

Philipp Hofer
Institute of
Networks and Security

@ philipp.hofer@ins.jku.at
🌐 <https://www.digidow.eu/>

August 2021

Face recognition: Combining embeddings



Technical Report

Christian Doppler Laboratory for
Private Digital Authentication in the Physical World

This work has been carried out within the scope of Digidow, the Christian Doppler Laboratory for Private Digital Authentication in the Physical World, funded by the Christian Doppler Forschungsgesellschaft, 3 Banken IT GmbH, Kepler Universitätsklinikum GmbH, NXP Semiconductors Austria GmbH, and Österreichische Staatsdruckerei GmbH.

Contents

| | |
|--|-----------|
| 1. Motivation | 3 |
| 2. Dataset | 3 |
| 3. Average | 4 |
| 3.1 Baseline - Face image | 4 |
| 3.2 Baseline - mask image | 6 |
| 3.3 Combination of mask + face image | 7 |
| 3.4 Different arithmetic methods of combining embeddings | 7 |
| 3.5 Multiple angles | 8 |
| 3.5.1 5 images | 8 |
| 3.5.2 25 images | 8 |
| 3.6 Amount of images limit | 9 |
| 4. Voting | 15 |
| 5. KMeans | 25 |
| 5.1 Summary | 27 |

1. Motivation

In order to increase the accuracy of SOTA face recognition pipelines, intuitively it would make sense to not only use a single image as reference embedding (*template*), but combine multiple embeddings from different images (different pose, angle, setting) to create a more accurate and robust template. In order to objectively evaluate our different proposed combinations of embeddings, we would benefit from having a single metric to tell how well the template is performing on our dataset. For certain applications (e.g. opening doors) a low false-positive rate is required, while in other situations (e.g. sensor contacting PIA's) a low false-negative rate is required. Therefore, in this document we try to balance these different approaches by using the harmonic mean of recall and precision. There are multiple ways of combining different embeddings:

1. Average (as proposed by e.g. ArcFace)
2. Voting (perform comparison on n different images, everyone votes whether there is a match; if $m \leq n$ votes say *match*, it is actually counted as match)

2. Dataset

For evaluation, we created a dataset by mounting a camera in our printer room and collecting images whenever a person is detected in front of the camera.

- There are 13 different people in the dataset.
- There are between 5 and 210 images of every person, with an average of 79.3 images/person.

From this dataset, we pick a person and copy the images to a new location:

```
1 person = "hofer"
2 person_path = os.path.join(base_path, person)
3 new_location = os.path.expanduser("~/datasets")
4 new_location = os.path.join(new_location, person)
5 !rm -rf $new_location
6 !cp -r $person_path $new_location
```

As discussed in *Simple heuristics to reduce false-positives in face recognition*, we remove all images where the head rotation is bad:

```
1 for file in os.listdir(new_location):
2     img_path = os.path.join(new_location, file)
3     if os.path.isfile(img_path):
4         if rec.get_eye_distance(img_path) < eye_distance:
5             ! rm $img_path
6         elif rec.get_eyes_mouth_distance(img_path) < eye_mouth_distance:
7             ! rm $img_path
```

We have 166 images of this person.

3. Average

In this setting, we simply calculate the embedding of multiple images, take their average and use this average as template.

3.1 Baseline - Face image

In order to be able to compare the performance of different ways of combining embeddings, we first calculate the baseline where we use a single, full-frontal image as template.

```
1 template_single_path = os.path.join(base_path, "atemplates_variations/hofer/2_face
  -5/2021-05-25-144227_30.jpg")
2 template_single = rec.get_emb(template_single_path)
```

After processing the image, the face looks like this:

```
1 from PIL import Image
2 # Image.fromarray(rec.vis._get_img(template_single_path, 1, "", True), 'RGB')
```

Next, we need to calculate the harmonic mean of a given template with respect to the whole dataset.

```
1 from prettytable import PrettyTable
2 def get_confusion(threshold, same_person, different_person):
3     tp, tn, fp, fn = 0, 0, 0, 0
4
5     for p in same_person:
6         if p > threshold:
7             fn += 1
8         else:
9             tp += 1
10    for p in different_person:
11        if p <= threshold:
12            fp += 1
13        else:
14            tn += 1
15    return fp/len(different_person), fn/len(same_person), tp/len(same_person), tn/
    len(different_person)
16 def print_confusion(fp, fn, tp, tn):
17    x = PrettyTable()
18    x.title = 'Predicted_condition'
19    x.field_names = ["_", "", "Positive", "Negative"]
20    x.add_row(["Actual", "Positive", "{:.2f}".format(tp), "{:.2f}".format(fn)])
21    x.add_row(["Condition", "Negative", "{:.2f}".format(fp), "{:.2f}".format(tn)])
22    print(x)
23
24 def get_harmonic_mean(threshold, same, different):
25    fp, fn, tp, tn = get_confusion(threshold, same, different)
26    if tp+fp == 0:
27        return 0
28    if tp + fn == 0:
29        return 0
30    prec = tp / (tp + fp)
31    recall = tp / (tp + fn)
32    harmonic_mean = 2*(prec*recall/(prec+recall))
33    return harmonic_mean
```

```

34 def optimize_harmonic_mean(same, different):
35     highest_h_mean = 0
36     threshold_highest_h_mean = -1
37     upper_limit = int(max(max(same), max(different))*1000)
38     for i in range(0, upper_limit):
39         i = i/1000
40         harmonic_mean = get_harmonic_mean(i, same, different)
41         if harmonic_mean > highest_h_mean:
42             highest_h_mean = harmonic_mean
43             threshold_highest_h_mean = i
44     return threshold_highest_h_mean, highest_h_mean
45 def harmonic_mean(template):
46     same = []
47     for file in os.listdir(new_location):
48         img_path = os.path.join(new_location, file)
49         if os.path.isfile(img_path):
50             emb = rec.get_emb(img_path)[0]
51             same.append(rec.get_score(template, emb))
52     different = []
53     for file in os.listdir(base_path):
54         person_path = os.path.join(base_path, file)
55         if os.path.isdir(person_path) and file != person and not file in
exclude_dirs:
56             for img in os.listdir(person_path):
57                 img_path = os.path.join(person_path, img)
58                 if rec.get_eye_distance(img_path) < eye_distance:
59                     continue
60                 if rec.get_eyes_mouth_distance(img_path) < eye_mouth_distance:
61                     continue
62                 emb = rec.get_emb(img_path)[0]
63                 different.append(rec.get_score(template, emb))
64     threshold_highest_h_mean, highest_h_mean = optimize_harmonic_mean(same,
different)
65     return highest_h_mean, same, different, threshold_highest_h_mean
66 def print_harmonic_mean(h_mean, threshold):
67     print("Harmonic mean: {:.3f} (with a threshold of {})".format(h_mean,
threshold))

```

```

1 h_mean, same, different, threshold = harmonic_mean(template_single)
2 print_harmonic_mean(h_mean, threshold)

```

```

1     Harmonic mean: 0.852 (with a threshold of 1.54)

```

We can also plot our confusion matrix:

```

1  +-----+
2  |               Predicted condition               |
3  +-----+-----+-----+-----+
4  |               | Positive | Negative |
5  +-----+-----+-----+-----+
6  | Actual | Positive | 0.89 | 0.11 |
7  | Condition | Negative | 0.20 | 0.80 |
8  +-----+-----+-----+-----+

```

In our current setup, we want to minimize FP, thus we lowered the threshold to 1.2. This would result in the following confusion matrix:

```

1  +-----+

```

| | | Predicted condition | |
|------------------|----------|---------------------|----------|
| | | Positive | Negative |
| Actual Condition | Positive | 0.30 | 0.70 |
| Actual Condition | Negative | 0.00 | 1.00 |

3.2 Baseline - mask image

Due to the current pandemic, most of the images in the dataset depicts a person wearing a face mask. Thus, it might be interesting, if using an image where the person wears a face mask increases our performance. Even though the template image more closely resembles the test images, we expect the performance to decrease, as the mask hides more than half of the face, and thus information is lost.

```

1 template_singlemask_path = os.path.join(base_path, "atemplates_variations/hofer/1
  _different-setting-50/2021-05-25-145821_26.jpg")
2 template_singlemask = rec.get_emb(template_singlemask_path)

```

After processing the image, the face looks like this:

```

1 from PIL import Image
2 Image.fromarray(rec.vis._get_img(template_singlemask_path, 1, "", True), 'RGB')

```



```

1 h_mean_mask, same_mask, different_mask, threshold_mask = harmonic_mean(
  template_singlemask)
2 print_harmonic_mean(h_mean_mask, threshold_mask)

```

```

1 Harmonic mean: 0.850 (with a threshold of 1.637)

```

| | | Predicted condition | |
|------------------|----------|---------------------|----------|
| | | Positive | Negative |
| Actual Condition | Positive | 0.88 | 0.12 |
| Actual Condition | Negative | 0.19 | 0.81 |

The harmonic mean decreased by 0.2 percent points, which supports our expectation of reduced accuracy due to lower amount of information.

3.3 Combination of mask + face image

```
1 emb_combined = np.mean( np.array([ template_single, template_singlemask ]), axis=0
    )
```

```
1 h_mean_comb, same_comb, different_comb, threshold_comb = harmonic_mean(
    emb_combined)
2 print_harmonic_mean(h_mean_comb, threshold_comb)
```

```
1     Harmonic mean: 0.870 (with a threshold of 1.267)
```

Face: 0.85 Mask: 0.85 Face+Mask: 0.87

3.4 Different arithmetic methods of combining embeddings

RoyChowdhury et al. proposed in *One-to-many face recognition with bilinear CNNs* instead of the classical usage of the average for each dimension of an embedding the use of the maximum. They argue, that this will reduce the overfit on the dominant angle and thus create a more robust embedding. Based on this, we evaluated four different arithmetic methods of combining embeddings.

```
1 embs = []
2 for file in os.listdir("/home/philipp/Nextcloud/printerroom-dataset/
    atemplates_variations/hofer/1_different-setting-50"):
3     abs_path = os.path.join("/home/philipp/Nextcloud/printerroom-dataset/
    atemplates_variations/hofer/1_different-setting-50", file)
4     if len(rec.get_emb(abs_path)) == 1:
5         embs.append(rec.get_emb(abs_path)[0])
```

```
1 print("Mean")
2 emb_combined = np.mean( np.array(embs), axis=0 )
3 h_mean_comb, same_comb, different_comb, threshold_comb = harmonic_mean(
    emb_combined)
4 print_harmonic_mean(h_mean_comb, threshold_comb)
```

```
1     Mean
2     Harmonic mean: 0.914 (with a threshold of 0.996)
```

```
1 print("Max")
2 emb_combined = np.max( np.array(embs), axis=0 )
3 h_mean_comb, same_comb, different_comb, threshold_comb = harmonic_mean(
    emb_combined)
4 print_harmonic_mean(h_mean_comb, threshold_comb)
```

```
1     Max
2     Harmonic mean: 0.868 (with a threshold of 4.178)
```

```

1 print("Median")
2 emb_combined = np.median( np.array(embs), axis=0 )
3 h_mean_comb, same_comb, different_comb, threshold_comb = harmonic_mean(
    emb_combined)
4 print_harmonic_mean(h_mean_comb, threshold_comb)

```

```

1     Median
2     Harmonic mean: 0.899 (with a threshold of 0.955)

```

```

1 print("Min")
2 emb_combined = np.min( np.array(embs), axis=0 )
3 h_mean_comb, same_comb, different_comb, threshold_comb = harmonic_mean(
    emb_combined)
4 print_harmonic_mean(h_mean_comb, threshold_comb)

```

```

1     Min
2     Harmonic mean: 0.764 (with a threshold of 3.999)

```

3.5 Multiple angles

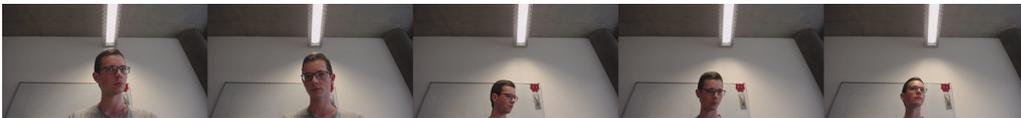
3.5.1 5 images

Let's see if the accuracy improves if we use 5 images with various angles.

```

1 comb_images(os.path.join(base_path, "atemplates_variations/hofer/2_face-5"))

```



```

1 folder_2 = os.path.join(base_path, "atemplates_variations/hofer/2_face-5")
2 emb_2 = rec.get_combined_embedding(folder_2)
3 h_mean_2, same_2, different_2, threshold_2 = harmonic_mean(emb_2)
4 print_harmonic_mean(h_mean_2, threshold_2)

```

```

1     Harmonic mean: 0.892 (with a threshold of 1.166)

```

Getting a higher accuracy with multiple pictures seems promising. What happens if we increase the amount of images?

3.5.2 25 images

```

1 folder_3 = os.path.join(base_path, "atemplates_variations/hofer/3_face-25")
2 emb_3 = rec.get_combined_embedding(folder_3)
3 h_mean_3, same_3, different_3, threshold_3 = harmonic_mean(emb_3)
4 print_harmonic_mean(h_mean_3, threshold_3)

```

```

1     Harmonic mean: 0.919 (with a threshold of 0.947)

```

```
1 print_harmonic_mean(h_mean_3, threshold_3)
```

```
1 Harmonic mean: 0.919 (with a threshold of 0.947)
```

Accuracy is increased again (by 2.63 percentage points).

3.6 Amount of images limit

So far, there has been a positive correlation between amount of images used for the template and its harmonic mean. In this section we test its limit. For this we need a lot of images of the same person, thus we decided to go with the LFW dataset. Let's look at some stats from the dataset.

```
1 from collections import defaultdict
2 amount_images = defaultdict(lambda: 0)
3 paths_many_images = []
4 path = "/home/philipp/datasets/lfw"
5 for file in os.listdir(path):
6     person_path = os.path.join(path, file)
7     amount_images[len(os.listdir(person_path))] += 1
8     if len(os.listdir(person_path)) > 100:
9         paths_many_images.append(person_path)
```

```
1 amount_people_below_50_images = 0
2 for i in sorted(amount_images.keys(), reverse=True):
3     if i < 50:
4         amount_people_below_50_images += amount_images[i]
5     else:
6         print("{}_{}_{}_images".format(amount_images[i], "people_have" if
7 amount_images[i] > 1 else "person_has", i))
8 print("{}_people_have_less_than_50_images".format(amount_people_below_50_images))
```

```
1 1 person has 530 images
2 1 person has 236 images
3 1 person has 144 images
4 1 person has 121 images
5 1 person has 109 images
6 1 person has 77 images
7 1 person has 71 images
8 1 person has 60 images
9 1 person has 55 images
10 1 person has 53 images
11 2 people have 52 images
12 5737 people have less than 50 images
```

We use the 5 people with more than 100 images and plot the amount of difference in their embeddings.

```
1 import matplotlib.pyplot as plt
2 plt.style.use('seaborn-whitegrid')
3 import numpy as np
4 def calc_distances(images, print_only_last_change = False, dont_print=False):
5     ret = []
6     embs = []
7     for image in images:
8         emb = rec.get_emb(image)
```

```

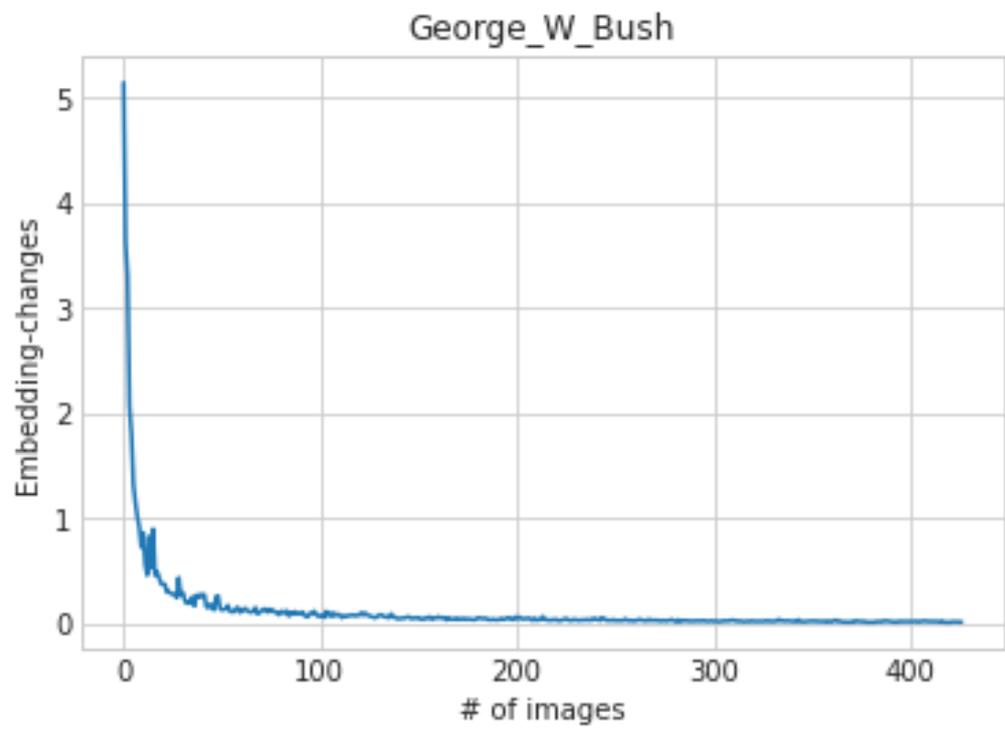
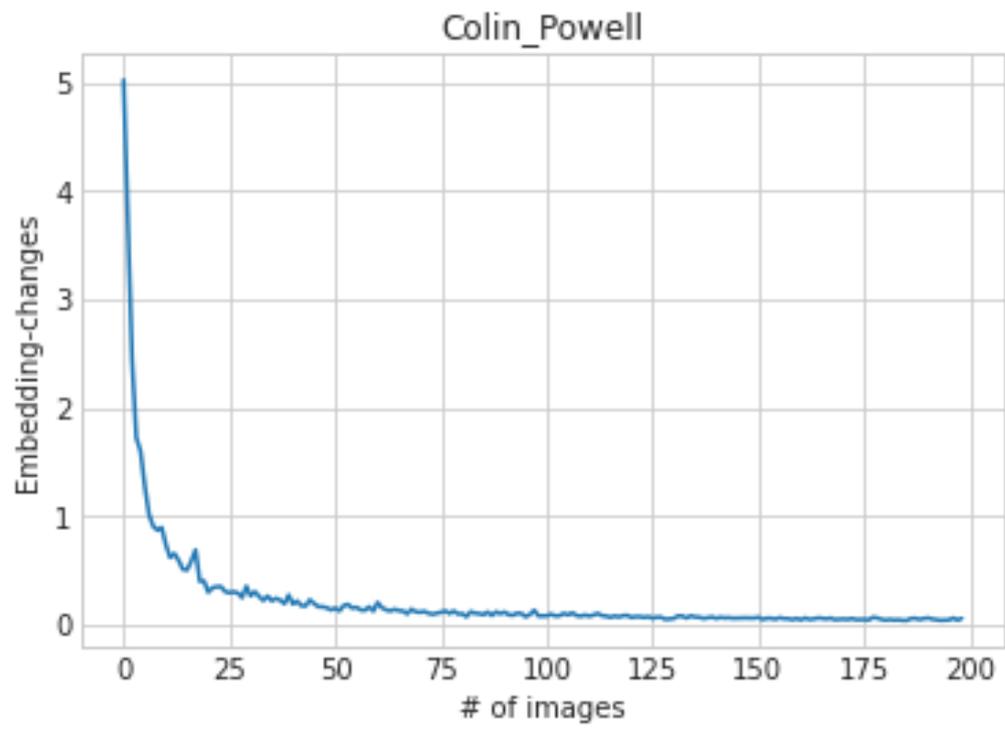
9         if len(emb) != 1:
10             if not dont_print:
11                 print("skip_{}_due_to_{}_people_detected".format(image, len(emb)))
12             else:
13                 embs.append(np.array(emb[0]))
14         combs = [embs[0]]
15         for i in range(1, len(embs)):
16             cur = []
17             for j in range(i+1):
18                 cur.append(embs[j])
19             emb_combined = np.mean( np.array(cur), axis=0 )
20             prev_combined = combs[-1]
21             combs.append(emb_combined)
22             dist = np.sum(np.absolute(emb_combined - prev_combined))
23             if not print_only_last_change or i == len(embs)-1:
24                 if not dont_print:
25                     print("After_{}_images:_Distance_to_previous:_{}".format(i+1, dist
26             ))
27             ret.append(dist)
28         return ret
29     def emb_change(path):
30         imgs = []
31         for img in sorted(os.listdir(path)):
32             img_path = os.path.join(path, img)
33             imgs.append(img_path)
34
35         dist = calc_distances(imgs, dont_print=True)
36         fig = plt.figure()
37         ax = plt.axes()
38         plt.title(os.path.basename(path))
39         plt.xlabel("#_of_images")
40         plt.ylabel("Embedding-changes")
41         ax.plot(range(len(dist)), dist)

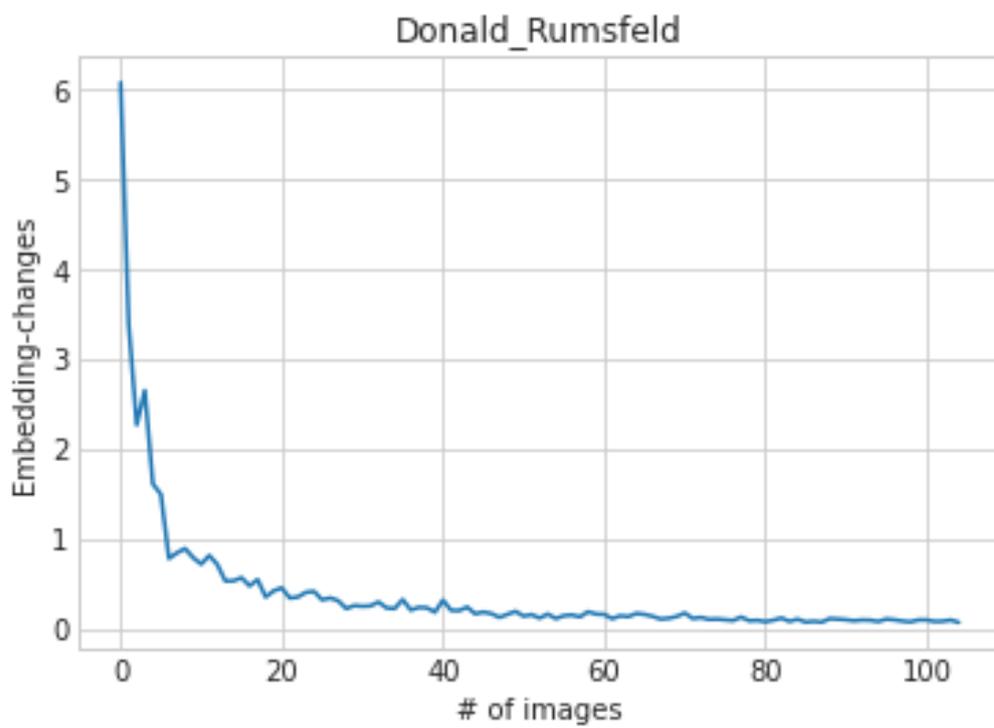
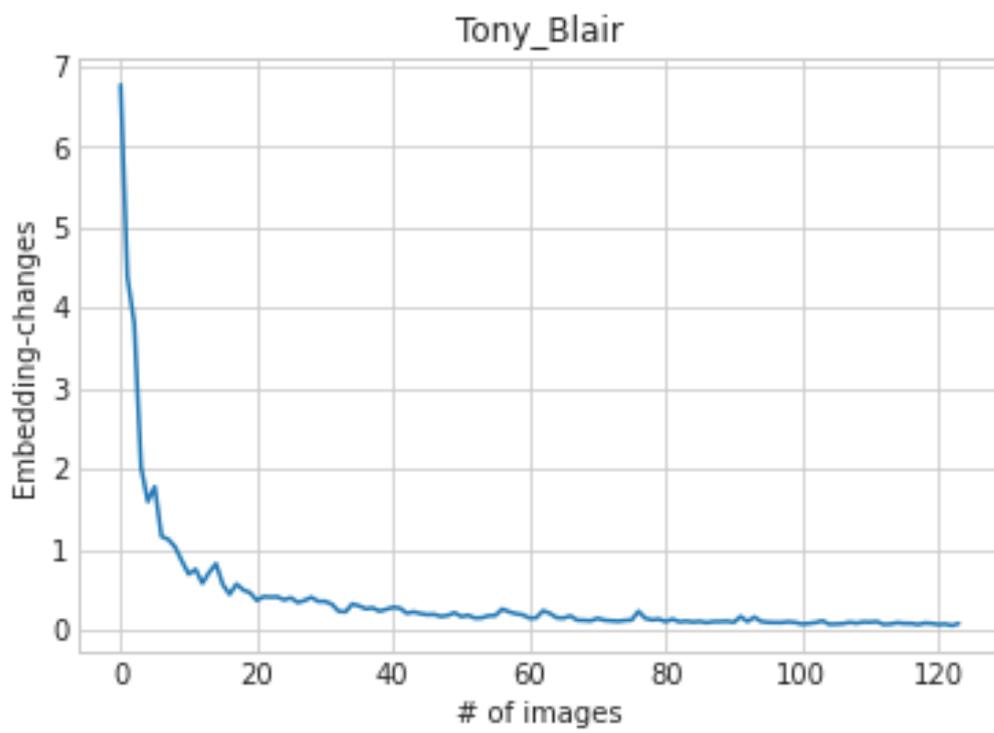
```

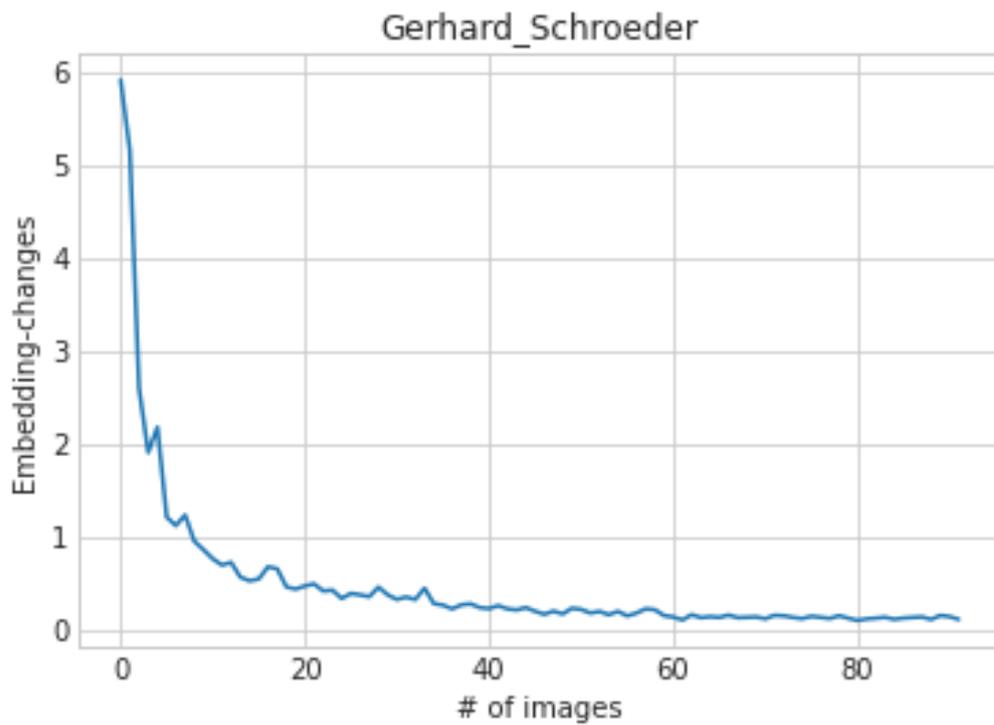
```

1     for path in paths_many_images:
2         emb_change(path)

```

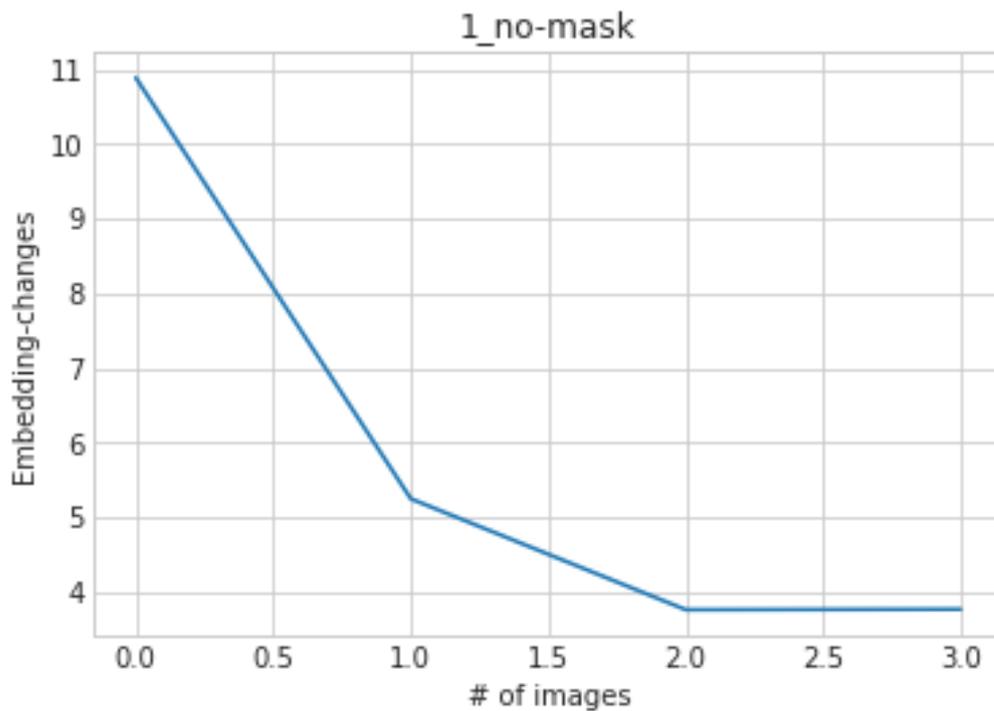




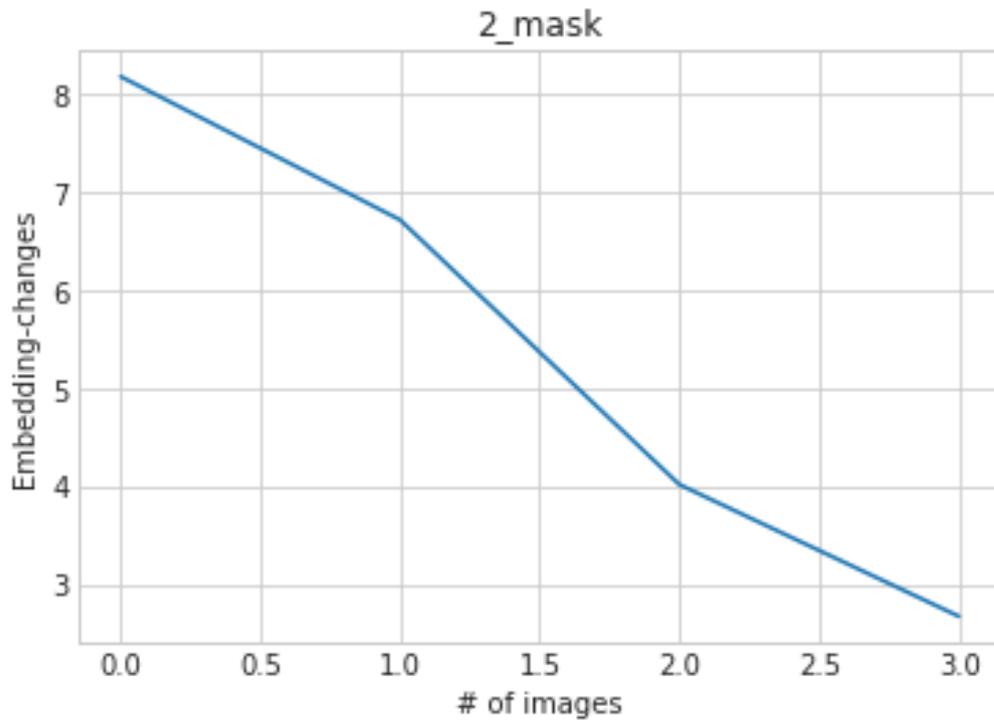


It is clearly visible, that the amount of change decreases with an increased amount of images. After roughly 20 images the combined embedding is not getting better anymore. How is the embedding changing with our templates?

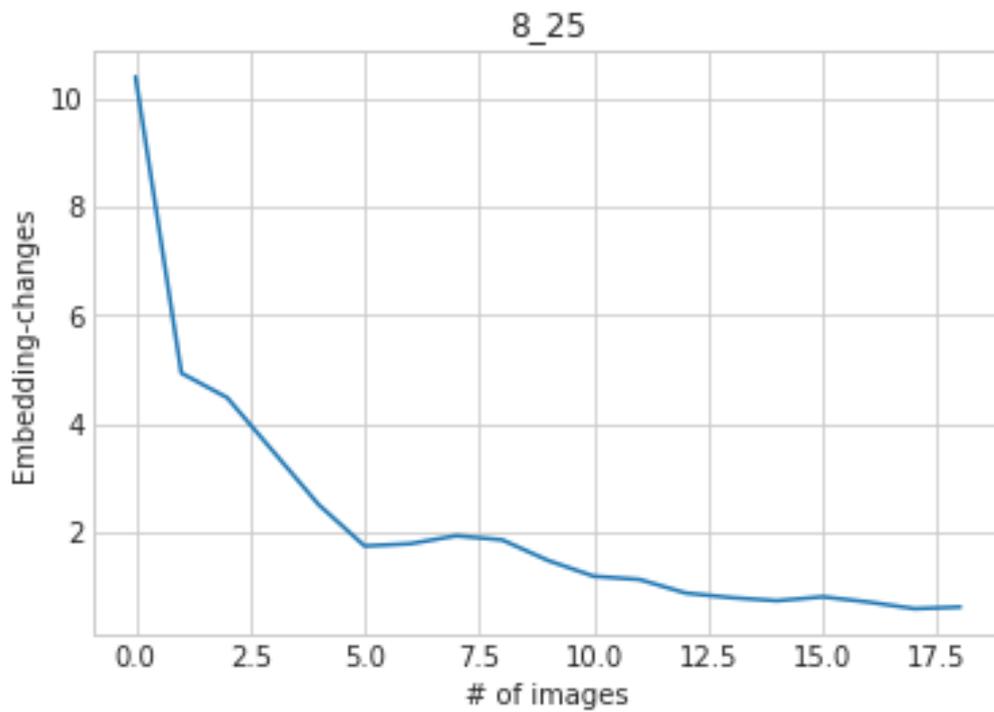
```
1 emb_change(os.path.join(base_path, "atemplates_variations/hofer2/1_no-mask"))
```



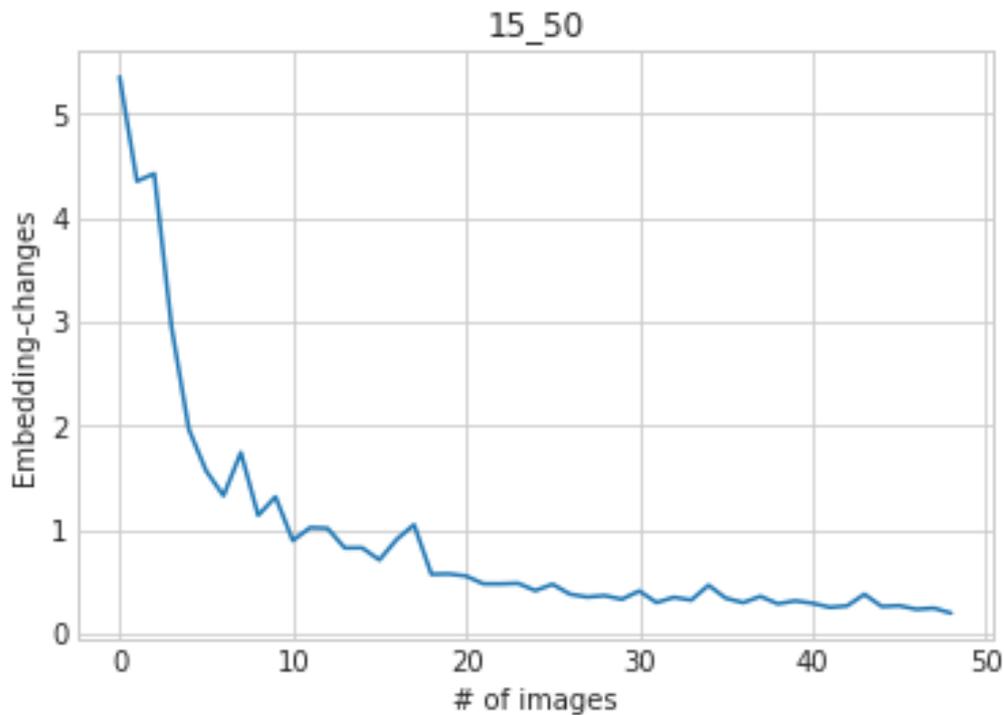
```
1 emb_change(os.path.join(base_path, "atemplates_variations/hofer2/2_mask"))
```



```
1 emb_change(os.path.join(base_path, "atemplates_variations/hofer2/8_25"))
```



```
1 emb_change(os.path.join(base_path, "templates_variations/hofer2/15_50"))
```



4. Voting

What happens if we switch from doing 1 comparison with the average template (created using n images) to doing n comparisons with the original images?

```
1 def harmonic_mean_from_conf(tp, tn, fp, fn):
2     if tp+fp == 0:
3         return 0
4     if tp + fn == 0:
5         return 0
6     prec = tp / (tp + fp)
7     recall = tp / (tp + fn)
8     harmonic_mean = 2*(prec*recall/(prec+recall))
9     return harmonic_mean
10 def optimize_amount_votes(same, different):
11     glob_min = min(min(same), min(different))
12     glob_max = max(max(same), max(different))
13
14     max_h_mean = 0
15     votesneeded_max_h_mean = 0
16
17     for votes_needed in range(glob_min, glob_max+1):
18         tp, tn, fp, fn = 0,0,0,0
19         for s in same:
20             if s >= votes_needed:
21                 tp += 1
22             else:
23                 fn += 1
```

```

24     for d in different:
25         if d >= votes_needed:
26             fp += 1
27         else:
28             tn += 1
29     h_mean = harmonic_mean_from_conf(tp, tn, fp, fn)
30     if h_mean > max_h_mean:
31         max_h_mean = h_mean
32         votesneeded_max_h_mean = votes_needed
33     return max_h_mean, votesneeded_max_h_mean

```

As a baseline, let's calculate the harmonic mean if we always predict match/non-match:

```

1 amount_same = 0
2 for file in os.listdir(new_location):
3     img_path = os.path.join(new_location, file)
4     if os.path.isfile(img_path):
5         amount_same += 1
6 amount_different = 0
7 for file in os.listdir(base_path):
8     person_path = os.path.join(base_path, file)
9     if os.path.isdir(person_path) and file != person and not file in exclude_dirs:
10        for img in os.listdir(person_path):
11            img_path = os.path.join(person_path, img)
12
13            if rec.get_eye_distance(img_path) >= eye_distance and rec.
14               get_eyes_mouth_distance(img_path) >= eye_mouth_distance:
15                amount_different += 1
16 # always predict true
17 tp=amount_same
18 fp=amount_different
19 fn,tn = 0,0
20 print("Always predicting true would yield a harmonic mean of {}".format(
21     harmonic_mean_from_conf(tp, tn, fp, fn)))
22 # always predict false
23 tn = amount_different
24 fn = amount_same
25 tp, fp = 0, 0
26 print("Always predicting false would yield a harmonic mean of {}".format(
27     harmonic_mean_from_conf(tp, tn, fp, fn)))

```

```

1     Always predicting true would yield a harmonic mean of 0.3528161530286929
2     Always predicting false would yield a harmonic mean of 0

```

```

1 local_cache = dict()

```

```

1 folder = os.path.join(base_path, "atemplates_variations/hofer/3_face-25")
2 def get_embs_from_path(path):
3     if path in local_cache:
4         return local_cache[path]
5
6     embs = []
7     for file in sorted(os.listdir(path)):
8         img_path = os.path.join(path, file)
9         embs.append(rec.get_emb(img_path)[0])
10
11     local_cache[path] = embs

```

```

12     return embs
13 def get_amount_votes(img, template_path, threshold):
14     amount_votes = 0
15     probe = rec.get_emb(img)[0]
16     count = 0
17     for emb in get_embs_from_path(template_path):
18         score = rec.get_score(probe, emb)
19         if score < threshold:
20             amount_votes += 1
21         count += 1
22     return amount_votes

```

```

1 output = []
2 h_means = []
3 for threshold in np.arange(0.5,2,0.1):
4     same = []
5     for file in os.listdir(new_location):
6         img_path = os.path.join(new_location, file)
7         if os.path.isfile(img_path):
8             same.append(get_amount_votes(img_path, folder, threshold))
9     different = []
10    for file in os.listdir(base_path):
11        person_path = os.path.join(base_path, file)
12        if os.path.isdir(person_path) and file != person and not file in
13        exclude_dirs:
14            for img in os.listdir(person_path):
15                img_path = os.path.join(person_path, img)
16                if rec.get_eye_distance(img_path) < eye_distance:
17                    continue
18                if rec.get_eyes_mouth_distance(img_path) < eye_mouth_distance:
19                    continue
20                different.append(get_amount_votes(img_path, folder, threshold))
21    h_mean, amount_votes = optimize_amount_votes(same,different)
22    h_means.append(h_mean)
23    output.append("Threshold={:.2f}: h_mean: {:.4f}_ (with {}_{}_needed)".format(
24        threshold, h_mean, amount_votes, "vote" if amount_votes==1 else "votes"))

```

```

1 - Threshold=0.50: h_mean: 0.3528 (with 0 votes needed)
2 - Threshold=0.60: h_mean: 0.3528 (with 0 votes needed)
3 - Threshold=0.70: h_mean: 0.3528 (with 0 votes needed)
4 - Threshold=0.80: h_mean: 0.3528 (with 0 votes needed)
5 - Threshold=0.90: h_mean: 0.6033 (with 1 vote needed)
6 - Threshold=1.00: h_mean: 0.8317 (with 1 vote needed)
7 - **Threshold=1.10: h_mean: 0.8580 (with 2 votes needed)**
8 - Threshold=1.20: h_mean: 0.8520 (with 6 votes needed)
9 - Threshold=1.30: h_mean: 0.8563 (with 15 votes needed)
10 - Threshold=1.40: h_mean: 0.8333 (with 28 votes needed)
11 - Threshold=1.50: h_mean: 0.8474 (with 40 votes needed)
12 - Threshold=1.60: h_mean: 0.8446 (with 45 votes needed)
13 - Threshold=1.70: h_mean: 0.8078 (with 48 votes needed)
14 - Threshold=1.80: h_mean: 0.7481 (with 49 votes needed)
15 - Threshold=1.90: h_mean: 0.6256 (with 50 votes needed)

```

As expected, the higher the threshold is set, the more votes are needed for optimal performance. Interestingly, the highest-performing combination of hyper-parameters sets the threshold quite low (1.05) and requires only a single vote. With the previous average embedding, we achieve an 94.3% harmonic

mean with this template. With the best hyper-parameters (threshold=1.05; votes_needed=1) we achieve about the same harmonic mean (93%). Thus, we conclude that the additional processing power during inference is not significantly impacting the harmonic mean. To get more intuition, we can plot a box plot of the amount of votes images get, grouped by if face recognition successfully detected the pair correctly. We use the hyper-parameters of our best output:

```

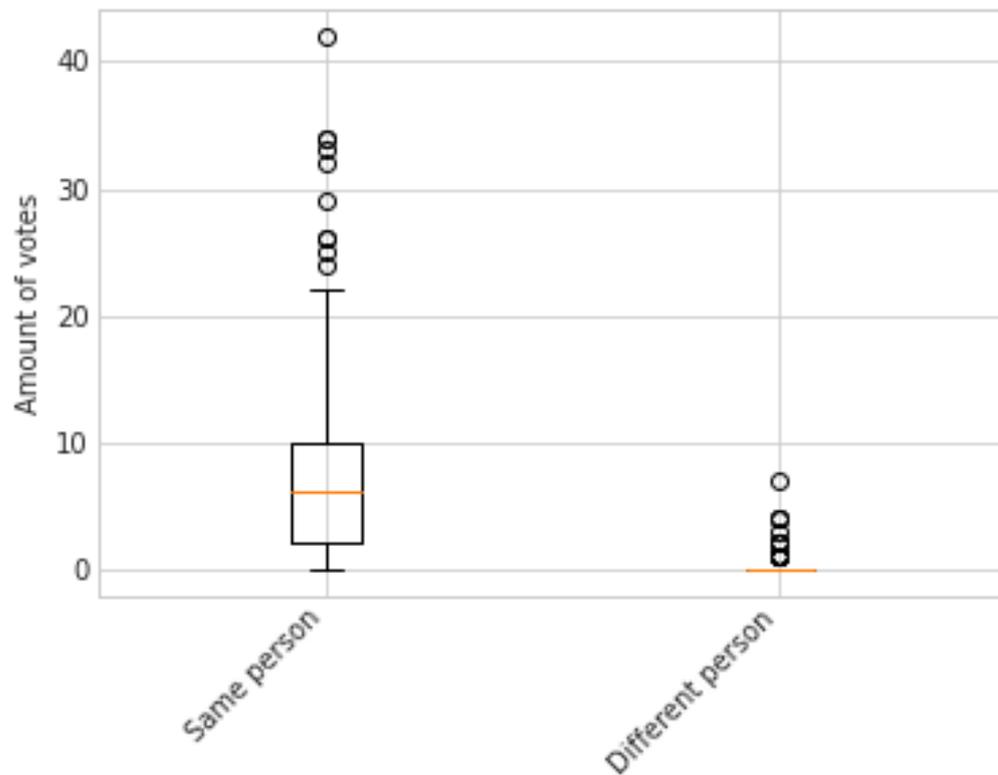
1 threshold = 1.05
2 same = []
3 for file in os.listdir(new_location):
4     img_path = os.path.join(new_location, file)
5     if os.path.isfile(img_path):
6         same.append(get_amount_votes(img_path, folder, threshold))
7 different = []
8 for file in os.listdir(base_path):
9     person_path = os.path.join(base_path, file)
10    if os.path.isdir(person_path) and file != person and not file in exclude_dirs:
11        for img in os.listdir(person_path):
12            img_path = os.path.join(person_path, img)
13            if rec.get_eye_distance(img_path) < eye_distance:
14                continue
15            if rec.get_eyes_mouth_distance(img_path) < eye_mouth_distance:
16                continue
17            different.append(get_amount_votes(img_path, folder, threshold))

```

```

1 import matplotlib.pyplot as plt
2 data = [same, different]
3 axes = plt.boxplot(data)
4 plt.xticks(np.arange(len(data))+1, ["Same_person", "Different_person"], rotation
5         =45, ha="right")
6 plt.ylabel("Amount_of_votes")
7 plt.show()

```

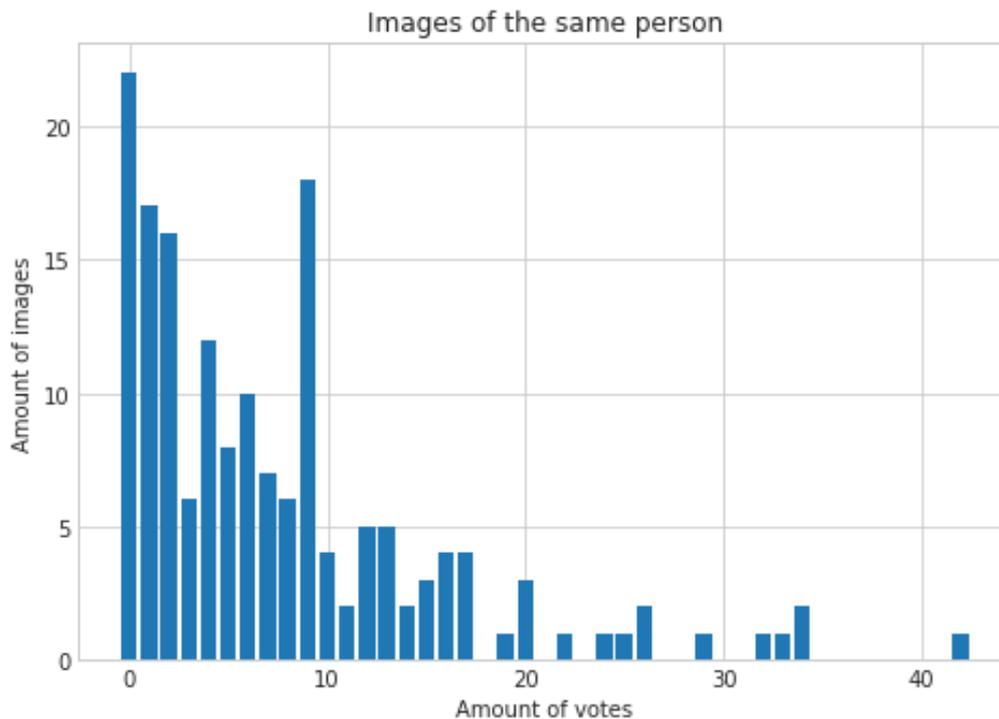


As expected, for the *same person* class most faces receive many votes. Although, there seems to be quite a few faces where only a few (or a single) vote matches the template. Let's plot this as bar chart:

```

1 from collections import defaultdict
2 amount_votes = defaultdict(lambda: 0)
3 for s in same:
4     amount_votes[s] += 1
5 data = []
6 for i in range(0, max(amount_votes.keys())+1):
7     data.append(amount_votes[i])
8 import matplotlib.pyplot as plt
9 fig = plt.figure()
10 ax = fig.add_axes([0,0,1,1])
11 ax.bar(range(len(data)), data)
12 ax.set_xlabel("Amount_of_votes")
13 ax.set_ylabel("Amount_of_images")
14 ax.set_title("Images_of_the_same_person")
15 plt.show()

```



(Interpretation: e.g. there are 22 images of the same person, which receive 0 votes) Let's look at the images with 0 positive votes:

```

1 zero_votes = []
2 for file in os.listdir(new_location):
3     img_path = os.path.join(new_location, file)
4     if get_amount_votes(img_path, folder, threshold) == 0:
5         zero_votes.append(img_path)
6 comb_images(zero_votes)

```



In all these images, the face is not really visible. Thus, with an additional pipeline step between face detection and face recognition which calculates the quality of the face, we might get rid of this example.

```

1 def conf_matrix(votes_needed):
2     tp, fn, fp, tn = 0,0,0,0
3     for s in same:
4         if s >= votes_needed:
5             tp += 1
6         else:
7             fn += 1
8     for d in different:
9         if d >= votes_needed:
10            fp += 1
11        else:
12            tn += 1
13    print_confusion(fp, fn, tp, tn)
14    return tp, tn, fp, fn
15 for votes_needed in range(10):
16    res = conf_matrix(votes_needed)

```

```

17 print("(votes_needed={};_h_mean={:.3f})".format(votes_needed,
18         harmonic_mean_from_conf(*res)))
19 print()

```

```

1  +-----+
2  |           Predicted condition           |
3  +-----+-----+-----+-----+
4  |           |           | Positive | Negative |
5  +-----+-----+-----+-----+
6  | Actual   | Positive | 166.00 | 0.00 |
7  | Condition| Negative | 609.00 | 0.00 |
8  +-----+-----+-----+-----+

```

```

9  (votes_needed=0; h_mean=0.353)

```

```

10
11 +-----+
12 |           Predicted condition           |
13 +-----+-----+-----+-----+
14 |           |           | Positive | Negative |
15 +-----+-----+-----+-----+
16 | Actual   | Positive | 144.00 | 22.00 |
17 | Condition| Negative | 22.00  | 587.00 |
18 +-----+-----+-----+-----+

```

```

19 (votes_needed=1; h_mean=0.867)

```

```

20
21 +-----+
22 |           Predicted condition           |
23 +-----+-----+-----+-----+
24 |           |           | Positive | Negative |
25 +-----+-----+-----+-----+
26 | Actual   | Positive | 127.00 | 39.00 |
27 | Condition| Negative | 11.00  | 598.00 |
28 +-----+-----+-----+-----+

```

```

29 (votes_needed=2; h_mean=0.836)

```

```

30
31 +-----+
32 |           Predicted condition           |
33 +-----+-----+-----+-----+
34 |           |           | Positive | Negative |
35 +-----+-----+-----+-----+
36 | Actual   | Positive | 111.00 | 55.00 |
37 | Condition| Negative | 5.00   | 604.00 |
38 +-----+-----+-----+-----+

```

```

39 (votes_needed=3; h_mean=0.787)

```

```

40
41 +-----+
42 |           Predicted condition           |
43 +-----+-----+-----+-----+
44 |           |           | Positive | Negative |
45 +-----+-----+-----+-----+
46 | Actual   | Positive | 105.00 | 61.00 |
47 | Condition| Negative | 4.00   | 605.00 |
48 +-----+-----+-----+-----+

```

```

49 (votes_needed=4; h_mean=0.764)

```

```

50
51 +-----+
52 |           Predicted condition           |
53 +-----+-----+-----+-----+
54 |           |           | Positive | Negative |
55 +-----+-----+-----+-----+
56 | Actual   | Positive | 93.00  | 73.00 |

```

```

57 | Condition | Negative | 1.00 | 608.00 |
58 +-----+-----+-----+-----+
59 (votes_needed=5; h_mean=0.715)
60
61 +-----+-----+-----+-----+
62 | Predicted condition |
63 +-----+-----+-----+-----+
64 | | Positive | Negative |
65 +-----+-----+-----+-----+
66 | Actual | Positive | 85.00 | 81.00 |
67 | Condition | Negative | 1.00 | 608.00 |
68 +-----+-----+-----+-----+
69 (votes_needed=6; h_mean=0.675)
70
71 +-----+-----+-----+-----+
72 | Predicted condition |
73 +-----+-----+-----+-----+
74 | | Positive | Negative |
75 +-----+-----+-----+-----+
76 | Actual | Positive | 75.00 | 91.00 |
77 | Condition | Negative | 1.00 | 608.00 |
78 +-----+-----+-----+-----+
79 (votes_needed=7; h_mean=0.620)
80
81 +-----+-----+-----+-----+
82 | Predicted condition |
83 +-----+-----+-----+-----+
84 | | Positive | Negative |
85 +-----+-----+-----+-----+
86 | Actual | Positive | 68.00 | 98.00 |
87 | Condition | Negative | 0.00 | 609.00 |
88 +-----+-----+-----+-----+
89 (votes_needed=8; h_mean=0.581)
90
91 +-----+-----+-----+-----+
92 | Predicted condition |
93 +-----+-----+-----+-----+
94 | | Positive | Negative |
95 +-----+-----+-----+-----+
96 | Actual | Positive | 62.00 | 104.00 |
97 | Condition | Negative | 0.00 | 609.00 |
98 +-----+-----+-----+-----+
99 (votes_needed=9; h_mean=0.544)

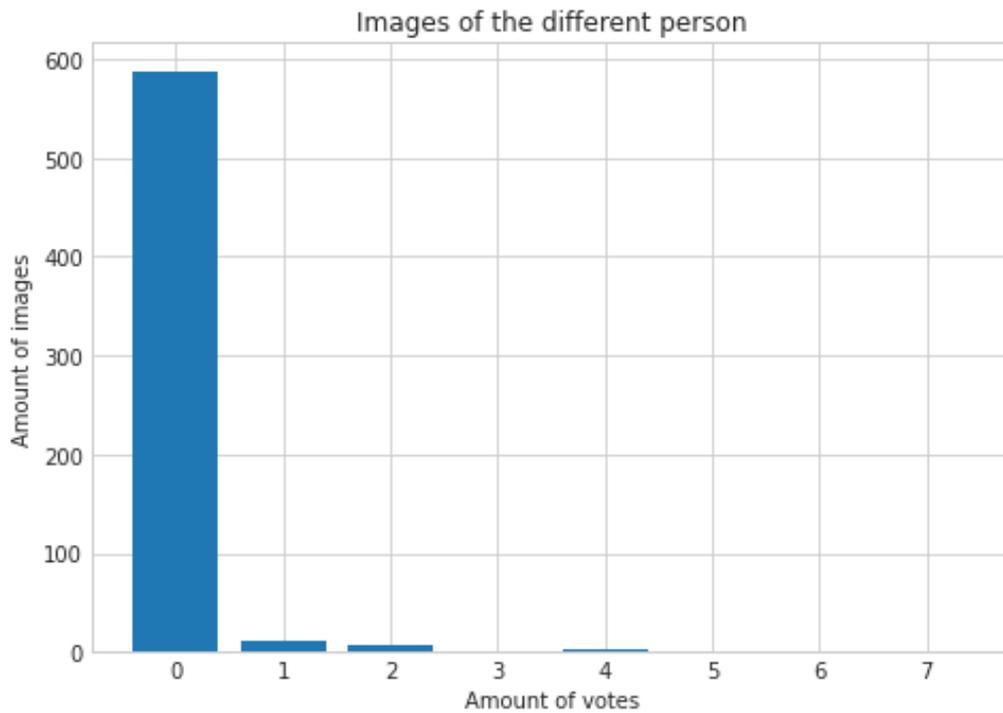
```

Similar to the previous analysis, we can create the same plot for all different people:

```

1 amount_votes = defaultdict(lambda: 0)
2 for d in different:
3     amount_votes[d] += 1
4 data = []
5 for i in range(0, max(amount_votes.keys())+1):
6     data.append(amount_votes[i])
7 fig = plt.figure()
8 ax = fig.add_axes([0,0,1,1])
9 ax.bar(range(len(data)), data)
10 ax.set_xlabel("Amount_of_votes")
11 ax.set_ylabel("Amount_of_images")
12 ax.set_title("Images_of_the_different_person")
13 plt.show()

```



For 96.4% of the images, no single template image has a distance smaller than the chosen threshold. There are 1 images where 7 template images have a distance smaller than the threshold. Thus, if we set `votes_needed` to 8 we would not get a single FP (in our dataset). Let's look at these 1 examples with 7 votes:

```

1 highest_errors = []
2 max_votes = len(data)-1
3 for file in os.listdir(base_path):
4     person_path = os.path.join(base_path, file)
5     if os.path.isdir(person_path) and file != person and not file in exclude_dirs:
6         for img in os.listdir(person_path):
7             img_path = os.path.join(person_path, img)
8             if rec.get_eye_distance(img_path) < eye_distance:
9                 continue
10            if rec.get_eyes_mouth_distance(img_path) < eye_mouth_distance:
11                continue
12            if get_amount_votes(img_path, folder, threshold) == max_votes:
13                highest_errors.append(img_path)

```

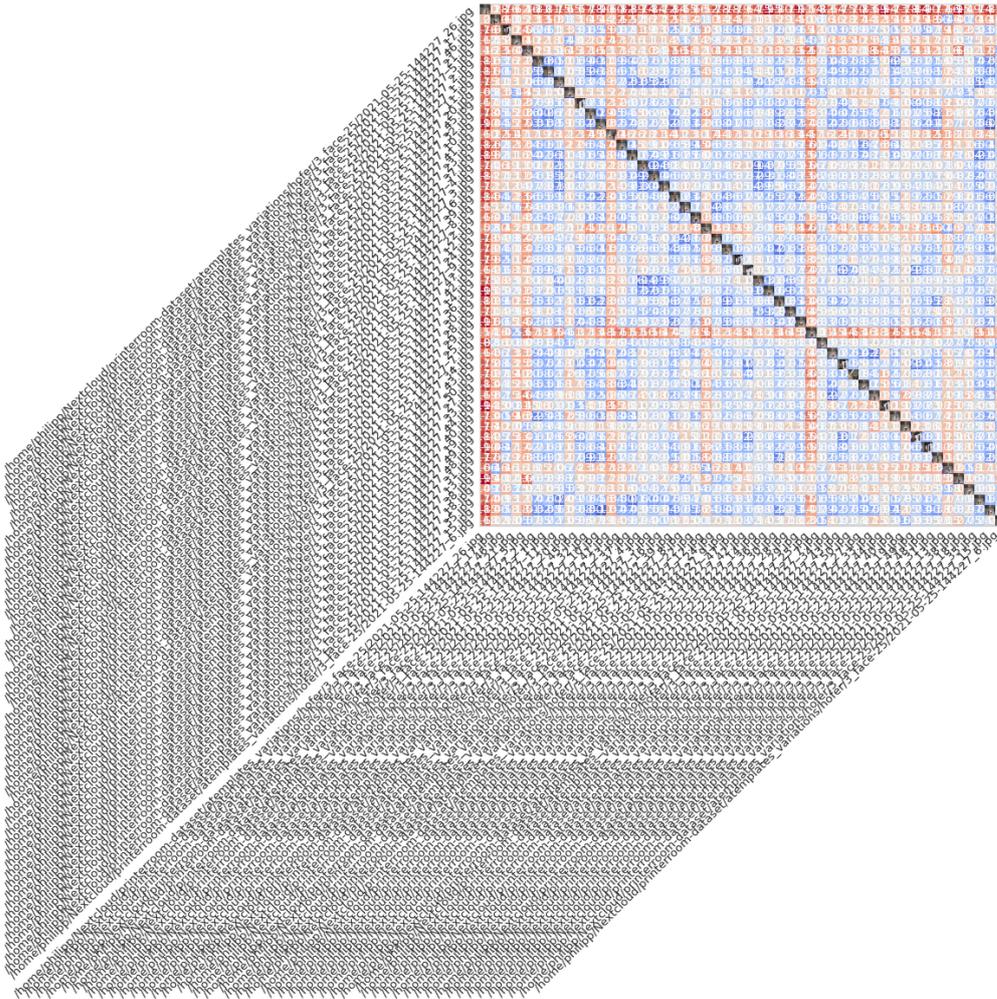


```
1 highest_errors
```

```
1     ['/home/philipp/Nextcloud/printerroom-dataset/sonntag/2021-06-14T15-20-44.png  
1     ']
```

```
1 for error in highest_errors:  
2     get_amount_votes(error, folder, 1.05)
```

```
1 data = []  
2 for file in os.listdir(folder):  
3     file_path = os.path.join(folder, file)  
4     data.append(file_path)  
5 rec.vis.get_confusion_vis(data)
```



5. KMeans

```
1 from sklearn.cluster import KMeans
2 import numpy as np
```

```
1 kmeans = KMeans(n_clusters=1, random_state=42).fit(get_embs_from_path(folder))
```

```
1 template = kmeans.cluster_centers_
```

```
1 h_mean_onecluster, same_onecluster, different_onecluster, threshold_onecluster =
  harmonic_mean(template)
2 print_harmonic_mean(h_mean_onecluster, threshold_onecluster)
```

```
1 Harmonic mean: 0.919 (with a threshold of 0.947)
```

```
1 kmeans = KMeans(n_clusters=5, random_state=42).fit(get_embs_from_path(folder))
```

```
1 emb_combined = np.mean( np.array(kmeans.cluster_centers_), axis=0 )
2 h_mean_fiveclustersavg, same_fiveclustersavg, different_fiveclustersavg,
  threshold_fiveclustersavg = harmonic_mean(emb_combined)
3 print_harmonic_mean(h_mean_fiveclustersavg, threshold_fiveclustersavg)
```

```
1 Harmonic mean: 0.932 (with a threshold of 0.93)
```

```
1 for amount_clusters in range(1,20+1):
2     print("{}_cluster:".format(amount_clusters, " if amount_clusters == 1 else "
3         kmeans = KMeans(n_clusters=amount_clusters, random_state=42).fit(
4             get_embs_from_path(folder))
5             emb_combined = np.mean( np.array(kmeans.cluster_centers_), axis=0 )
6             h_mean_fiveclustersavg, same_fiveclustersavg, different_fiveclustersavg,
7             threshold_fiveclustersavg = harmonic_mean(emb_combined)
8             print_harmonic_mean(h_mean_fiveclustersavg, threshold_fiveclustersavg)
9             print("")
```

```
1 1 cluster:
2 Harmonic mean: 0.919 (with a threshold of 0.947)
3
4 2 clusters:
5 Harmonic mean: 0.919 (with a threshold of 0.944)
6
7 3 clusters:
8 Harmonic mean: 0.932 (with a threshold of 0.958)
9
10 4 clusters:
11 Harmonic mean: 0.925 (with a threshold of 0.943)
12
13 5 clusters:
14 Harmonic mean: 0.932 (with a threshold of 0.93)
15
16 6 clusters:
17 Harmonic mean: 0.928 (with a threshold of 0.945)
18
19 7 clusters:
20 Harmonic mean: 0.929 (with a threshold of 0.932)
21
22 8 clusters:
23 Harmonic mean: 0.929 (with a threshold of 0.91)
24
25 9 clusters:
26 Harmonic mean: 0.920 (with a threshold of 0.918)
27
28 10 clusters:
29 Harmonic mean: 0.927 (with a threshold of 0.906)
30
31 11 clusters:
32 Harmonic mean: 0.932 (with a threshold of 0.907)
33
34 12 clusters:
35 Harmonic mean: 0.929 (with a threshold of 0.909)
36
37 13 clusters:
38 Harmonic mean: 0.932 (with a threshold of 0.918)
39
40 14 clusters:
41 Harmonic mean: 0.927 (with a threshold of 0.921)
42
43 15 clusters:
44 Harmonic mean: 0.936 (with a threshold of 0.923)
45
```

```

46     16 clusters:
47     Harmonic mean: 0.931 (with a threshold of 0.937)
48
49     17 clusters:
50     Harmonic mean: 0.931 (with a threshold of 0.941)
51
52     18 clusters:
53     Harmonic mean: 0.935 (with a threshold of 0.938)
54
55     19 clusters:
56     Harmonic mean: 0.934 (with a threshold of 0.946)
57
58     20 clusters:
59     Harmonic mean: 0.929 (with a threshold of 0.942)

```

```

1 kmeans = KMeans(n_clusters=4, random_state=42).fit(get_embs_from_path(folder))

```



```

1 import sys
2 from PIL import Image
3 def comb_images_face(path):
4     images = []
5     if isinstance(path, list):
6         for p in path:
7             images.append(Image.open(p))
8     else:
9         for file in os.listdir(path):
10            images.append(Image.open(os.path.join(path, file)))
11     faces = []
12     for image in images:
13         faces.append(rec.extractor.get_faces(np.array(image))[0])
14
15     images = faces
16
17     widths, heights = zip(*(i.size for i in images))
18     total_width = sum(widths)
19     max_height = max(heights)
20     new_im = Image.new('RGB', (total_width, max_height))
21     x_offset = 0
22     for im in images:
23         new_im.paste(im, (x_offset,0))
24         x_offset += im.size[0]
25     return new_im

```

5.1 Summary

91.9% for average embedding, 85.8% for voting; 93.6% for k_means (with 15 clusters)