

CBR tagging of emotions from facial expressions

Jose L. Jorro-Aragoneses, Belén Díaz-Agudo, and Juan A. Recio-García

Department of Software Engineering and Artificial Intelligence
Universidad Complutense de Madrid, Spain
email: jljorro@ucm.es, belend@ucm.es, jareciog@fdi.ucm.es

Abstract. In our current research we use CBR to identify the emotional state of a user during her interaction with a recommender system by analysing pictures of her momentary facial expression. In a previous work [2] we introduced PhotoMood, a CBR system that uses gestures to identify emotions from the user face self-pictures, and presented preliminary experiments analysing only the external mouth contour of the faces. In this paper we extend the experiments to analyse other gestures in the faces, like eyes and eyebrows. We demonstrate that different people express emotions with different gestures and apply a GA to compute the specific set of weights that maximize the precision results for each user. We experiment with different configurations of case bases and set of weights and demonstrate that the system obtains good precision results even in the cold start situation using only a general case base made of pictures of anonymous people.

1 Introduction

Facial expression analysis is probably the best non-intrusive way to estimate the emotional state [7]. However the development of an automated system that interprets facial expressions is rather difficult as explained in [23], where authors identify three common problems: the face detection, the extraction of the facial feature information and the expression classification.

There are many applications that can use the recognition of emotions. In [19], we implemented a module, PhotoMood, to estimate the emotional state of the user regarding a recommendation about leisure activities and events in Madrid. We use the emotional state to infer in a non-intrusive way if the user likes the activity that the application has recommended. As our recommender system has been implemented in mobile devices, when it can show a recommended activity to the user it takes a photography using the front camera and detects his emotion. This emotion serves as the system's feedback to improve future recommendations. Although the PhotoMood CBR system has been integrated into our recommender application, there are many other scenarios where it can be exploited. For example, a music recommender can use the emotional context to show a type of song depending of the user's mood.

In our previous work we showed that the CBR approach was a suitable solution to estimate the user's mood. We evaluated a classification method that

compares tagged pictures to a new picture of the user. It is a binary classification system to detect two basic emotions: I like and I dislike. Here, the similarity metric was based on the external contour of the mouth. Although we showed in [19] that this approach was valid we also found several limitations that should be addressed. There are people that don't show their emotions only using the mouth, and it could be useful to include additional gestures such as eyebrows or cheeks into the similarity function. But this extension of the similarity metric implies additional considerations about the proper weighting of each gesture to obtain an optimal classification of a new picture. If we compute a weighting for each user individually it would ensure optimal results. However, this strategy suffers the cold-start problem as it requires a certain number of previously tagged self-pictures. An alternative solution is to compute an average weighting for all the users and use it when there are no previous tagged pictures of a concrete user. Leaving aside the weighting used to compare pictures, the case base used to estimate the mood should have a significant impact in the performance. Again, an individual case base containing exclusively tagged pictures of the user should obtain better results than a generic case base with pictures from other users. Here, the cold-start problem appears again and both approaches should be compared.

In this paper we make experimental tests to prove the following hypothesis:

- H_1 The inclusion of additional gestures into the similarity metric of PhotoMood would improve the performance of the system.
- H_2 The inclusion of additional gestures into the similarity metric implies the search of their optimal weighting to maximize the performance.
- H_3 The optimal weighting can be computed globally or individually. The global weighting could be a satisfactory solution to the cold-start problem although the individual approach should obtain better performance.
- H_4 The use of specific case base for each user should obtain better results than a global case base including pictures of other users. But, again, the global case base could be a satisfactory solution to the cold-start problem.

To test these hypotheses this paper presents: first, an extension of the similarity metric that includes additional gestures; second, the use of a genetic algorithm to compute the optimal weighting; and third, an experimental evaluation to validate the proposed approach.

The paper runs as follow. Section 2 reviews the related work in the field of tagging emotions from pictures. In Section 3 we describe the extension of PhotoMood CBR module. In section 4 we explain how we obtain an optimal configuration of the weights that maximizes the performance of the system. Next, in Section 5 we explain the experiments that we have performed and their results including a discussion of the limitations found. Finally, Section 6 presents the conclusions and some ideas for future work.

2 Related work

Despite its difficulty there are many approaches to face detection, which could be generally classified into the following categories: Template matching methods look for certain patterns in pictures [8]. Feature-based methods use features such as colour [26, 15] or shapes [30] to identify expressions and the face contour. Geometry-based techniques use sizes and relative positions of the components of face [24, 3]. Knowledge-based methods try to encode human knowledge of what constitutes a typical face [16, 31]. Appearance based techniques use linear transformation and statistical methods to determine the vectors that represent the face [6]. There are also several machine learning methods that use training samples [25, 27]. And finally, hybrid approaches that combine some of the previous techniques [17].

We can also find many face detection algorithms in the literature that can be used: Principal Component Analysis [29], Independent Component Analysis [12], Kernel methods [18], Support Vector Machines [11], or Hidden Markov Models [21] are some examples of the approaches being applied.

Regarding face detection applied to tagging of emotions we can find several proposals in the literature [9, 22]. For example, Maglogiannis et al [20] identify the emotions using edge detection and measuring the gradient of eye's or mouth's regions. They consider five major emotions: neutral, happy, sad, surprised and angry. Alternatively, Chawan et al [3] use the nearest matching pattern after applying the bezier curve on eyes and mouth.

These algorithmic techniques for emotion's recognition provide promising approaches in stationary settings [1, 4, 23]. However, mobile settings introduce important complications in capturing facial expressions. Recently some novel solutions have been proposed. For example, in Teeters et al [28] authors present a solution based on a chest mounted self camera detecting 24 feature points and a dynamic bayesian model; Gruebler et al [10] use an interface device to detect facial bioelectrical signals. Not intrusive techniques based on smartphones have been also explored. For example, mobile applications that captures pictures of one's facial expressions throughout the day [14]. There are many other mobile applications using emotion recognition, mostly for personal skills reinforcement or autism disorder therapy.

We used the emotion's recognition to detect the feedback of a recommendation. To do it we created PhotoMood [19]. PhotoMood is a CBR system to detect emotions with facial recognition. In [19], we used the external contour of the mouth to the similarity function of the CBR system. We observed that the system need more gestures to detect the emotion correctly. In this paper, we add more gestures in the similarity and investigate the importance of each gesture in each user.

Next we present our proposal for face emotion recognition based on the CBR paradigm.

3 PhotoMood: a CBR approach to face emotion recognition

PhotoMood is a module that uses the facial expression to detect the emotions. The system uses pictures taken by the front camera of a mobile device. For each image, PhotoMood identifies the gestures that serve as the query of a CBR system, managing a case base of previously tagged pictures. After retrieving the k nearest neighbours a voting schema returns an estimation of the mood associated to the query picture.

Therefore we can distinguish 2 different stages in PhotoMood:

1. *Image pre-processing*: Infers the gestures of the user in the picture.
2. *CBR process*: Using the gestures as the query, the most similar pictures in the case base are used to estimate the user's mood.

Next, we explain both stages of the PhotoMood module.

3.1 Photography processing

PhotoMood's picture preprocessing consists mainly in a face and gesture detection stages. We use the Luxand SDK to detect gestures of the user because it offers a high quality according to comparative studies [5]. Concretely, our algorithm follows the steps shown in Figure 1:

1. Face Detection from the picture through Luxand SDK.
2. Gesture detection: Luxand SDK provides up to 66 facial feature points (eyes, eyebrows, mouth, nose...)
3. Selection of the gestures used by our similarity function.

The gesture selection obtains 46 points, organized in 8 vectors:

- \bar{v}_1 : External outline of mouth.
- \bar{v}_2 : Internal outline of mouth.
- \bar{v}_3 : Outline of right cheek.
- \bar{v}_4 : Outline of left cheek.
- \bar{v}_5 : Outline of right eye.
- \bar{v}_6 : Outline of left eye.
- \bar{v}_7 : Outline of right eyebrow.
- \bar{v}_8 : Outline of left eyebrow.

It is important to note that our previous work in [19] only takes into account \bar{v}_1 . Although Luxand SDK provides up to 66 we decided to ignore the points from the nose and the chin as our preliminary experiments shown that they are not relevant for detecting emotions.

These eight vectors comprise the query of the CBR system that is implemented as explained next.

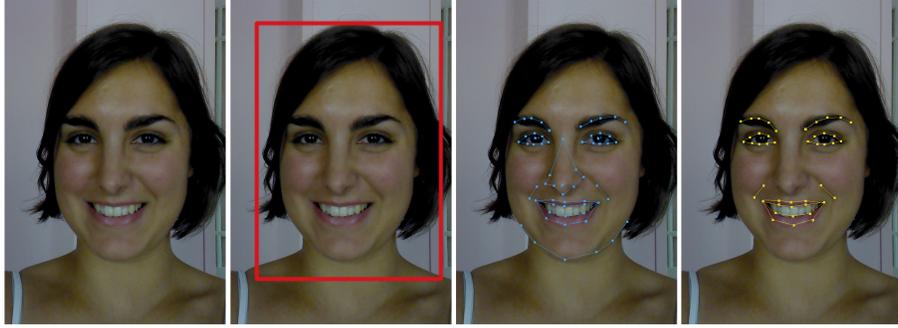


Fig. 1. PhotoMood steps (left to right): 1) original picture; 2) Face detection; 3) Gesture Detection (66 points); 4) Query points (46 points).

3.2 The PhotoMood CBR process

Each query (\mathcal{Q}) is composed of 8 vectors (\overline{v}_i) representing the different gestures of the face. Each vector includes a list of points that are the coordinates of the gesture in the image:

$$\mathcal{Q} = \langle \overline{v}_1^q, \dots, \overline{v}_8^q \rangle \quad (1)$$

Case base uses the same representation. Each case contains a description with the corresponding solutions, concretely, the emotion tags t associated to each picture.

$$\mathcal{C} = \langle \mathcal{D}_c, \mathcal{S}_c \rangle \quad (2)$$

where

$$\mathcal{D}_c = \langle \overline{v}_1^c, \dots, \overline{v}_8^c \rangle \quad (3)$$

$$\mathcal{S}_c = \{t_1, \dots, t_m\} \quad (4)$$

We assume that several emotion tags could be used to classify a picture. However, up to now, our system uses a binary classification (like and dislike), so $m = 2$.

Having this representation of the cases, next we explain the four standard stages of the CBR system:

Retrieval. The retrieval stage returns an ordered list with the k cases that are the most similar to the query. We have implemented a k -Nearest Neighbour algorithm which returns the k elements most similar to the query. The similarity value is computed by weighting the local similarity of each vector \overline{v}_i . Therefore each pair of vectors $\langle \overline{v}_i^q, \overline{v}_i^c \rangle$ is compared and the resulting value

is weighted with a value w_i that represents the relevance of the corresponding gesture in the global similarity computation:

$$Sim(\mathcal{Q}, \mathcal{D}_c) = \sum_{n=1}^8 w_i * Sim_i(\overline{v}_i^q, \overline{v}_i^c) \quad (5)$$

where

$$\sum_{i=1}^8 w_i = 1 \quad (6)$$

Having each vector composed of several points $\overline{v}_i = \langle p_1, \dots, p_z \rangle$ we obtain the angle between each pair of points and the horizontal axis. Following the contour of a gesture, the arctangent between a pair of points is computed to obtain their angle. Each angle of the query \mathcal{Q} is compared to the corresponding angle of the case \mathcal{D}_c producing z (angle-level) similarity values:

$$Sim_i(\overline{v}_i^q, \overline{v}_i^c) = \frac{1}{z} \sum_{j=1}^z 1 - \sqrt{(\arctan(\widehat{p_j^q p_l^q}) - \arctan(\widehat{p_j^c p_l^c}))^2} \quad (7)$$

where $l = (j + 1) \bmod z$

Note that the z limit could be different for each vector \overline{v}_i .

Reuse. To obtain a solution for the query, PhotoMood applies a weighted voting schema according to the similarity of the retrieved cases. Given the scoring function:

$$score(t_i) = \sum sim(\mathcal{Q}, \mathcal{D}_c) \forall c | \mathcal{S}_c = t_i \quad (8)$$

The solution assigned to the query is:

$$t_i = arg\ max\{score(t_i), i = 1, \dots, m\} \quad (9)$$

Revise. The revision of the proposed solution is directly performed by the user. Thus, it is external to the PhotoMood CBR module. The user should correct the system when the emotion detected is wrong.

Retain. The retain policy of PhotoMood only stores those cases that were revised by the user. That is, solutions incorrectly computed by the CBR process.

The values used to weight on the retrieval stage are key components. Next, we explain the algorithm that obtains the optimal weights.

4 Obtaining optimal weights

In this paper, we use a Genetic Algorithm (GA) to calculate the weight that maximizes the performance of the system right answer rate. This is a satisfactory

solution to learn the optimal weights [2, 13]. This algorithm uses a population of individuals representing different weightings. This population evolves until the algorithm obtains the individual (i.e. the weights) that returns the best performance. Additionally, each individual of the GA not only contains vectors' weights, it includes the k parameter of the K-Nearest Neighbour algorithm to estimate which are the best number of cases that must be retrieved to perform the classification of a new picture. The GA runs as follows:

First, it generates an initial population of 500 individuals. Each individual contains the weights of each vector and the k that the CBR system uses in the retrieval stage. When we create the individuals, we assigned a random value for each weight, and then they are normalized to sum 1. On the other hand, we assigned a random k value between 1, 3, 5 or 7.

Then the system repeats the following cycle until there is no improvement in the performance achieved by the best individual in the population:

1. *Evaluation stage:* The genetic algorithm executes a cross-validation of our CBR system configured with the weights and k value contained by every individual of the population. The resulting performance is the fitness of the individual.
2. *Remove stage:* After every individual of the population has been evaluated, the 25% of the population with the worst fitness is removed.
3. *Combination stage:* To replace the 25% of the population that has been removed, the algorithm combines the individuals the best individuals. These individuals are taken in pairs (parents) and they are combined to obtain a new individual (child). The child individual will contain for each weight the average of the parents' weight (properly normalized). The k value is computed analogously.
4. *Mutation stage:* Once the new population has been generated there is a mutation stage to avoid local maximal values. The algorithm proceeds to modify the 5% of the population. It chooses random individuals and applies a mutation function. This function modifies the individual's weights randomly.

For the cross-validation in the evaluation stage we use the leave-one-out function. Here the evaluation of the CBR system can be performed in two different ways:

1. Using the pictures of all the users as queries. This configuration of the GA obtains the optimal global weighting that should be applied when there are no previous tagged pictures of a new user. In this case, the system uses a general case base (*GCB*) that contains 300 images of anonymous people that we obtain with search processes in Google Images
2. Using just the pictures of a concrete user as queries. This configuration of the GA obtains the optimal personal weighting for that user. These weights should be used when there are already pictures of the user in the system. In this scenario, the case base only contains the pictures for each user. We will refer to this case base as the user case base (*PCB*).

Next, we explain the different experiments that we have performed and the results obtained.

5 Experimental evaluation

This section describes the experimental evaluation performed in order to validate the CBR process being proposed. We explain the experiment design, its results and some proposals to enhance the obtained results.

5.1 Experimental setup

We have designed the test to demonstrate the hypothesis that we have explained in section 1. We have used 2 different types of case bases to do it. Next, we explain each type:

- *General Case Base (GCB)*: It consists of 300 images of anonymous people that we obtained with search processes in Google Images.
- *Personal Case Base (PCB)*: Each user has a personal case base. It (This Case Bases) contains the images of each user. Each case base has a different number of images. This is because we want to see the performance of the CBR system. From $User_1$ to $User_{10}$, the case bases are compound with 30 images per user. The rest of the users have between 13 and 21 images per user.

All case bases have 50% of images that are tagged with like and the other 50% with a dislike tag.

Now we explain the tests that we have applied:

1. *Test 1*: This test only takes into account \bar{v}_1 , and ignores the other gestures from \bar{v}_2 to \bar{v}_8 . It uses the general case base *GCB*. This way, we reproduce exactly the experimental setup in [19]. The associated results serve as the baseline to compare the improvements presented in this paper.
2. *Test 2*: Evaluate the benefits of including additional gestures. This test does not weight the gesture vectors. It assumes that all of them have the same relevance ($w_i = \frac{1}{8}$) using *GCB*.
3. *Test 3*: Evaluates the performance of the CBR system using the optimal global weighting (gw) obtained by the genetic algorithm and the general case base *GCB*.
4. *Test 4*: We evaluate the impact of using an optimal personal weighting pw for each user (again obtained by the GA) and the general case base *GCB*.
5. *Test 5*: Finally, we evaluate the use of a specific case base for each user *PCB*. We use a optimal personal weighting (pw) for each user with his personal case base. The weight configuration is obtained by the GA.

Next, we explain the results that we have got for each test.

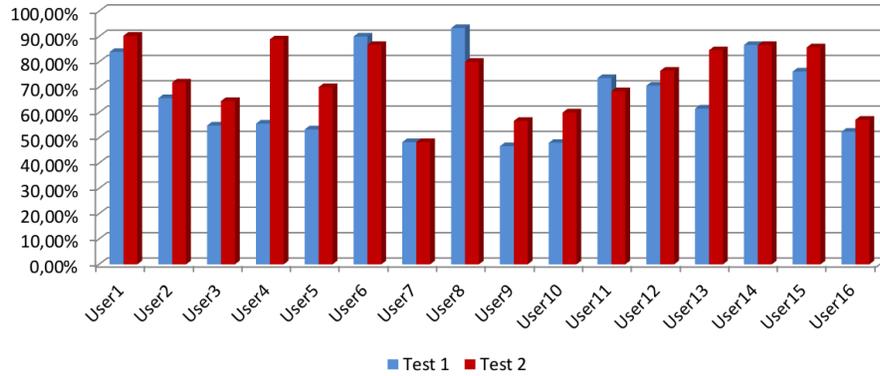


Fig. 2. Compare the use of more gestures than \bar{v}_1 .

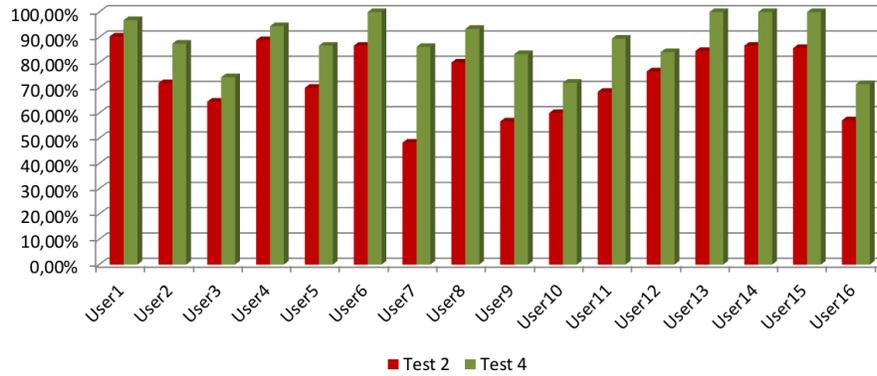


Fig. 3. Compare the use of an optimal weighting for each user.

5.2 Results

Figure 2 shows a comparison of the precision obtained by *Test 1* and *Test 2*. In this Figure we can observe that when we include additional gestures to the similarity metric (not only \bar{v}_1) the precision in most users improves. This demonstrates our first hypothesis H_1 that was statistically confirmed by a Wilcoxon test ($p - value < 0.05$). Also, we can observe that the Test 1's average ($\mu = 66.28\%$ with $\sigma = 0.16$) is worse than Test 2 ($\mu = 73.52\%$ with $\sigma = 0.13$).

Also, Figure 2 shows there are users' precisions that are worse than the baseline (\bar{v}_1 only). The reason is that \bar{v}_1 is very important when comparing these users, so the inclusion of additional gestures in *Test 2* adds noise to the similarity estimation. These users need a specific weighting configuration to indicate that \bar{v}_1 is more important than other gestures in their similarities.

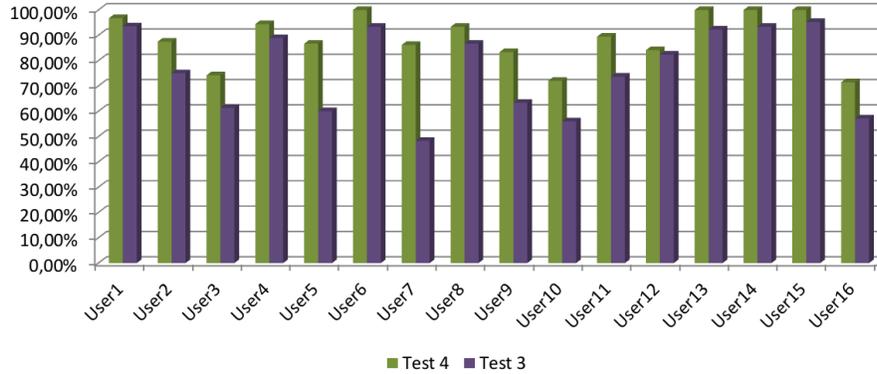


Fig. 4. Compare the use of a global configuration for all users in relation to an individual configuration for each user.

To demonstrate the correctness of H_2 we have evaluated the impact of optimal weights for each vector. Results are shown in the Figure 3. As we can observe, the use of an optimal configuration weighting for each user improves precision results for all users. In addition, the average also improves ($\mu = 89.34\%$ with $\sigma = 0.09$). Again, this result was statistically confirmed by using the Wilcoxon test ($p - value < 0.05$). Consequently, we can assert that each user expresses his emotions in a different way, and the system should be aware of the importance of each gesture to adapt the similarity metric.

Figure 4 shows our hypothesis H_3 that using a global weighting for all the users is a good approximation for most users. Our principal problems are $User_5$ and $User_7$. They are not expressive at all and when we apply a global configuration they get worse results. They need an individual configuration because they have similar gestures to express *like* and *dislike*, and the configuration that we apply for other users are not suitable for these users. In spite of these users, the global configuration is a satisfactory solution to the cold-start problem. When the system cannot be trained with previous user's pictures it can use the global configuration as an acceptable solution. Again this result was statistically confirmed by the Wilcoxon test ($p - value < 0.05$). The average of Test 3 confirms that it is a satisfactory solution to the cold-start problem ($\mu = 76.27\%$ with $\sigma = 0.16$).

To obtain the optimal weighting for each user we use a GA as it is explained in the section 4. Figure 4 shows that GA is a good tool to find the optimal configuration of each weight. On the other hand, GA is an algorithm that has a high computational cost because it takes a lot of time to find the optimal configuration for each user. For these reason, we have looked for an alternative approach to obtain an acceptable performance without this high computational costs. We have used a global configuration for all users.

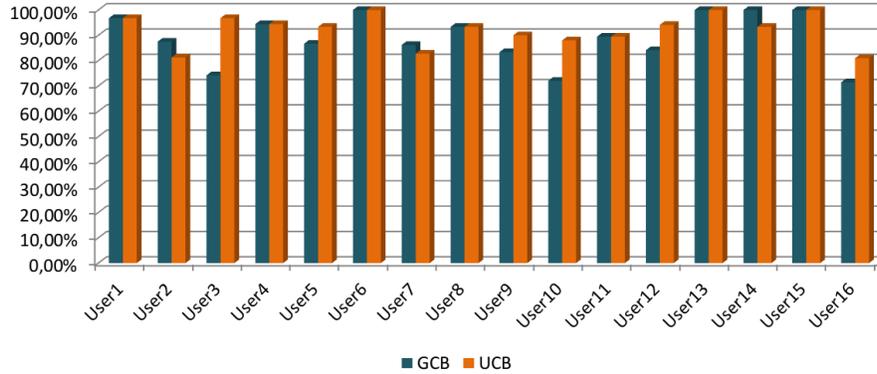


Fig. 5. Compare the uses *GCB* in relation to uses *PCB* for each user. Both use an individual configuration weightings.

Finally, Figure 5 compares the precision when we use *GCB* for all users and the user case base *PCB*. When the system uses *PCB*, we get better results for most users. We can see it with their averages ($\mu_{GCB} = 89.34\%$ with $\sigma_{GCB} = 0.09$ and $\mu_{UCB} = 92.16\%$ with $\sigma_{UCB} = 0.06$). The statistical significance of the results (Wilcoxon Test, $p - value < 0.07$) confirms our hypothesis H_4 .

However, there are some users whose precision gets worse, like *User₂*, *User₇* and *User₁₄*. Next we discuss the causes of these problems. *User₂* has a problem of emotion tagging. She tags images wrongly, and the system confuses these images with an incorrect emotion. She needs to correct her images to obtain the correct precision. The problem of *User₇* is the same that we explained before. She is not expressive at all, and her images are similar when she expresses *like* and *dislike*. To improve her precision, she should insert more pictures in her case base. Finally, *User₁₄* gets worse in his precision because his personal case base is very small. The system doesn't have enough cases to detect the user's emotion correctly. This is an example of cold-start problem and the using of *GCB* with a global configuration weighting is a good approximation.

Except these users, the using of the personal case base with a individual configuration weighting for each user is the best solution, as we proposed in H_4 .

Next, we explain the conclusions we have obtained from these results and the future work for PhotoMood.

6 Conclusions and future work

In this paper we have described an extension of PhotoMood system. PhotoMood is a CBR system that compares the gestures of the face with other that, were previously tagged. Before, we used the top of the external contour of the mouth to detect the user's emotion, but now we have added new face elements.

In the first part of the paper we explained four hypotheses:

- H_1 The inclusion of additional gestures improves the performance of the system.
- H_2 An optimal weighting of each gesture maximizes the performance.
- H_3 The global weighting is a satisfactory solution to the cold-start problem.
- H_4 The specific case base for each user gets better results than a global case base. The global case base is a satisfactory solution to the cold-start problem.

In the Section 5.2, we demonstrated these hypotheses with different tests. We obtained the best results using a personal case base (PCB) and a personal weighting configuration (pw). In addition, we demonstrated that using a global case base (GCB) and a global weighting configuration (gw) is a satisfactory solution to the cold-start problem for most users.

By the results that we have obtained, the performance of each configuration on the system is ordered as:

$$prf(GCB + gw) < prf(GCB + pw) < prf(UCB + pw) \quad (10)$$

With these results, we can watch that precision rate is very high. The reason for this is that the users have a forced expressions in their pictures. As future work, we could apply the PhotoMood extension in an environment where the system makes the photography without warning the user, and then, we could get different gestures for a same emotion.

Another problem is the images quality. For detecting the gestures we need good images. Bad images make that image processing doesn't work correctly. The typical problems are: the incorrect angle of the photograph, a low resolution, lighting conditions, difference on image proportions. For future work, we could apply algorithms to improve the images quality.

In this paper we consider that the user change his face expression when he like or dislike the recommendation. For a future work, it should determine the relation between the user feedback and his face expression.

Finally, a problem that we have found is the computing cost of the genetic algorithm. When the algorithm looks for the optimal configuration of each user, it slows because repeats the process many times. The algorithm can take a long time before it returns the best solution. In a future work, it may be possible looking for a solution so we can get the optimal configuration for each user in less time. A suitable solution could be divide the case base in clusters and the genetic algorithm would need less comparisons to detect an optimal configuration.

References

1. Anderson, S., Conway, M.: Investigating the structure of autobiographical memories. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 19(5), 1178–1196 (1993)
2. Benou, P., Bitos, V.: Context-aware query processing in ad-hoc environments of peers. *JECO* 6(1), 88–62 (2008)
3. Chawan, P.M., Jadhav, M.M.C., Mashruwala, J.B., Nehete, A.K., Panjari, P.A.: Real time emotion recognition through facial expressions for desktop devices. *International Journal of Emerging Science and Engineering (IJESE)* 1(7), 104–108 (2013)

4. Cohen, I., Sebe, N., Garg, A., Chen, L.S., Huang, T.S.: Facial expression recognition from video sequences: temporal and static modeling. *Computer Vision and Image Understanding* 91(12), 160 – 187 (2003), special Issue on Face Recognition
5. Degtyarev, N., Seredin, O.: Comparative testing of face detection algorithms. In: *Proceedings of the 4th International Conference on Image and Signal Processing*. pp. 200–209. ICISP'10, Springer-Verlag, Berlin, Heidelberg (2010)
6. Draper, B.A., Baek, K., Bartlett, M.S., Beveridge, J.: Recognizing faces with PCA and ICA. *Computer Vision and Image Understanding* 91(12), 115 – 137 (2003), special Issue on Face Recognition
7. Eckman, P.: Facial expression and emotion. *American Psychologist* 48(4), 384–392 (1993)
8. Feris, R.S., de Campos, T.E., Junior, R.M.C.: Detection and tracking of facial features in video sequences. In: Cairó, O., Sucar, L.E., Cantu, F.J. (eds.) *MICAI. Lecture Notes in Computer Science*, vol. 1793, pp. 127–135. Springer (2000)
9. Fragopanagos, N., Taylor, J.: Emotion recognition in humancomputer interaction. *Neural Networks* 18(4), 389 – 405 (2005), emotion and Brain
10. Gruebler, A., Suzuki, K.: Analysis of social smile sharing using a wearable device that captures distal electromyographic signals. In: Stoica, A., Zarzhitsky, D., Howells, G., Frowd, C.D., McDonald-Maier, K.D., Erdogan, A.T., Arslan, T. (eds.) *EST*. pp. 178–181. IEEE Computer Society (2012)
11. Guo, G., Li, S.Z., Chan, K.L.: Support vector machines for face recognition. *Image and Vision Computing* 19(910), 631 – 638 (2001)
12. Hyvärinen, A., Oja, E.: Independent component analysis: Algorithms and applications. *Neural Netw.* 13(4-5), 411–430 (May 2000)
13. Jarmulak, J., Craw, S., Rowe, R.: Genetic algorithms to optimise cbr retrieval. In: Blanzieri, E., Portinale, L. (eds.) *Advances in Case-Based Reasoning, Lecture Notes in Computer Science*, vol. 1898, pp. 136–147. Springer Berlin Heidelberg (2000)
14. Karapanos, E.: *Modeling Users' Experiences with Interactive Systems, Studies in Computational Intelligence*, vol. 436. Springer (2013)
15. Kovac, J., Peer, P., Solina, F.: Human skin color clustering for face detection. In: *EUROCON 2003. Computer as a Tool. The IEEE Region 8*. vol. 2, pp. 144–148 vol.2. IEEE (2003)
16. Lin, C., Fan, K.C.: Triangle-based approach to the detection of human face. *Pattern Recognition* 34(6), 1271 – 1284 (2001)
17. Lin, H.J., Yen, S.H., Yeh, J.P., Lin, M.J.: Face detection based on skin color segmentation and svm classification. In: *SSIRI*. pp. 230–231. IEEE Computer Society (2008)
18. Liu, Q., Huang, R., Lu, H., Ma, S.: Face recognition using kernel-based fisher discriminant analysis. In: *Automatic Face and Gesture Recognition, 2002. Proceedings. Fifth IEEE International Conference on*. pp. 197–201 (May 2002)
19. López-de-Arenosa, P., Díaz-Agudo, B., Recio-García, J.A.: CBR tagging of emotions from facial expressions. In: *ICCB* (2014)
20. Maglogiannis, I., Vouyioukas, D., Aggelopoulos, C.: Face detection and recognition of natural human emotion using markov random fields. *Personal and Ubiquitous Computing* 13(1), 95–101 (2009)
21. Nefian, A.V., Hayess, M.H.: Hidden Markov Models For Face Recognition. In: *Proc. International Conf. on Acoustics, Speech and Signal Processing (ICASSP'98)*. vol. 5, pp. 2721–2724 (May 1998)

22. Pantic, M., Rothkrantz, L.J.: Facial action recognition for facial expression analysis from static face images. *Trans. Sys. Man Cyber. Part B* 34(3), 1449–1461 (Jun 2004)
23. Pantic, M., Rothkrantz, L.J.M.: Automatic analysis of facial expressions: The state of the art. *IEEE Trans. Pattern Anal. Mach. Intell.* 22(12), 1424–1445 (Dec 2000)
24. Paul, S.K., Uddin, M.S., Bouakaz, S.: Extraction of facial feature points using cumulative histogram. *CoRR* abs/1203.3270 (2012)
25. Rowley, H.A., Baluja, S., Kanade, T.: Neural network-based face detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 20(1), 23–38 (Jan 1998)
26. Solina, F., Peer, P., Batagelj, B., Juvan, S., Kovac, J.: Color-based face detection in the "15 seconds of fame" art installation. In: in *Proceedings of Mirage 2003*, (INRIA Rocquencourt. pp. 38–47 (2003)
27. Sung, K.K., Poggio, T.: Example-based learning for view-based human face detection. *IEEE Trans. Pattern Anal. Mach. Intell.* 20(1), 39–51 (Jan 1998)
28. Teeters, A., El Kaliouby, R., Picard, R.: Self-cam: Feedback from what would be your social partner. In: *ACM SIGGRAPH 2006 Research Posters*. SIGGRAPH '06, ACM, New York, NY, USA (2006)
29. Turk, M., Pentland, A.: Eigenfaces for recognition. *J. Cognitive Neuroscience* 3(1), 71–86 (Jan 1991)
30. Wang, J., Tan, T.: A new face detection method based on shape information. *Pattern Recognition Letters* 21(67), 463 – 471 (2000)
31. Yang, G., Huang, T.S.: Human face detection in a complex background. *Pattern Recognition* 27(1), 53 – 63 (1994)