

A Personalized Affective Memory Model for Improving Emotion Recognition

Pablo Barros¹ German I. Parisi^{1,2} Stefan Wermter¹

Abstract

Recent models of emotion recognition strongly rely on supervised deep learning solutions for the distinction of general emotion expressions. However, they are not reliable when recognizing on-line and personalized facial expressions, e.g., for person-specific affective understanding. In this paper, we present a neural model based on a conditional adversarial autoencoder to learn how to represent and edit general emotion expressions. We then propose Grow-When-Required networks as personalized affective memories to learn individualized aspects of emotion expressions. Our model achieves state-of-the-art performance on emotion recognition when evaluated on *in-the-wild* datasets. Furthermore, our experiments include ablation studies and neural visualizations in order to explain the behavior of our model.

1. Introduction

Automatic facial expression recognition became a popular topic in the past years due to the success of deep learning techniques. What was once the realm of specialists on adapting human-made descriptors of facial representations (Zhao et al., 2003), became one of many visual recognition tasks for deep learning enthusiasts (Schmidhuber, 2015). The result is a continuous improvement in the performance of automatic emotion recognition during the last decade (Sariyanidi et al., 2015) reflecting the progress of both hardware and different deep learning solutions.

However, the performance of deep learning models for facial expression recognition has recently stagnated (Soleymani et al., 2017; Mehta et al., 2018). The most common cause is the high dependence of such solutions on balanced, strongly labeled and diverse data. Recent approaches propose to address these problems by introducing transfer learning

techniques (Ng et al., 2015; Kaya et al., 2017), neural activation and data distribution regularization (Ding et al., 2017b; Pons & Masip, 2018), and unsupervised learning of facial representations (Zen et al., 2014; Kahou et al., 2016). Most of these models present an improvement of performance when evaluated on specific datasets. Still, most of these solutions require strongly supervised training which may bias their capacity to generalize emotional representations. Recent neural models based on adversarial learning (Kim et al., 2017; Saha et al., 2018) overcome the necessity of strongly labeled data. Such models rely on their ability to learn facial representations in an unsupervised manner and present competitive performance on emotion recognition when compared to strongly supervised models.

Furthermore, it is very hard to adapt end-to-end deep learning-based models to different emotion recognition scenarios, in particular, real-world applications, mostly due to very costly training processes. As soon as these models need to learn a novel emotion representation, they must be re-trained or re-designed. Such problems occur due to understanding facial expression recognition as a typical computer vision task instead of adapting the solutions to specific characteristics of facial expressions (Adolphs, 2002).

We address the problem of learning adaptable emotional representations by focusing on improving emotion expression recognition based on two facial perception characteristics: the diversity of human emotional expressions and the personalization of affective understanding. A person can express happiness by smiling, laughing and/or by moving the eyes depending on who he or she is talking to, for example (Mayer et al., 1990). The diversity of emotion expressions becomes even more complex when we take into consideration inter-personal characteristics such as cultural background, intrinsic mood, personality, and even genetics (Russell, 2017). Besides diversity, the adaptation to learn personalized expressions is also very important. When humans already know a specific person they can adapt their own perception to how that specific person expresses emotions (Hamann & Canli, 2004). The interplay between recognizing an individual's facial characteristics, i.e. how they move certain muscles in their face, mouth and eye positions, blushing intensity, and clustering them into emotional states is the key for modeling a human-like emotion recognition system (Sprengelmeyer et al., 1998).

¹Knowledge Technology, Department of Informatics, University of Hamburg, Germany ²Apprentice Inc, Mountain View, CA, USA. Correspondence to: Pablo Barros <barros@informatik.uni-hamburg.de>.

The problem of learning diversity on facial expressions was addressed with the recent development of what is known as *in-the-wild* emotion expression datasets (Zadeh et al., 2016; Mollahosseini et al., 2017; Dhall et al., 2018). These corpora make use of a large amount of data to increase the variability of emotion representations. Although deep learning models trained with these corpora improved the performance on different emotion recognition tasks (de Bittencourt Zavan et al., 2017; Kollias & Zafeiriou, 2018), they still suffer from the lack of adaptability to personalized expressions.

Unsupervised clustering and dynamic adaptation were explored for solving the personalization of emotion expression problems (Chen et al., 2014; Zen et al., 2014; Valenza et al., 2014). Most of the recent solutions use a collection of different emotion stimuli (EEG, visual and auditory, for example) and neural architectures, but the principle is the same: to train a recognition model with data from one individual person to improve affective representation. Although performing well in specific benchmark evaluations, these models are not able to adapt to online and continuous learning scenarios found on real-world applications.

The affective memory model (Barros & Wermter, 2017) addresses the problem of continuous adaptation by proposing a Growing-When-Required (GWR) network to learn emotional representation clusters from one specific person. One affective memory is created per person and it is populated with representations from the last layer of a convolutional neural network (CNN). This model creates prototypical neurons that represent each of the perceived facial expressions. Although improving emotion expression recognition when compared to state-of-the-art models, this model is not reliable on online scenarios. Prototype neurons will be formed only after a certain expression is perceived, and it relies heavily on a supervised CNN, so the model will only produce good performance once all the possible emotion expressions were perceived.

In this paper, we propose a personalized affective memory framework which improves the original affective memory model with respect to online emotion recognition. At the beginning of an interaction, humans tend to rely heavily on their own prior knowledge to understand facial expressions, with such prior knowledge learned over multiple episodes of interaction (Nook et al., 2015). We propose the use of a novel generative adversarial autoencoder model to learn the prior knowledge of emotional representations. We adapt the model to transfer the learned representations to unknown persons by generating edited faces with controllable emotion expressions (Huang et al., 2018; Wang et al., 2018). We use the generated faces of a single person to initialize a personalized affective memory. We then update the affective memory with the perceived faces on an online manner and use it to recognize facial expressions online.

We evaluate our model in two steps: first, how well the model can learn a general representation of emotion expressions, using the AffectNet (Mollahosseini et al., 2017) dataset. And second, how the novel personalized affective memory performs when recognizing emotion expressions of specific persons using the continuous expressions on the OMG-Emotion dataset (Barros et al., 2018). We compare our model with state-of-the-art solutions on both datasets. Besides performance alone, we also provide an exploratory study on how our model works based on its neural activities using activation visualization techniques.

2. The Personalized Affective Memory

The Personalized Affective Memory (P-AffMem) model is composed of two modules: the prior-knowledge learning module (PK) and the affective memory module. The prior-knowledge learning module consists of an adversarial autoencoder that learns facial expression representations (z) and generates images using controllable emotional information, represented as continuous arousal and valence values (y).

After the autoencoder is trained, we use it to generate a series of edited images with different expressions from a person using different values of y . We use this collection of generated faces as initialization to a growing-when-required (GWR) network which represents the personalized affective memory for that specific person. We use the clusters of the personalized affective memory to recognize the emotion expressions of that person. Figure 1 illustrates the topological details of the P-AffMem model.

2.1. Prior-Knowledge Adversarial Autoencoder

Our PK model is an extended version of the Conditional Adversarial Autoencoder (CAAE) (Zhang et al., 2017), which was developed to learn how to represent and edit age information into facial images. We choose the CAAE as the basis for our prior-knowledge module due to its capability of working without paired-data (images of the same person with different ages) and to its robustness against statistical variations in the input data distribution.

The PK model consists of an encoder and generator architecture (E and G), and three discriminator networks. The first one learns arousal/valence representations (D_{em}), the second guarantees that the facial expression representation has a uniform distribution (D_{prior}) and the third one assures that the generated image is photo-realistic and expresses the desired emotion (D_{real}). The model receives as input an image (x) and a continuous arousal/valence label (y). It produces a facial representation (z) and an edited image expressing the chosen arousal/valence (x_{gen}).




Figure 1. The P-AffMem model is composed of the Prior-Knowledge Adversarial autoencoder (PK) and the affective memories. The PK implements an encoder/decoder architecture and specific discriminators to detail arousal and valence (D_{em}), to ensure a prior distribution to the encoded representation (D_{prior}) and to guarantee a photo-realistic image generation (D_{real}). The affective memories are individual Growing-When-Required(GWR) networks which learn personalized aspects of the facial expression.

2.1.1. ENCODER AND GENERATOR

The encoder network (E) consists of four convolution layers and one fully connected output layer. We make use of a convolution stride of 2 to avoid the use of pooling, and thus learning its own spatial down-sampling (Radford et al., 2015). E receives as input an RGB image with a dimension of $96 \times 96 \times 3$ and outputs the facial expression representation (z) with a dimension of (50×1) . z is concatenated with the desired emotional representation (y) and used as input to the decoder network (G) which generates an image (x_{gen}) with the desired emotional expression y . The decoder is composed of six transposed convolution layers. The first four have a stride of 2, while the last two have a stride of 1. All the convolution and transposed convolution of E and G use a kernel size of 5×5 . All the layers of E and G use ReLu activation functions. Both E and G are trained with an image reconstruction loss (L_{rec}):

$$\min_{E,G} L_{rec} = L_1(x, G(E(x), z)) \quad (1)$$

where L_1 is the mean absolute error.

One of the problems of the CAAE is the appearance of ghost artifacts on the reconstructed images. The CAAE authors address this problem by using a total variation minimization loss function (Mahendran & Vedaldi, 2015), but with questionable results. We address this problem by using the identity-preserving loss (L_{iden}) used by the ExprGan (Ding et al., 2017a) and the G2-GAN (Song et al., 2017) on the reconstructed image. The identity preserving loss is computed using a pre-trained VGG face model (Parkhi et al., 2015). The VGG face model proposes a topological update on the VGG16 model and is trained to identify persons based on their faces. To compute L_{iden} , we compare the activations of the five first convolutional layers of the VGG Face for x and x_{gen} :

$$\min_{E,G} L_{iden} = \sum_l L_1(\phi_l(x), \phi_l(G(E(x), z))) \quad (2)$$

where ϕ_l is the activation of the l th layer of the VGG face.

2.1.2. AROUSAL/VALENCE REPRESENTATION DISCRIMINATOR

The arousal/valence representation discriminator (D_{em}) enforces the encoder to learn facial representations which are able to be used on an emotion recognition task. This discriminator is important to guarantee that z is suitable to represent emotion expressions as z will be used to populate the affective memory. D_{em} consists of two fully connected hidden layers, the first one implementing a ReLu activation function and the last one a \tanh linear function. The last layer is used as output to represent continuous values of arousal and valence. The loss is calculated as the

$$L_{em} = \text{MSE}(D_{em}(x)^a, x^a) + \text{MSE}(D_{em}(x)^v, x^v) \quad (3)$$

where MSE is the Mean-Squared error, $D_{em}(x)^a$ and $D_{em}(x)^v$ represent the arousal and valence output of D_{em} respectively, and x^a and x^v the associated arousal and valence value to the image input x .

2.1.3. UNIFORM DISTRIBUTION DISCRIMINATOR

The uniform distribution discriminator (D_{prior}) is imposed on z and enforces it to be uniformly distributed. It is important that z is uniformly distributed to increase the generalization of the model when representing a facial expression. It receives as input z itself or a randomly chosen sample from the uniform distribution $p_{prior}(z)$ and it has as goal to distinguish between them. D_{prior} is composed of four fully connected layers, each one implementing a ReLu ac-

tivation function. The adversarial loss function between E and D_{prior} is defined as:

$$\min_E \max_{D_z} L_{prior} = \mathbb{E}_{z_{prior} \sim p_{prior}(z)} [\log D_z(z_{prior})] + \mathbb{E}_{x \sim p_{data}(x)} [\log(1 - D_z(E(x)))] \quad (4)$$

where E is the likelihood, $p_{prior}(z)$ the prior distribution imposed on the internal representation z and $p_{data}(x)$ the distribution of the training images.

2.1.4. PHOTO-REALISTIC IMAGE DISCRIMINATOR

The photo-realistic image discriminator (D_{real}) is basically a normal real/fake discriminator on the typical Generative Adversarial Network (GAN). In our case it also serves as a mechanism to impose the decoder (G) to produce photo-realistic images with the desired emotion expression (y). D_{real} implements four convolution layers, with stride of 2 and kernel size of 5x5, followed by two fully connected layers. To enforce the desired arousal and valence (y) into the generated image (x_{gen}), we perform a re-shape and zero-padding operation to y and concatenate it with each convolutional layer of D_{real} . The concatenation of y to all the convolutional layers proved to be an important step in our experiments towards generating a photo-realistic edited facial expression. To enforce the desired y into x_{gen} produced by G , we use the following adversarial loss function:

$$\min_G \max_{D_{img}} L_{img} = \mathbb{E}_{x, y \sim p_{data}(x, y)} [\log D_{img}(x, y)] + \mathbb{E}_{x, y \sim p_{data}(x, y)} [\log(1 - D_{img}(G(E(x), y), y))] \quad (5)$$

where $p_{data}(x, y)$ is the distribution of the training data.

2.1.5. OVERALL LOSS FUNCTION

To train our model we use an overall loss-function defined as:

$$\min_{E, G} \max_{D_z, D_{img}} L_{total} = \lambda_1 L_{rec} + \lambda_2 L_{em} + \lambda_3 L_{iden} + \lambda_4 L_z + \lambda_5 L_{img} \quad (6)$$

in which the coefficients λ_1 , λ_2 , λ_3 , λ_4 and λ_5 are used to balance the facial expression discrimination, the high-fidelity of the generated images and the presence of the desired emotion on the generated image.

2.2. Affective Memory

Growing-When-Required (GWR) networks (Marsland et al., 2002) were deployed recently to address the problems of

continuous learning (Parisi et al., 2017; 2018). The capability of creating prototype neurons using an online manner allows the GWR to adapt quickly to changes on the input data. Which makes it ideal for our online learning mechanism.

Each neuron of the GWR consists of a weight vector w_j representing prototypical information of the input data. A newly perceived emotion expression will be associated with a best-matching unit (BMU) b , which is calculated by minimizing the distances between the facial expression and all the neurons on the GWR. Given a set of N neurons, b with respect to the input $x \in \mathbb{R}^n$ is computed as:

$$b = \arg \min_{j \in N} (\|x - w_j\|^2). \quad (7)$$

New connections are created between the BMU and the second-BMU with relation to input. When a BMU is computed, all the neurons the BMU is connected to are referred to as its topological neighbors. Each neuron is equipped with a habituation counter $h_i \in [0, 1]$ expressing how frequently it has been fired based on a simplified model of how the efficacy of a habituating synapse reduces over time.

The habituation rule is given by $\Delta h_i = \tau_i \cdot \kappa \cdot (1 - h_i) - \tau_i$, where κ and τ_i are constants that control the decreasing behavior of the habituation counter (Marsland et al., 2002). To establish whether a neuron is habituated, its habituation counter h_i must be smaller than a given habituation threshold t_h .

The network is initialized with two neurons and, at each learning iteration, it inserts a new neuron whenever the activity of the network ($a(i) = \exp(-(\|x - w_j\|^2))$), when an expression i is perceived, of a habituated neuron is smaller than a given threshold t_a , i.e., a new neuron is created if $a(i) < t_a$ and $h_i < t_h$.

The P-AffMem model uses the GWR as affective memories to perform personalization learning of prototype neurons. The GWR is created when the first facial expression of a person is perceived. The PK module generates 200 edited images with combinations of arousal and valence within the interval of $[-1, 1]$, with increments of 0.01. These images are used to initialize the affective memory and act as a transfer knowledge from the PK initial estimations. To avoid that the generated samples dominate the affective memory update over time, we stop generating samples after a certain number of expressions were perceived.

To guarantee that the PK does not dominate the training of the affective memory over time, we use a novel update function to modulate the impact of the PK on the GWR. The training of the network is carried out by adapting the BMU according to:

$$\Delta \mathbf{w}_j = \epsilon_i \cdot h_i \cdot (\mathbf{x} - \mathbf{w}_j) \cdot (1 - a(i)), \quad (8)$$

where ϵ_i is a constant learning rate. If real faces are perceived, and they are different from the ones encoded by the PK, the activation of the network will be smaller and the impact on the weight update will be higher. The affective memory will be encouraged to create new neurons to represent newly perceived expressions instead of the ones coming from the PK.

To allow the GWR to perform classification of emotion expressions, we implement an associative labeling (Parisi et al., 2017) to each neuron. During the training phase, we assign to each neuron two continuous values, representing arousal and valence. When training with an example that comes from the generated images, we update the labels of the BMU using the desired arousal and valence. The update is modulated by a labeling learning factor γ which is defined during training. To categorize a newly perceived expression, we just read the labels of the BMU associated with it.

2.3. Parameter Optimization

We trained the PK model using L_{total} and optimized the model’s topology, coefficients, the batch size and the training parameters using a TPE optimization strategy (Bergstra et al., 2013). The optimization was based on two characteristics: the objective performance on the arousal/valence discriminator and the minimization of L_{total} . The final coefficient values are as follows: $\lambda_1 = 1$, $\lambda_2 = 0.02$, $\lambda_3 = 0.3$, $\lambda_4 = 0.01$ and $\lambda_5 = 0.01$. The batch size is 48. The normal distribution with mean 0 and standard deviation 0.02 is employed for the initialization of the weights of all layers. All biases are initially set to 0. For optimization, the Adam Optimizer with learning rate $\alpha = 0.0002$, $\beta_1 = 0.5$ and $\beta_2 = 0.999$ is employed.

Table 1. Training parameters of each affective memory.

Parameter	Value
Epochs	10
Activity threshold (t_a)	0.4
Habituation threshold (t_h)	0.2
Habituation modulation (τ_i and κ)	0.087, 0.032
Labeling factor (γ)	0.4

We also use TPE optimization to tune the parameters of the affective memory. Although they were created individually for one specific person, all of them have the same hyperparameters to simplify our evaluation task and allow the model to be used online. Given that the GWR adapts to the input data distribution, and we maintain the same data nature, we do not believe that fine-tuning each GWR to each individual person would obtain a large gain on recognition performance. Table 1 displays the final affective memory

parameters.

3. Experimental Setup

To evaluate and understand better the individual impact of the PK and the affective memory on the recognition of emotion expressions, we perform two different types of experiments. First, we run a series of ablation studies to assess the contribution of each mechanism on the PK for general emotion recognition. Second, we run an emotion recognition experiment with the entire framework to assess the impact of the personalization on the emotion recognition performance.

3.1. Datasets

The AffectNet dataset (Mollahosseini et al., 2017) is composed of more than 1 million images with facial expressions collected from different *in-the-wild* sources, such as Youtube and Google Images. More than 400 thousands images were manually annotated with continuous arousal and valence. The dataset has its own separation between training, testing and validation samples. The labels for the testing samples are not available, thus, all our experiments are performed using the training and validation samples. The large labelled data distribution of the AffectNet is important to guarantee that the PK learns general emotion recognition. It provides an ideal corpus to assess the impact of each proposed mechanism in the PK using an objectively comparable measure.

Furthermore, the AffectNet does not discriminate between images from the same person, so personalization does not play an important role. To evaluate the contributions of personalization, we provide final experiments on the One-Minute Gradual-Emotional Behavior dataset (OMG-Emotion) (Barros et al., 2018). The OMG-Emotion is composed of Youtube videos which are about one minute long and are annotated taking into consideration a continuous emotional behavior. The videos were selected using a crawler technique that uses specific keywords based on long-term emotional scenes such as “monologues”, “auditions”, “dialogues” and “emotional scenes”, which guarantees that each video has only one person performing an emotional display. A total of 675 videos were collected, which sums up to around 10h of data. Each utterance on the videos is annotated with two continuous labels, representing arousal and valence. The emotion expressions displayed in the OMG-Emotion dataset are heavily impacted by person-specific characteristics which are highlighted by the gradual change of emotional behavior over the entire video.

3.2. Experiments

Our experiments are divided into two categories: the ablation studies (A) and the personalized emotion recogni-

tion (P). The A category is divided into a series of experiments, each of them evaluating one of the discriminators on the PK: the photo-realistic image discriminator (D_{real}), the uniform distribution (D_{prior}) and the arousal/valence discriminator (D_{em}). We train the PK in each of these experiments with the training subset of the AffectNet corpus and evaluate it using the validation subset.

We first train the PK without any of the previously mentioned discriminators (PK_{base}) to guarantee an unbiased baseline. Then, we repeat the training, now adding each discriminator individually. We report experimental results by combining the presence of each discriminator. Finally, we add all the discriminators (PK_{all}) and train the PK again.

To provide a standard evaluation metric, we use the encoder representation (z) as input to an emotion recognition classifier. The classifier is composed of the same topology as the arousal/valence discriminator, and it is post-trained using the same training subset of the AffectNet dataset. When training the PK with the arousal/valence discriminator, we do not use the emotion recognition classifier, so that we can assess the performance of this specific discriminator.

The P category is divided into two experiments: first, the PK is pre-trained with the AffectNet dataset and used to evaluate the test set of the OMG-Emotion dataset. Then we use the entire P-AffMem framework with the affective memories and repeat the experiment, now using one affective memory for each video of the test set. During our optimization routine we found that if the PK generates faces for more than 1s of the videos, the affective memory did improve the final results. As the OMG-Emotion dataset has a framerate of 25 frames per second, we turn off the PK image generation after the first 25 frames.

To evaluate the dimensional arousal and valence recognition, we use the Concordance Correlation Coefficient (CCC) (Lawrence & Lin, 1989) between the outputs of the model and the true labels. The CCC can be computed as:

$$CCC = \frac{2\rho\sigma_x\sigma_y}{\sigma_x^2 + \sigma_y^2 + (\mu_x - \mu_y)^2} \quad (9)$$

where ρ is the Pearson’s Correlation Coefficient between model prediction labels and the annotations, μ_x and μ_y denote the mean for model predictions and the annotations and σ_x^2 and σ_y^2 are the corresponding variances.

4. Results

4.1. Ablation Studies (A)

Our ablation studies, summarized in Table 2, help us to understand and quantify the impact of each of the PK modules on recognizing emotion expressions. The baseline model,

without any extra discriminator, can be evaluated as a simple Generative Adversarial Network (GAN). Without surprise, the baseline model did not perform well on emotion recognition, achieving the worst result in our experiments. This reflects the inability of the baseline model to provide a good facial expression discrimination capability. Evaluating the model by introducing each discriminator allows us to assess their impact. The emotion recognition discriminator (D_{em}) seems to be the one which impacts most on the PK performance, which not surprising, as it enforces the PK encoder to learn specific characteristics for facial expression recognition. Nevertheless, we also observe a strong impact on the photo-realistic discriminator (D_{real}), indicating that the network benefits greatly from enforcing the encoding and generation of characteristics that discriminate between individuals. The prior distribution discriminator (D_{prior}) has a smaller individual contribution than all the others. However, when combined with the other discriminators, it provides a great improvement on the recognition.

Table 2. Concordance Correlation Coefficient (CCC) for arousal and valence when evaluating the different discriminator of the PK on the validation subset of the AffectNet dataset.

Model	Arousal	Valence
PK_{base}	0.03	0.05
$PK_{base} + D_{prior}$	0.03	0.08
$PK_{base} + D_{real}$	0.08	0.09
$PK_{base} + D_{em}$	0.18	0.22
$PK_{base} + D_{prior} + D_{real}$	0.11	0.18
$PK_{base} + D_{prior} + D_{em}$	0.25	0.35
$PK_{base} + D_{real} + D_{em}$	0.29	0.45
PK_{all}	0.38	0.67
(Mollahosseini et al., 2017)	0.34	0.60

Training the PK with all the discriminators yield the best results. To the best of our knowledge, the only reported CCC for arousal and valence on the AffectNet corpus comes from the authors of the dataset themselves (Mollahosseini et al., 2017). They use an AlexNet convolutional neural network (Krizhevsky et al., 2012) re-trained to recognize arousal and valence. Our PK model presents a better general performance improving the CCC in more than 0.04 for arousal and 0.07 for valence. As our ablation studies are only intended to shed light on the impact of our PK modules, we did not pursue a pure benchmark study with other datasets.

4.2. Personalized Recognition (P)

we perform the experiments with and without the affective memory in order to quantify its contribution to the P-AffMem model. Table 3 summarizes the achieved CCC and the current state-of-the-art results on the OMG-Emotion dataset. The presence of the affective memories greatly

improve the performance of the model with an increase on the achieved CCC of 0.9 and 0.7 for arousal and valence, respectively, when compared to the PK with all discriminators.

The P-AffMem achieved the best results up to this point on the OMG-Emotion dataset. This dataset was recently used as part of a challenge and different solutions, mostly based on pre-trained deep learning models, were presented. Our model achieved a CCC improvement of 0.07 on arousal and 0.04 in valence, when compared to the winner of the challenge (Zheng et al., 2018), which made use of audio/visual processing. The same occurred with the second best model (Peng et al., 2018). The best model which used only facial expressions (Deng et al., 2018) achieved an arousal and valence CCC which were 0.16 and 0.18 smaller than our model. Further results are available at the OMG-Emotion website ¹.

Table 3. Concordance Correlation Coefficient (CCC) for arousal and valence when evaluating our model on the testing subset of the OMG-Emotion dataset.

Model	Arousal	Valence
PK_{base}	-0.06	-0.10
$PK_{base} + D_{prior}$	0.02	0.01
$PK_{base} + D_{real}$	0.04	0.02
$PK_{base} + D_{em}$	0.09	0.12
$PK_{base} + D_{prior} + D_{real}$	0.13	0.13
$PK_{base} + D_{prior} + D_{em}$	0.21	0.29
$PK_{base} + D_{real} + D_{em}$	0.27	0.36
PK_{all}	0.32	0.46
$P - AffMem$	0.43	0.53
(Zheng et al., 2018)	0.36	0.49
(Peng et al., 2018)	0.24	0.43
(Deng et al., 2018)	0.27	0.35

5. Discussions

The PK module must be pre-trained to allow a better emotion expression recognition generalization. With 16 million parameters, the PK model has an extensive and computational expensive training procedure, as common to any adversarial network training. Our hypothesis was that, once trained, the PK would provide a general emotion recognition capability. Our experiments with the OMG-Emotion dataset demonstrate that the PK achieved a performance similar to other pre-trained deep learning models on recognizing general emotion expressions.

The affective memory contributes to the P-AffMem by providing a self-adaptable mechanism that improves the performance of the PK by clustering the individual characteristics of how a person expresses emotions. This assumption was

¹<https://bit.ly/2XUJBjg>

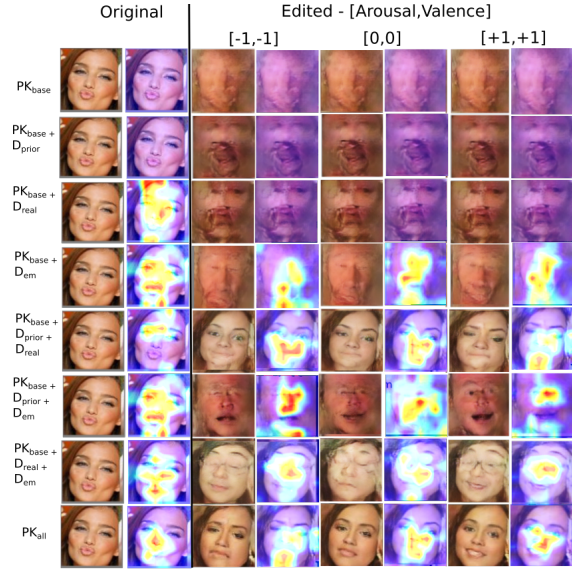


Figure 2. The first pair of the column represents the perceived emotion expression and the neural activation map of the last layer of the encoder. The remaining pairs of columns display the edited face using the indicated arousal and valence value, and its respective neural activation map.

demonstrated objectively as the P-AffMem has a higher CCC than the PK alone.

To extend our understanding of the P-AffMem and to pinpoint our claims on general versus specific emotion recognition, we will discuss in the next sessions the individual contributions of both the PK and the affective memories. We also present an insight into how the models contribute to the performance increase.

5.1. The Impact of the Discriminator Mechanisms

As an adversarial autoencoder, the PK is trained to both encode and edit a facial image. Each discriminator imposes a specific characteristic to both the encoding and editing tasks, and together they showed to be optimal, in our setup, for emotion expression recognition. While the impact of the discriminators on encoding a facial expression was objectively measured by our experiments, its contributions to face editing were just implicitly evaluated. By demonstrating that the affective memory improved the performance of the model, we guarantee that the edited faces used to initialize and train the affective memory actually have the desired facial expressions.

As the model is optimized to recognize facial expressions, an analysis of the pixel-wise quality of the edited faces can be misleading. To illustrate this assumption, we fed one image to each of the PK combinations proposed in our A category experiments and produced three edited faces

with the extreme values of arousal and valence $[-1,-1]$ and $[1,1]$, and a neutral value $[0,0]$. We then calculated the mean neural activation maps (Zhou et al., 2016) for the last convolution layers of the discriminator for each of these examples. Figure 2 exhibits four pairs of columns, where the first is composed of the original image fed to the PKs and the obtained encoder activation visualizations. The other tree columns exhibit the edited images with the indicated arousal and valence values and the encoder activation. As the training of the encoder is directly affected by the discriminators, the visualizations give us an indication on how each discriminator impacts the learned representations, and help us to explain the network objective performance.

Clearly, the encoder trained only with PK_{base} does not learn any meaningful information, which is reflected in the edited images and activation. This is also backed up by our objective results. The combination of the discriminators impact the edited faces and encoder in a distinguishable manner. While PK_{real} enforces a realistic characteristic on the edited faces, it does impact the encoder on focusing facial structures which are not reliable for face expression representation. This effect translates on a lower CCC when compared to D_{em} , for example. When D_{em} is present, the encoder clearly filtered facial expression-rich regions, such as eyes and mouth regions. This is clearly perceptible in the $PK_{base} + PK_{prior} + PK_{photo}$ combination, which presents a distinguishable facial editing but lower CCC than $PK_{base} + D_{real} + D_{em}$, for example. The encoder activation of the former indicates that it ignores some facial characteristics, which could explain the lower performance.

Finally, the combination of all discriminators gave the PK the best objective performance, and this reflects on an edited image and activation. The images contain present clear facial features, without the formation of distorting artifacts, maintaining the general facial proportions and characteristics. The activation shows that the encoder clearly focuses on the eyes-to-mouth region most of the times. These visualizations demonstrate the reliability of the PK in both encoding and editing facial expressions, which is very important for the optimal functioning of the affective memory.

5.2. Affective Memory Behavior

The videos on the OMG-Emotion dataset are heavily impacted by personalization characteristics, as they contain one person performing a monologue-like act for more than 1 minute. As the P-AffMem is initialized with the images edited by the PK, it contains, from the beginning, prototype neurons with a wide range of associated arousal and valence labels. As the video goes on, the neurons of the affective memory are much more influenced by newly perceived facial expressions which make the neural activation increase. Important to note that a newly perceived expression does not change the neuron’s labels, only its prototypical infor-

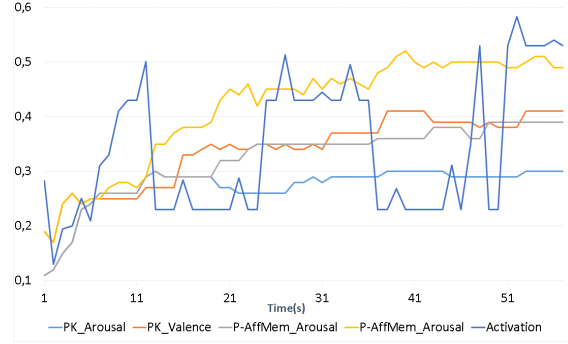


Figure 3. Performance of the model, measured in CCC for both PK and P-AffMem and neural activation when processing the "fd41c38b2" video of the OMG-Emotion dataset.

mation. This way, we guarantee that the network maintains a reliable classification.

As the PK already gives the affective memory general prototype neurons with reliable labels, the neural update makes them much more accurate towards specific characteristics of the newly perceived expressions. After some seconds have passed, the affective memory starts to perform better than the PK. This effect is demonstrated in Figure 3, which illustrates the performance of the PK and the P-AffMem, measured as arousal and valence CCC, and the neural activity evolution of the affective memory when processing the video "fd41c38b2" of the OMG-Emotion dataset.

6. Conclusion and Future Work

The development of our Personalized Affective Memory (P-AffMem) model was inspired by two facial perception characteristics: the understanding of generalized emotional concepts and the online adaptation of individualized aspects of facial expressions. The general prior knowledge (PK) adversarial autoencoder was trained to transfer facial characteristics, represented as arousal and valence, to unknown persons. The affective memories contributed to the learned representations by creating prototypical representations of emotions for a specific person in an online fashion. Our evaluation demonstrated that our model achieved state-of-the-art performance on facial expression recognition, and we presented different insights on how, and why, our model works.

A clear limitation of our model is the processing of instantaneous emotion expressions. To address this problem, recent work on recurrent self-organization and sequence generation would be encouraged (Parisi et al., 2017; 2018). Another direction which could be explored is the integration of multisensory information, but with particular care for asynchronous affective perception, e.g., from prosody in speech and language understanding.

Acknowledgments

The authors gratefully acknowledge partial support from the German Research Foundation DFG under project CML (TRR 169).

References

- Adolphs, R. Recognizing emotion from facial expressions: psychological and neurological mechanisms. *Behavioral and cognitive neuroscience reviews*, 1(1):21–62, 2002.
- Barros, P. and Wermter, S. A self-organizing model for affective memory. In *Neural Networks (IJCNN), 2017 International Joint Conference on*, pp. 31–38. IEEE, 2017.
- Barros, P., Churamani, N., Lakomkin, E., Siqueira, H., Sutherland, A., and Wermter, S. The omg-emotion behavior dataset, Jul 2018.
- Bergstra, J., Yamins, D., and Cox, D. D. Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms. In *Proceedings of the 12th Python in Science Conference*, pp. 13–20. Citeseer, 2013.
- Chen, Y.-A., Wang, J.-C., Yang, Y.-H., and Chen, H. Linear regression-based adaptation of music emotion recognition models for personalization. In *Acoustics, Speech and Signal Processing (ICASSP), 2014 IEEE International Conference on*, pp. 2149–2153. IEEE, 2014.
- de Bittencourt Zavan, F. H., Gasparin, N., Batista, J. C., e Silva, L. P., Albiero, V., Bellon, O. R. P., and Silva, L. Face analysis in the wild. In *Graphics, Patterns and Images Tutorials (SIBGRAPI-T), 2017 30th SIBGRAPI Conference on*, pp. 9–16. IEEE, 2017.
- Deng, D., Zhou, Y., Pi, J., and Shi, B. E. Multimodal utterance-level affect analysis using visual, audio and text features. *arXiv preprint arXiv:1805.00625*, 2018.
- Dhall, A., Kaur, A., Goecke, R., and Gedeon, T. EmotiW 2018: Audio-video, student engagement and group-level affect prediction. In *Proceedings of the 2018 on International Conference on Multimodal Interaction*, pp. 653–656. ACM, 2018.
- Ding, H., Sricharan, K., and Chellappa, R. Exprgan: Facial expression editing with controllable expression intensity. *arXiv preprint arXiv:1709.03842*, 2017a.
- Ding, H., Zhou, S. K., and Chellappa, R. Facenet2expnet: Regularizing a deep face recognition net for expression recognition. In *Automatic Face & Gesture Recognition (FG 2017), 2017 12th IEEE International Conference on*, pp. 118–126. IEEE, 2017b.
- Hamann, S. and Canli, T. Individual differences in emotion processing. *Current opinion in neurobiology*, 14(2):233–238, 2004.
- Huang, B., Chen, W., Wu, X., Lin, C.-L., and Suganthan, P. N. High-quality face image generated with conditional boundary equilibrium generative adversarial networks. *Pattern Recognition Letters*, 111:72–79, 2018.
- Kahou, S. E., Bouthillier, X., Lamblin, P., Gulcehre, C., Michalski, V., Konda, K., Jean, S., Froumenty, P., Dauphin, Y., Boulanger-Lewandowski, N., et al. Emonets: Multimodal deep learning approaches for emotion recognition in video. *Journal on Multimodal User Interfaces*, 10(2):99–111, 2016.
- Kaya, H., Gürpınar, F., and Salah, A. A. Video-based emotion recognition in the wild using deep transfer learning and score fusion. *Image and Vision Computing*, 65:66–75, 2017.
- Kim, Y., Yoo, B., Kwak, Y., Choi, C., and Kim, J. Deep generative-contrastive networks for facial expression recognition. *arXiv preprint arXiv:1703.07140*, 2017.
- Kollias, D. and Zafeiriou, S. Training deep neural networks with different datasets in-the-wild: The emotion recognition paradigm. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8. IEEE, 2018.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Lawrence, I. and Lin, K. A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, pp. 255–268, 1989.
- Mahendran, A. and Vedaldi, A. Understanding deep image representations by inverting them. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 5188–5196, 2015.
- Marsland, S., Shapiro, J., and Nehmzow, U. A self-organising network that grows when required. *Neural Networks*, 15(8–9):1041–1058, 2002.
- Mayer, J. D., DiPaolo, M., and Salovey, P. Perceiving affective content in ambiguous visual stimuli: A component of emotional intelligence. *Journal of personality assessment*, 54(3–4):772–781, 1990.
- Mehta, D., Siddiqui, M. F. H., and Javaid, A. Y. Facial emotion recognition: A survey and real-world user experiences in mixed reality. *Sensors*, 18(2):416, 2018.

- Mollahosseini, A., Hasani, B., and Mahoor, M. H. Affectnet: A database for facial expression, valence, and arousal computing in the wild. *arXiv preprint arXiv:1708.03985*, 2017.
- Ng, H.-W., Nguyen, V. D., Vonikakis, V., and Winkler, S. Deep learning for emotion recognition on small datasets using transfer learning. In *Proceedings of the 2015 ACM on international conference on multimodal interaction*, pp. 443–449. ACM, 2015.
- Nook, E. C., Lindquist, K. A., and Zaki, J. A new look at emotion perception: Concepts speed and shape facial emotion recognition. *Emotion*, 15(5):569, 2015.
- Parisi, G., Tani, J., Weber, C., and Wermter, S. Lifelong learning of spatiotemporal representations with dual-memory recurrent self-organization. In *arXiv:1805.10966*, 2018.
- Parisi, G. I., Tani, J., Weber, C., and Wermter, S. Lifelong learning of humans actions with deep neural network self-organization. *Neural Networks*, 96:137–149, 2017.
- Parkhi, O. M., Vedaldi, A., Zisserman, A., et al. Deep face recognition. In *BMVC*, volume 1, pp. 6, 2015.
- Peng, S., Zhang, L., Ban, Y., Fang, M., and Winkler, S. A deep network for arousal-valence emotion prediction with acoustic-visual cues. *arXiv preprint arXiv:1805.00638*, 2018.
- Pons, G. and Masip, D. Supervised committee of convolutional neural networks in automated facial expression analysis. *IEEE Transactions on Affective Computing*, 9(3):343–350, 2018.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- Russell, J. A. Cross-cultural similarities and differences in affective processing and expression. In *Emotions and Affect in Human Factors and Human-Computer Interaction*, pp. 123–141. Elsevier, 2017.
- Saha, S., Navarathna, R., Helminger, L., and Weber, R. M. Unsupervised deep representations for learning audience facial behaviors. *arXiv preprint arXiv:1805.04136*, 2018.
- Sariyanidi, E., Gunes, H., and Cavallaro, A. Automatic analysis of facial affect: A survey of registration, representation, and recognition. *IEEE transactions on pattern analysis and machine intelligence*, 37(6):1113–1133, 2015.
- Schmidhuber, J. Deep learning in neural networks: An overview. *Neural networks*, 61:85–117, 2015.
- Soleymani, M., Garcia, D., Jou, B., Schuller, B., Chang, S.-F., and Pantic, M. A survey of multimodal sentiment analysis. *Image and Vision Computing*, 65:3–14, 2017.
- Song, L., Lu, Z., He, R., Sun, Z., and Tan, T. Geometry guided adversarial facial expression synthesis. *arXiv preprint arXiv:1712.03474*, 2017.
- Sprengelmeyer, R., Rausch, M., Eysel, U. T., and Przuntek, H. Neural structures associated with recognition of facial expressions of basic emotions. *Proceedings of the Royal Society of London B: Biological Sciences*, 265(1409): 1927–1931, 1998.
- Valenza, G., Citi, L., Lanatà, A., Scilingo, E. P., and Barbieri, R. Revealing real-time emotional responses: a personalized assessment based on heartbeat dynamics. *Scientific reports*, 4:4998, 2014.
- Wang, X., Li, W., Mu, G., Huang, D., and Wang, Y. Facial expression synthesis by u-net conditional generative adversarial networks. In *Proceedings of the 2018 ACM on International Conference on Multimedia Retrieval*, pp. 283–290. ACM, 2018.
- Zadeh, A., Zellers, R., Pincus, E., and Morency, L.-P. Multimodal sentiment intensity analysis in videos: Facial gestures and verbal messages. *IEEE Intelligent Systems*, 31(6):82–88, 2016.
- Zen, G., Sangineto, E., Ricci, E., and Sebe, N. Unsupervised domain adaptation for personalized facial emotion recognition. In *Proceedings of the 16th international conference on multimodal interaction*, pp. 128–135. ACM, 2014.
- Zhang, Z., Song, Y., and Qi, H. Age progression/regression by conditional adversarial autoencoder. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 2, 2017.
- Zhao, W., Chellappa, R., Phillips, P. J., and Rosenfeld, A. Face recognition: A literature survey. *ACM computing surveys (CSUR)*, 35(4):399–458, 2003.
- Zheng, Z., Cao, C., Chen, X., and Xu, G. Multimodal emotion recognition for one-minute-gradual emotion challenge. *arXiv preprint arXiv:1805.01060*, 2018.
- Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., and Torralba, A. Learning deep features for discriminative localization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2921–2929, 2016.