



The Role of Coherence in Facial Expression Recognition

Lisa Graziani, Stefano Melacci and Marco Gori

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

August 21, 2018

The Role of Coherence in Facial Expression Recognition

Lisa Graziani¹, Stefano Melacci², and Marco Gori²

¹ DINFO, University of Florence, Italy

`lisa.graziani@unifi.it`

² DIISM, University of Siena, Italy

`{mela,marco}@diism.unisi.it`

Abstract. Recognizing facial expressions from static images or video sequences is a widely studied but still challenging problem. The recent progresses obtained by deep neural architectures, or by ensembles of heterogeneous models, have shown that integrating multiple input representations leads to state-of-the-art results. In particular, the appearance and the shape of the input face, or the representations of some face parts, are commonly used to boost the quality of the recognizer. This paper investigates the application of Convolutional Neural Networks (CNNs) with the aim of building a versatile recognizer of expressions in static images that can be further applied to video sequences. We first study the importance of different face parts in the recognition task, focussing on appearance and shape-related features. Then we cast the learning problem in the Semi-Supervised setting, exploiting video data, where only a few frames are supervised. The unsupervised portion of the training data is used to enforce two types of coherence, namely temporal coherence and coherence among the predictions on the face parts. Our experimental analysis shows that coherence constraints can improve the quality of the expression recognizer, thus offering a suitable basis to profitably exploit unsupervised video sequences.

Keywords: Facial Expression Recognition, Convolutional Neural Networks, Learning from Constraints, Coherence Constraints

1 Introduction

Facial expression recognition is the problem of detecting emotions in facial images or videos. The research activity on this problem involves the scientific community that is about psychology but also the one that is about computer science and artificial intelligence. Although this task is widely studied and much progress has been made, it still remains a challenging problem, due to the variability and complexity of facial expressions. As a matter of fact, facial expressions can be categorized with respect to multiple classes of emotions. The most widely followed approach consists in considering six basic emotions plus the neutral case, and it is due to the studies of Paul Ekman [2], while other scientists provided more fine grained descriptions [14]. Facial features of expressions are mostly located

around mouth, nose, and eyes, and their locations are essential in explaining and categorizing expressions [1]. Despite the large number of advanced psychological experiments about the human perception and recognition of emotions, we can trivially figure out that different face parts have a different impact in the way humans recognize emotions: the role of eyebrows when we are angry or the way we treat our mouth when we are happy or surprised, for example.

We can find several approaches that exploit Machine Learning with the aim of learning to categorize emotions from examples. Most of them are about using still images [10, 13], while several more recent works also consider video sequences where actors start with a neutral expression and generate a non-neutral one [9, 16]. The learning framework is usually fully supervised, and supervision is either about each training image or about each video sequence. Works that exploit video data focus on the importance of the temporal evolution of the input face. The system proposed by Fan and Tjahjadi [3] processes four sub-regions of the face: forehead, eyes/eyebrows, nose and mouth. They used an extension of the spatial pyramid histogram of gradients and dense optical flow to extract spatial and dynamic features from video sequences, and adopted a multi-class SVM-based classifier with one-to-one strategy to recognise facial expressions. Jung et al. [7] propose a neural-network-based method where two different networks are exploited: the first one extracts appearance features from image sequences, learning temporal correlations, while the other network extracts shape features from a set of facial landmarks. The two nets are combined to yield the final decision on the emotion class. Happy and Routray [5] identify salient areas with generalized discriminative features for expression classification. They only use appearance-based features, and they do not consider the time domain. The framework from Jain et al. [6] recognizes facial expressions from video sequences by modeling temporal variations within shapes. They show that shape provides important information that is sometimes hard to grasp from appearance only. Zhang and Huang [16] propose a mixed model which include a “temporal” and a “spatial” network. The former captures dynamic features from consecutive frames, while the latter is about extracting static features from still frames. More generally, we can roughly characterize the popular trends in the existing literature by the usage of (*i.*) appearance-related (i.e., visual) features, (*ii.*) shape-related features, (*iii.*) features from face parts, (*iv.*) the temporal domain (i.e., video data).

This paper investigates the application of a pool of Convolutional Neural Networks (CNNs) with the aim of building recognizers of expressions in static images, that can be further applied to video sequences. We consider both (*i.*) appearance and (*ii.*) shape features, but, differently from most of the existing works, we do not hand-engineer shape features, and we let the CNNs learn the right representations from special shape-only images. We propose a model that considers (*iii.*) sub-parts of the face in addition to the entire face, motivated by the need of gaining deeper insights in the role of each component. Then, we move to the Semi-Supervised setting, exploiting (*iv.*) video data. The unsupervised portion of the training data is used to enforce “temporal coherence” among consecutive frames, and “part coherence” in each frame, i.e., a coherent

prediction among the CNNs that operate on the different face parts. Our experimental analysis shows that coherence constraints can improve the quality of the expression recognizer, thus offering a suitable basis to profitably exploit unsupervised video sequences.

This paper is organized as follows. The next Section formalizes the problem of facial expression recognition. Section 3 introduces our model. The role of coherence is described in Section 4, while experiments and conclusions (and future work) are collected in Section 5 and Section 6, respectively.

2 Facial Expression Recognition

The task of facial expression recognition that we consider in this paper consists in building a classifier that predicts one of the six universal emotions [2], that are *anger*, *disgust*, *fear*, *happiness*, *sadness*, *surprise*, plus the *neutral* case, and that we collect into the set Y , codified with indices from 1 to 7. The most popular inputs of the recognizer are images of faces, represented in foreground, usually with frontal orientation. When video data is considered, the recognition problem focusses on short video clips where a transition from the *neutral* state toward one the six emotions is recorded. Processing videos instead of still images can improve the recognition performance because facial expressions involve variations of the facial muscles along the temporal dimension. However, classifiers that are specifically trained to build a latent representation from a video clip \mathcal{V} before taking a decision [7], cannot be immediately applied to classify images. Differently, image-based classifiers can process single frames $\{\mathcal{I}_t\}$ of a video (being t the time index) to produce a final decision over a time window, so they are more versatile from the point of view of easiness of deployment in different real-world applications. The facial expression recognition problem is usually faced in the “Fully-Supervised” setting, and, in the case of videos, the available datasets are composed of labeled video clips where we do not have access to the labelings of the single frames³. Nonetheless, obtaining supervised data is costly, while nowadays is pretty easy to have access to collections of unsupervised frontal view faces (web, social networks, smartphones, ...) or unsupervised video recordings (video conference/call applications). This suggests that studying the “Semi-Supervised” setting, where a portion of the training data is labeled and a larger portion is unsupervised, can be a promising way to approach the recognition task.

Motivated by the need of building a versatile emotion recognition system, we focus on a predictor that operates on still images and that we can use to make predictions on video data. The system can be trained exploiting both video and image data in a Semi-Supervised setting, taking advantage of the temporal evolution described by the video format. In detail, we consider a classifier $f(\cdot)$ that produces a decision $y \in Y$ for each input image \mathcal{I} , or for a set of consecutive

³ See CK+ <http://www.consortium.ri.cmu.edu/ckagree/>, Oulu-CASIA <http://www.cse.oulu.fi/CMV/Downloads/Oulu-CASIA>, MMI <https://mmifacedb.eu/>.

frames belonging to a time window W (that covers a video clip, for example),

$$y = f(\mathcal{I}) \tag{1}$$

$$y = \text{majority}_{t \in W} \{f(\mathcal{I}_t)\} , \tag{2}$$

where **majority** is the majority-voting function, that returns the most frequent prediction in the time window W . Differently from the existing approaches, our system can be trained using labeled and unlabeled image databases, collected in $\mathcal{D}_{\mathcal{I}}$, or labeled and unlabeled frames extracted from the previously described labeled video sequences, collected in $\mathcal{D}_{\mathcal{V}}$. Due to the aforementioned properties of the existing video datasets (containing transitions from *neutral* to a certain emotion), we can artificially generate $\mathcal{D}_{\mathcal{V}}$ by labeling as *neutral* the very first frames of each video clip, and by assigning the provided video label to the last frames of the sequence. The frames in the internal portion of the sequence are not labeled. Formally, we have

$$\mathcal{D}_{\mathcal{I}} = \{(\mathcal{I}_i, y_i), i = 1, \dots, l\} \cup \{(\mathcal{I}_i, \text{none}), i = l + 1, \dots, l + u\} ,$$

where $y_i \in Y$ is the image label, and the rightmost set is fully unlabeled. Then,

$$\mathcal{D}_{\mathcal{V}} = \{\mathcal{D}_{\mathcal{V}_z}, z = 1, \dots, v\} ,$$

where v is the number of available video clips and $\mathcal{D}_{\mathcal{V}_z}$ is a sequence extracted from the z -th clip,

$$\begin{aligned} \mathcal{D}_{\mathcal{V}_z} = & ((\mathcal{I}_{z,t}, \text{neutral}), t = 1, \dots, \alpha|\mathcal{V}_z|) \oplus \\ & ((\mathcal{I}_{z,t}, \text{none}), t = \alpha|\mathcal{V}_z| + 1, \dots, \beta|\mathcal{V}_z|) \oplus \\ & ((\mathcal{I}_{z,t}, y_z), t = \beta|\mathcal{V}_z| + 1, \dots, |\mathcal{V}_z|) , \end{aligned}$$

being \oplus the sequence concatenation operator, $\mathcal{I}_{z,t} \in \mathcal{V}_z$ the t -th frame of the z -th video, and $0 < \alpha < \beta < 1$, arbitrarily chosen. In this case $y_z \in Y \setminus \{\text{neutral}\}$ is the label provided with the video clip \mathcal{V}_z (**neutral** is the identifier of the *neutral* class). We notice that $\mathcal{D}_{\mathcal{V}}$ is more informed than $\mathcal{D}_{\mathcal{I}}$, since it also stores the image/frame order and the frame grouping with respect to the videos. For this reason, we can consider $\mathcal{D}_{\mathcal{I}}$ to be an instance of the more general representation $\mathcal{D}_{\mathcal{V}}$, and in the rest of the paper we will focus on data represented as in $\mathcal{D}_{\mathcal{V}}$ without reducing the generality of what we described so far, and we will compactly indicate it with \mathcal{D} .

3 Model

Our model is based on CNNs that process two categories of representations of the input image/frame \mathcal{I} . Such categories consist in *appearance*-based (i.e, visual) representations and a *shape*-based representations.

In both the cases, we do not consider the whole \mathcal{I} , but only the rectangular area that is covered by the target face. We localize the face first, and then we

crop the image accordingly. This choice is crucial when processing inputs with multiple faces or when the face is not well positioned at the center of the image (or more generally, at a position incoherent with the training data). The *appearance*-based representation of the face is simply a grayscale instance of the cropped face. In the case of the *shape*-based representation, we still focus on the same cropped region, but we extract a set of shape features that essentially describe the contours of the face parts, and that, in this work, consist of a set of facial landmark points. However, instead of stacking their 2D coordinates into a vector (that is only possible if the set of points is consistent among different faces), we consider a more generic approach in which the shape is simply represented by an artificial image with uniform background and in which the landmarks points are depicted at their coordinates. This allows us to treat the shape in a way that is similar to what we do with the appearance, and it opens the possibility of providing different shape “sketches” that are not only based on landmark points (but also on contour lines, for example).

In order to study the effects of the different face parts in the recognition process, we computed the appearance and shape representations for the face (as just described) and for all the face parts: mouth, nose, eyes, eyebrows. We localized the face area and a set of 68 landmark points using the localizer of Viola and Jones [15] and a landmark detector [8]⁴. The detector uses the classic Histogram of Oriented Gradients (HOG) features combined with a linear classifier, an image pyramid, and a sliding window detection scheme. Cropping around each set of part-related landmarks (adding a small padding), we obtained 7 instances of appearance-based representations of the input \mathcal{I} and 8 shape-based ones, since in the case of shape we also included the landmarks associated to the jaw contour. Figure 1 shows the overall 15 representations that we generate. We resized these representations to the following sizes: face area 200×200 , mouth area 80×50 , eye area 60×30 , eyebrow area 100×30 , nose area 60×100 pixels, jaw area 200×170 .




Fig. 1. Representations extracted from an input image. On the left there are the 7 appearance-based representations. On the right there are the 8 shape-based representations, that we implement by sketching landmark points in artificial images.

We implemented a pool of 15 CNNs, each of them processing one of the aforementioned representations (Figure 2). The generic CNN_h associated to the

⁴ We used OpenCV <https://opencv.org/> and the “dlib” library <http://dlib.net/>

h -th representation has two convolutional layers followed by max pooling, and some fully connected layers terminated with a softmax activation that outputs a probability distribution over the emotions in Y . We indicate with $p_h(\cdot)$ the function computed by such CNN_h . All the hidden neural units have ReLU activation functions. The face-related CNNs have 32 and 64 filters on the two convolutional layers, respectively, and two fully connected layers (64 and $|Y| = 7$ neurons). The other CNNs, that are based on inputs with smaller sizes, exploit 16 and 32 filters, and a single fully connected layer ($|Y| = 7$ neurons).

The output of each of the 15 CNNs, when followed by an arg max operation (assuming 1-based indexing), is a possible instance of the function f in Eq. (1) and Eq. (2). Formally, for a given h ,

$$\begin{aligned} x_h &= \text{representation}_h(\mathcal{I}) \\ p_h(x_h) &= \text{CNN}_h(x_h) \\ f(\mathcal{I}) &= \arg \max p_h(x_h), \end{aligned}$$

where x_h is the h -representation of the input, and $p_h(x_h)$ outputs a vector of size $|Y|$ that sums to 1. Even if our final goal is to focus on the case in which h is the index of the full-face-based classifier, in Section 5 we will evaluate the quality of multiple instances of f , considering the predictors on the face parts too. In the next Section we will introduce a link between the full-face and face parts.

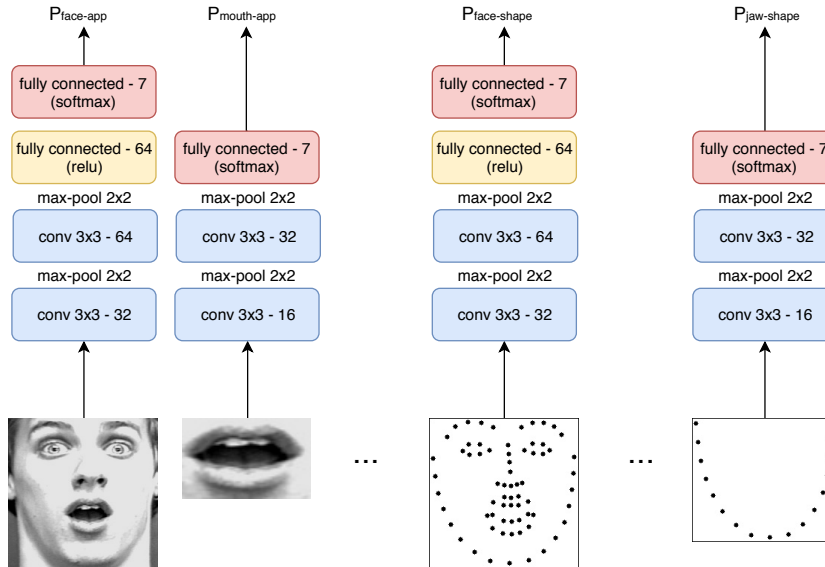


Fig. 2. Structure of CNNs employed.

4 Learning by Enforcing Coherence

We trained the pool of CNNs by minimizing an objective function involving the cross-entropy $L(p_h(x_h), y)$ between the outputs of the networks and the available labels (one-hot encoding), considering the training data $\mathcal{T} \subset \mathcal{D}$. The cross-entropy only exploits the labeled pairs in \mathcal{T} . However, our objective function is also composed by the penalties associated to the fulfilment of “coherence constraints” that we enforce on all the samples of \mathcal{T} , being them labeled or not. We have considered two types of coherence, namely “temporal coherence” and “coherence among the predictions on the face parts”. The former enforces the CNNs to be coherent over time for each video sequence, i.e., it enforces the predictions to smoothly change along the time axis. This constraint introduces a regularizing effect, since it prevents the system from developing unstable models that abruptly change their decisions among consecutive frames⁵. The part-based coherence enforces each full-face-representation-based classifiers to take decisions that are coherent with the ones taken (*on average*) by the other part-based classifiers (and vice-versa). The idea behind this constraint is that the committee of the local (i.e. part-based) predictors could provide important fine-grained information that the global (face-based) predictor might not have been able to capture. We already experimented some related constraints in the case of multi-view object recognition [12], and these ideas are borrowed by the generic framework of “Learning from Constraints” [4], where a predictor is constrained exploiting high-level knowledge on the task at hand, bridging the symbolic and sub-symbolic worlds.

In detail, given two scalars $\lambda_t, \lambda_c \geq 0$ that weigh the importance of the coherence (soft) constraints, our objective function for the appearance-based classifiers (or, equivalently, for the shape-based classifiers) is

$$\sum_h \sum_{\substack{i=1 \\ y_i \neq \text{none}}} w_i \cdot L(p_h(x_{h,i}), y_i) + \lambda_t \overbrace{\sum_h \sum_{z=1}^v \sum_{t=2}^{|\mathcal{V}_z|} (1 - p_h(x_{h,(z,t-1)})' \cdot p_h(x_{h,(z,t)}))}^{\text{temporal coherence}} + \lambda_c \underbrace{\sum_{h \neq \text{face}} \sum_i (1 - p_{\text{face}}(x_{h,i})' \cdot p_h(x_{h,i}))}_{\text{part coherence}} \quad (3)$$

where the index h spans over the 7 appearance-based classifiers (or the 8 shape-based classifiers). The index i spans over all the pairs in \mathcal{T} , and, for the sake of simplicity, we used the notation $y_i \neq \text{none}$ to indicate that we consider only the labeled examples. The scalar weights w_i are used to give custom weights to the examples, and we used them to give more importance to the classes that are less represented in \mathcal{T} . The notation (z, t) is the index of the t -th frame in the

⁵ We remark that the enforcement of both the coherence constraints only happens at training time.

z -th video sequence belonging to \mathcal{T} . Finally, **face** is used to indicate the index associated with the full-face input, and $'$ is the transpose operator.

We notice that since $p(\cdot)$ is a probability distribution, the dot products involving two instances of $p(\cdot)$ are 1 when such instances are equivalent (and the coherence constraints are fulfilled). The temporal constraint involves dot products between the predictions on pairs of consecutive frames in the same video clip. We kept the same structure to build the part-based constraint, where the averaging operation on the part-based classifiers is evident when $\sum_{h \neq \mathbf{face}}$ is moved right before the second term of the dot product $p_{\mathbf{face}}(x_{h,i})' \cdot p_h(x_{h,i})$.

5 Experimental Results

In order to validate our model, we used the popular Extended Cohn-Kanade dataset (CK+) [11]. It consists of 593 frames belonging to a set of short video sequences, where 120 subjects (different age and gender) generate expressions belonging to the following list: anger, contempt, disgust, fear, happiness, sadness and surprise. We excluded the sequences associated to “contempt”, which is not included into the six universal emotions. The video sequences are composed of 10-60 frames, they start with a neutral expression and they end with the peak of one of the previously listed expressions. Each sequence is associated with an emotion label.

In order to build the Semi-Supervised set \mathcal{D} described in Section 2, we selected $\alpha = 0.1$ and $\beta = 0.7$. We generated 5 randomizations of the whole dataset, and divided each of them into training (70%), validation (15%), and test sets (15%), keeping the original distribution of the classes in each set. The validation data was used to validate the model parameters and excluded from training. The test partition was used to measure the quality of the model, and the results presented in this Section are averaged over the 5 test partitions (when available, we also report the standard deviation in brackets). Each collection of training data consists of about $\approx 4,000$ frames, out of which $\approx 1,500$ are labeled, and they are organized into ≈ 200 sequences, while the validation data is composed of ≈ 600 frames, out of which ≈ 200 are labeled, and organized into ≈ 30 sequences. Since examples from the “neutral” class are much more represented with respect to other examples, we set $w_i = 0.1$ in Eq. (3) if i is an example from the neutral class, $w_i = 1$ otherwise. We selected the optimal λ_c, λ_t by a grid-search in $\{10^{-10}, 10^{-8}, 10^{-7}, 10^{-6}, 10^{-4}, 10^{-2}\}$, measuring frame-level accuracy (i.e., only the labeled validation frames are considered). We implemented our model using TensorFlow, and we minimized Eq. (3) by the Adam-based optimizer (starting learning rate 0.001), mini-batches of size 96, and we have trained the model for multiple epochs, stopping the procedure when the validation error started increasing.

We performed experiments comparing a system with no-coherence-constraints ($\lambda_c = \lambda_t = 0$) with other models that include either temporal or part-based coherence. We compared the cases of single-frame-level predictions (where only the labeled portion of the test set is considered) and the case of video-sequence-

level predictions, following the decision rules of Eq. (1) and Eq. (2), respectively (where W covers the full video sequence). Since examples of the different classes are not balanced in the given dataset, and in order to provide a more informative set of results, we measured two types of accuracies, namely **Micro** and **Macro** accuracies. The former is simply the percentage of correctly labeled frames/sequences, while the latter is the average of the percentages of correctly labeled frames/sequences in each emotion class.

Table 1 shows the results we obtain when testing the classifiers that operate on the full-face inputs, considering both appearance and shape representations. We also report results of an additional classifier obtained by averaging the outputs of the full set of 15 classifiers (thus mixing appearance and shape data).

Table 1. Micro and macro accuracies (std dev. in brackets) at image and video (sequence) level of the full-face-based classifiers (appearance and shape representations) and of an ensemble of the 15 classifiers (average of 15 outputs, both shape and appearance). Results without coherence constraints (NONE), with PART-based coherence and TEMP-oral coherence (results where coherence improves the accuracy are in bold).

	IMAGES					
	% Micro Acc			% Macro Acc		
	NONE	PART	TEMP	NONE	PART	TEMP
Face_{app.}	78.9 (3.6)	78.0 (2.0)	81.1 (3.0)	71.2 (2.8)	72.8 (2.2)	72.2 (7.4)
Face_{shape}	71.8 (3.0)	71.9 (3.1)	72.5 (2.9)	61.1 (2.9)	61.3 (3.0)	62.1 (2.7)
Avg_{all}	73.7 (4.1)	71.4 (3.1)	72.1 (4.8)	71.9 (3.9)	70.2 (3.3)	69.7 (3.7)

	VIDEOS					
	% Micro Acc			% Macro Acc		
	NONE	PART	TEMP	NONE	PART	TEMP
Face_{app.}	75.3 (5.1)	77.0 (3.4)	80.0 (2.9)	64.0 (3.2)	66.8 (3.1)	64.4 (10.3)
Face_{shape}	68.5 (3.0)	68.1 (3.1)	69.4 (2.9)	54.0 (2.9)	53.5 (3.0)	55.5 (2.7)
Avg_{all}	78.3 (4.9)	77.9 (2.5)	80.4 (5.5)	65.6 (6.5)	65.9 (3.9)	64.8 (7.4)

Temporal coherence always improves the quality of the face-based classifiers, up to 5% in the case of sequences (micro). In the case of macro-accuracy we observe larger standard deviations, that are due to the effects of the predictions on the classes with a smaller number of examples. Such classes are less-frequently predicted, and asking for a strong temporal regularization sometimes further reduces such frequency. Coherence among parts helps in a less evident manner, especially when using shapes. Shape is less informative than appearance, resulting in a performance drop of $\approx 10\%$. The average-based classifier is only in some cases better than the face-based ones. Constraints are less effective in this case (even if we get a strong micro accuracy in videos + temporal coherence). This suggests that mixing the 15 classifiers together is not a promising direction, mostly because some of them have low performances that can degrade the average quality of the system.

To gain better insights about the last comment, Table 2 reports the accuracies for all the part-based classifiers. The mouth area is a very effective input for facial expression recognition, that can sometimes compete with the full-face. This is

Table 2. Micro and macro accuracies (std dev. in brackets) at image and video level of all the part-based classifiers (appearance and shape representation). Results without coherence constraints (NONE), with PART-based coherence and TEMP-oral coherence (results where coherence improves the accuracy are in bold).

	IMAGES					
	% Micro Acc			% Macro Acc		
	NONE	PART	TEMP	NONE	PART	TEMP
Mouth _{app.}	70.5 (3.5)	68.6 (3.0)	72.8 (2.6)	71.5 (6.7)	70.8 (5.8)	73.3 (4.4)
Left-eye _{app.}	42.3 (6.0)	41.4 (6.0)	40.0 (4.2)	41.3 (6.5)	39.1 (4.9)	38.5 (3.9)
Right-eye _{app.}	42.0 (5.6)	42.0 (7.3)	40.6 (5.2)	40.8 (5.7)	40.5 (5.7)	38.8 (5.6)
Left-eyebrow _{app.}	40.5 (6.8)	37.7 (7.3)	38.4 (9.1)	40.1 (6.1)	37.4 (7.5)	37.6 (8.4)
Right-eyebrow _{app.}	40.1 (2.5)	39.7 (2.4)	40.4 (2.9)	40.1 (3.5)	39.5 (2.8)	40.3 (3.1)
Nose _{app.}	43.6 (2.9)	44.1 (5.5)	43.4 (4.0)	41.6 (3.4)	42.4 (4.8)	42.0 (3.7)
Mouth _{shape}	64.3 (2.3)	63.8 (3.5)	63.4 (3.2)	64.4 (4.7)	63.4 (4.8)	66.2 (4.9)
Left-eye _{shape}	35.8 (3.4)	34.5 (3.7)	35.2 (2.6)	33.2 (3.9)	33.0 (3.4)	32.5 (2.3)
Right-eye _{shape}	40.7 (3.2)	40.6 (2.7)	41.5 (3.0)	36.9 (2.4)	37.2 (2.1)	37.9 (2.0)
Left-eyebrow _{shape}	31.2 (4.4)	31.0 (3.8)	30.1 (3.5)	31.8 (1.8)	31.9 (2.0)	31.7 (3.7)
Right-eyebrow _{shape}	34.3 (4.2)	33.9 (3.7)	34.1 (3.5)	34.3 (5.2)	33.4 (4.5)	33.6 (4.9)
Nose _{shape}	30.8 (3.7)	30.4 (3.2)	30.9 (4.2)	30.6 (5.6)	31.0 (5.0)	31.6 (5.2)
Jaw _{shape}	37.4 (3.7)	37.2 (3.7)	37.0 (3.5)	34.1 (4.6)	34.9 (4.3)	33.8 (4.0)

	VIDEOS					
	% Micro Acc			% Macro Acc		
	NONE	PART	TEMP	NONE	PART	TEMP
Mouth _{app.}	77.5 (7.7)	72.3 (9.0)	75.7 (6.4)	73.0 (9.5)	66.4 (8.4)	69.9 (8.7)
Left-eye _{app.}	49.4 (8.4)	50.6 (4.1)	47.2 (5.9)	42.7 (5.8)	41.3 (2.7)	40.2 (6.3)
Right-eye _{app.}	46.8 (2.3)	47.2 (4.9)	47.7 (2.9)	39.8 (1.7)	39.2 (3.0)	38.9 (3.7)
Left-eyebrow _{app.}	43.0 (9.7)	41.7 (9.2)	42.1 (11.1)	35.2 (7.7)	34.3 (9.1)	34.3 (9.6)
Right-eyebrow _{app.}	43.4 (4.6)	42.5 (5.5)	43.8 (3.2)	36.5 (6.6)	35.6 (6.8)	35.9 (4.0)
Nose _{app.}	44.3 (4.9)	47.7 (5.1)	47.2 (2.8)	35.4 (4.3)	38.8 (4.3)	38.9 (3.1)
Mouth _{shape}	71.9 (2.5)	74.0 (3.7)	70.6 (2.8)	64.3 (4.2)	66.1 (6.0)	67.3 (5.0)
Left-eye _{shape}	45.1 (5.8)	44.7 (8.5)	45.1 (4.5)	36.6 (7.1)	37.2 (6.1)	38.3 (4.1)
Right-eye _{shape}	51.9 (2.2)	52.8 (3.7)	56.2 (3.7)	39.4 (3.1)	41.5 (3.3)	44.9 (3.9)
Left-eyebrow _{shape}	36.2 (6.7)	34.5 (3.4)	34.9 (3.5)	28.7 (5.1)	28.7 (3.0)	29.3 (4.1)
Right-eyebrow _{shape}	40.4 (5.0)	40.0 (5.9)	41.3 (6.7)	33.9 (5.6)	33.1 (5.0)	33.8 (7.0)
Nose _{shape}	37.5 (5.0)	35.7 (3.7)	34.0 (1.4)	31.4 (5.4)	28.5 (5.6)	31.8 (4.4)
Jaw _{shape}	40.9 (2.5)	40.9 (2.1)	40.0 (3.7)	30.5 (2.5)	31.3 (2.7)	29.8 (2.7)

more evident in the case of videos, when comparing shape-based representations of face and mouth. As expected, the other parts are worse than the full-face, since they are just local views. The addition of both coherences sparsely helps in improving the local classifiers, with a preference toward temporal coherence. The worst results are obtained by eyebrows and nose in shape-based classification. Interestingly, the eye-based predictors score the most effective results after face and mouth in video sequences. While their appearance representation is altered when eyes get closed, their shape representation is more stable. This analysis suggests that an accurate choice of a sub-portion of the face parts could significantly help the part-based coherence constraint (since some of the parts are not very informative).

We deepened the analysis on the temporal-constrained classifiers in the case of making predictions in video sequences. Since the number of sequences is small, we selected the optimal λ_t using image-level predictions on the validation data (as already stated), leading to $\lambda_t = 10^{-8}$ and $\lambda_t = 10^{-2}$ in the case of micro and macro accuracy, respectively. Figure 3 reports the performances on videos with different values of λ_t (appearance only). We can see that the distributions

of the performances are multimodal, and if we focus on the macro accuracy we observe that we could have obtained much better results with different values of λ_t . This suggests that the validation procedure has room for being improved in the case of video data.




Fig. 3. Micro and macro accuracies in the case of video data, full-face-based classifier (appearance), for different values of λ_t . The black-bordered bars are the results we reported in Table 1.

In Table 3 we show the results on single emotion classes for face and mouth appearance-based classification, focussing on the case where no-coherence is introduced and the ones with a selection of the best $\lambda_c > 0$ and $\lambda_t > 0$ from the previously described experiments. “Fear” and “sadness” classes are difficult to classify because they do not involve strong facial movements, while “happiness” and “surprise” are easy to recognize. The mouth-based model has difficulties in the “neutral” class, since some emotions do not evidently alter the mouth area (the face model does not show this issue). In the “sadness” class, where the face-based model scores low accuracies, the mouth-based classifier is much more performant. This suggests that the face-related network has difficulties in developing a generalizable representation for the whole face to identify “sadness”. Larger training data could help in this case.

Temporal coherence shows better performance in “neutral”, “anger” (image-level only), and “disgust” emotions. It is also helpful in the “happiness” class, where the face model performs a close-to-flawless classification. Introducing coherence among parts improves the recognition of “disgust”, “fear” (face only), “happiness” (image-level only), and it slightly improves the accuracy of “surprise” for the face-based predictor.

In addition to these results, we report that eye-based recognition reaches very good results for the “surprise” class; the accuracy of right-eye classifier with temporal coherence is 88.2%. This is due to the fact that the eyes in surprise expressions are wide open, so easily recognizable. Differently, the “neutral” class is not recognizable at all from the eyebrows. Nose-based classification (appearance) reaches an accuracy of 79.4% with temporal coherence in the “disgust” class, where the nose is wrinkled. As a final comment, we have also tried to

Table 3. Accuracies (std dev. in brackets) on each class of full-face and mouth classifiers (appearance). Results without coherence constraints, with PART-based coherence and TEMPORAL coherence (results where coherence improves the accuracy are in bold).

		IMAGES						
		<i>Anger</i>	<i>Disgust</i>	<i>Fear</i>	<i>Happiness</i>	<i>Sadness</i>	<i>Surprise</i>	<i>Neutral</i>
face _{app.}	NONE	73.7 (8.8)	69.2 (12.4)	56.1 (28.9)	92.5 (9.8)	29.5 (27.7)	96.1 (1.3)	81.1 (14.9)
face _{app.}	PART	68.0 (28.0)	78.2 (10.0)	75.2 (11.8)	98.2 (2.9)	24.3 (21.0)	97.4 (0.7)	68.6 (9.3)
face _{app.}	TEMP	77.1 (18.6)	81.8 (14.2)	50.0 (27.7)	97.5 (4.9)	26.2 (21.5)	95.5 (2.7)	81.8 (15.1)
mouth _{app.}	NONE	66.4 (18.5)	69.6 (19.1)	59.4 (9.5)	92.7 (7.0)	75.6 (16.0)	96.6 (2.5)	40.2 (11.8)
mouth _{app.}	PART	66.4 (31.8)	81.8 (14.5)	65.1 (13.3)	95.0 (3.8)	59.6 (27.5)	95.5 (3.1)	32.2 (12.2)
mouth _{app.}	TEMP	67.8 (19.8)	80.4 (12.3)	58.8 (9.7)	94.8 (5.9)	72.0 (20.8)	95.2 (3.4)	44.1 (6.8)

		VIDEOS						
		<i>Anger</i>	<i>Disgust</i>	<i>Fear</i>	<i>Happiness</i>	<i>Sadness</i>	<i>Surprise</i>	<i>Neutral</i>
face _{app.}	NONE	77.1 (7.0)	62.2 (13.3)	33.3 (21.1)	90.9 (14.1)	25.0 (27.4)	95.4 (6.2)	–
face _{app.}	PART	68.6 (31.8)	71.1 (16.6)	53.3 (16.3)	90.9 (11.5)	20.0 (18.7)	96.9 (3.8)	–
face _{app.}	TEMP	77.1 (26.5)	73.3 (11.3)	40.0 (13.3)	98.2 (3.6)	25.0 (15.8)	96.9 (3.8)	–
mouth _{app.}	NONE	77.1 (17.1)	73.3 (16.6)	46.7 (26.7)	78.2 (14.8)	75.0 (27.4)	87.7 (11.5)	–
mouth _{app.}	PART	62.9 (34.5)	77.8 (14.1)	46.7 (16.3)	74.6 (15.6)	55.0 (36.7)	81.5 (9.2)	–
mouth _{app.}	TEMP	74.3 (21.0)	77.8 (14.1)	40.0 (24.9)	74.6 (14.6)	65.0 (33.9)	87.7 (10.4)	–

perform some preliminary experiments involving both temporal and part-based coherences activated, and while we obtained good results on average (showing improvements with respect to the unconstrained case), they were not better than the “best” ones we obtained by activating only one of the two coherences. However, we will further investigate this point with a more detailed cross-validation and larger data collections.

6 Conclusions and Future Work

We presented a Convolutional Neural Network (CNN)-based approach to Facial Expression Recognition. Our model is based on a pool of CNNs that process distinct face parts, represented using visual (appearance) or shape-only features. In the latter case, we treated shape as a generic input of the learnable model, without manually engineering its representation. We studied the importance of the different representations on the task at hand, showing an analysis that involved all the considered face parts, and reporting results of experiments on a popular dataset composed of six basic emotions, plus the neutral case. We proposed the introduction of coherence constraints among the face-part predictors and between predictions on consecutive time instants, casting the learning problem in the Semi-Supervised setting and using video data. Our results have shown that using unsupervised training data paired with coherence constraints improves the quality of the recognizer, especially in the case of temporal coherence. Our future work will include a more detailed study on the face-part coherence, selecting only on the most promising face parts, according to the results of this study. We will also consider introducing coherence between shape and appearance-based predictors, and the use of larger collections of data, to grasp the importance of large-scale unsupervised data obtained from video conferences.

References

1. Duchenne, G.B., de Boulogne, G.B.D.: The mechanism of human facial expression. Cambridge University press (1990)
2. Ekman, P., Friesen, W.V.: Constants across cultures in the face and emotion. *Journal of Personality and Social Psychology* 17(2), 124 (1971)
3. Fan, X., Tjahjadi, T.: A spatial-temporal framework based on histogram of gradients and optical flow for facial expression recognition in video sequences. *Pattern Recognition* 48(11), 3407–3416 (2015)
4. Gnecco, G., Gori, M., Melacci, S., Sanguineti, M.: Foundations of support constraint machines. *Neural Computation* 27(2), 388–480 (2015)
5. Happy, S., Routray, A.: Automatic facial expression recognition using features of salient facial patches. *IEEE Transactions on Affective Computing* 6(1), 1–12 (2015)
6. Jain, S., Hu, C., Aggarwal, J.K.: Facial expression recognition with temporal modeling of shapes. In: *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*. pp. 1642–1649. IEEE (2011)
7. Jung, H., Lee, S., Yim, J., Park, S., Kim, J.: Joint fine-tuning in deep neural networks for facial expression recognition. In: *Computer Vision (ICCV), 2015 IEEE International Conference on*. pp. 2983–2991. IEEE (2015)
8. Kazemi, V., Sullivan, J.: One millisecond face alignment with an ensemble of regression trees. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 1867–1874 (2014)
9. Long, F., Bartlett, M.S.: Video-based facial expression recognition using learned spatiotemporal pyramid sparse coding features. *Neurocomputing* 173, 2049–2054 (2016)
10. Lopes, A.T., de Aguiar, E., De Souza, A.F., Oliveira-Santos, T.: Facial expression recognition with convolutional neural networks: coping with few data and the training sample order. *Pattern Recognition* 61, 610–628 (2017)
11. Lucey, P., Cohn, J.F., Kanade, T., Saragih, J., Ambadar, Z., Matthews, I.: The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In: *Computer Vision and Pattern Recognition Workshops (CVPRW), 2010 IEEE Computer Society Conference on*. pp. 94–101. IEEE (2010)
12. Melacci, S., Maggini, M., Gori, M.: Semi-supervised learning with constraints for multi-view object recognition. In: *International Conference on Artificial Neural Networks*. pp. 653–662. Springer (2009)
13. Mollahosseini, A., Chan, D., Mahoor, M.H.: Going deeper in facial expression recognition using deep neural networks. In: *Applications of Computer Vision (WACV), 2016 IEEE Winter Conference on*. pp. 1–10. IEEE (2016)
14. Plutchik, R.: The nature of emotions: Human emotions have deep evolutionary roots, a fact that may explain their complexity and provide tools for clinical practice. *American Scientist* 89(4), 344–350 (2001)
15. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: *Computer Vision and Pattern Recognition, 2001. CVPR 2001. Proceedings of the 2001 IEEE Computer Society Conference on*. vol. 1, pp. I–I. IEEE (2001)
16. Zhang, K., Huang, Y., Du, Y., Wang, L.: Facial expression recognition based on deep evolutionary spatial-temporal networks. *IEEE Transactions on Image Processing* 26(9), 4193–4203 (2017)