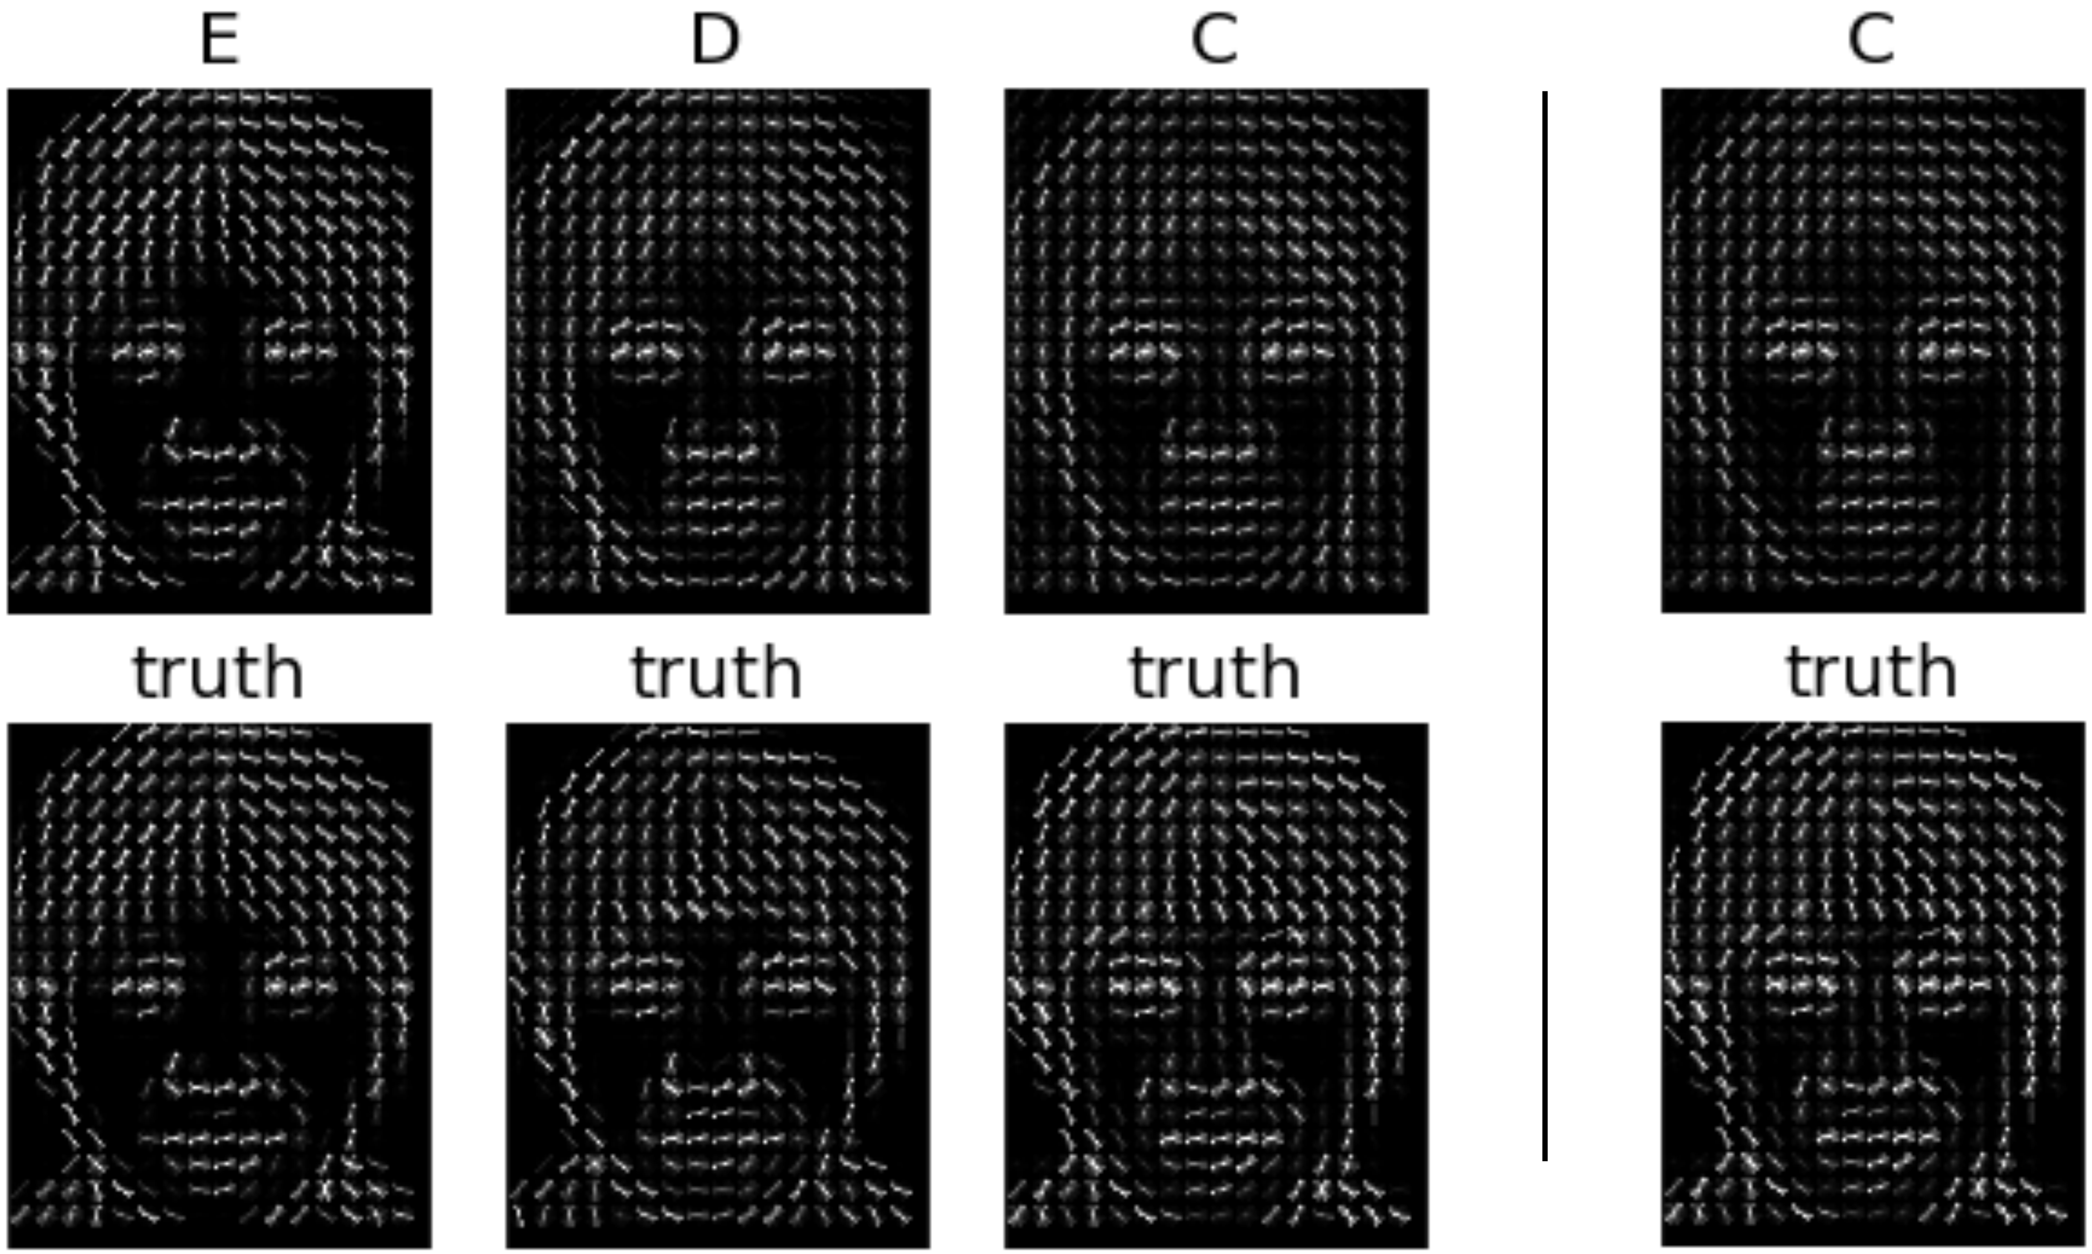

E→D→C

E→C

Chaining

- Differences are too slight to see a difference in HOG images.

- Error is ~ 100. Difference in error ~3. Most extreme gains and losses are only about 3% different.

- I’m not convinced chaining significantly improves results.

Conclusions

- Gaussian Processes can be used to learn the relation between sketch images.
- It's not perfect. More data or a different feature space may help.
- The authors' use of multi-task learning helped alleviate the problem of small data.