

# Context Based Emotion Recognition using EMOTIC Dataset

Ronak Kosti, Jose M. Alvarez, Adria Recasens, Agata Lapedriza

**Abstract**—In our everyday lives and social interactions we often try to perceive the emotional states of people. There has been a lot of research in providing machines with a similar capacity of recognizing emotions. From a computer vision perspective, most of the previous efforts have been focusing in analyzing the facial expressions and, in some cases, also the body pose. Some of these methods work remarkably well in specific settings. However, their performance is limited in natural, unconstrained environments. Psychological studies show that the scene context, in addition to facial expression and body pose, provides important information to our perception of people's emotions. However, the processing of the context for automatic emotion recognition has not been explored in depth, partly due to the lack of proper data. In this paper we present EMOTIC, a dataset of images of people in a diverse set of natural situations, annotated with their apparent emotion. The EMOTIC dataset combines two different types of emotion representation: (1) a set of 26 discrete categories, and (2) the continuous dimensions *Valence*, *Arousal*, and *Dominance*. We also present a detailed statistical and algorithmic analysis of the dataset along with annotators' agreement analysis. Using the EMOTIC dataset we train different CNN models for emotion recognition, combining the information of the bounding box containing the person with the contextual information extracted from the scene. Our results show how scene context provides important information to automatically recognize emotional states and motivate further research in this direction.

**Index Terms**—Emotion recognition, Affective computing, Pattern recognition

## 1 INTRODUCTION

Over the past years, the interest in developing automatic systems for recognizing emotional states has grown rapidly. We can find several recent works showing how emotions can be inferred from cues like text [1], voice [2], or visual information [3], [4]. The automatic recognition of emotions has a lot of applications in environments where machines need to interact or monitor humans. For instance, automatic tutors in an online learning platform would provide better feedback to a student according to her level of motivation or frustration. Also, a car with the capacity of assisting a driver can intervene or give an alarm if it detects the driver is tired or nervous.

In this paper we focus on the problem of emotion recognition from visual information. Concretely, we want to recognize the apparent emotional state of a person in a given image. This problem has been broadly studied in computer vision mainly from two perspectives: (1) facial expression analysis, and (2) body posture and gesture analysis. Section 2 gives an overview of related work on these perspectives and also on some of the common public datasets for emotion recognition.

Although face and body pose give lot of information on the affective state of a person, our claim in this work is that scene context information is also a key component for understanding emotional states. Scene context includes the



Fig. 1: How is this kid feeling? Try to recognize his emotional states from the person bounding box, without scene context.

surroundings of the person, like the place category, the place attributes, the objects, or the actions occurring around the person. Fig. 1 illustrates the importance of scene context for emotion recognition. When we just see the kid it is difficult to recognize his emotion (from his facial expression it seems he is feeling *Surprise*). However, when we see the context (Fig. 2.a) we see the kid is celebrating his birthday, blowing the candles, probably with his family or friends at home. With this additional information we can interpret much better his face and posture and recognize that he probably feels *engaged*, *happy* and *excited*.

The importance of context in emotion perception is well supported by different studies in psychology [5], [6]. In general situations, facial expression is not sufficient to determine the emotional state of a person, since the perception of the emotion is heavily influenced by different types of context, including the scene context [2], [3], [4].

In this work, we present two main contributions.

- Ronak Kosti & Agata Lapedriza are with Universitat Oberta de Catalunya, Spain. Email: rkosti@uoc.edu, alapedriza@uoc.edu.
- Adria Recasens is with the Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, USA. Email: recasens@mit.edu.
- Jose M. Alvarez is with NVIDIA, USA. Email: jalvarez.research@gmail.com
- Project Page: <http://sunai.uoc.edu/emotic/>

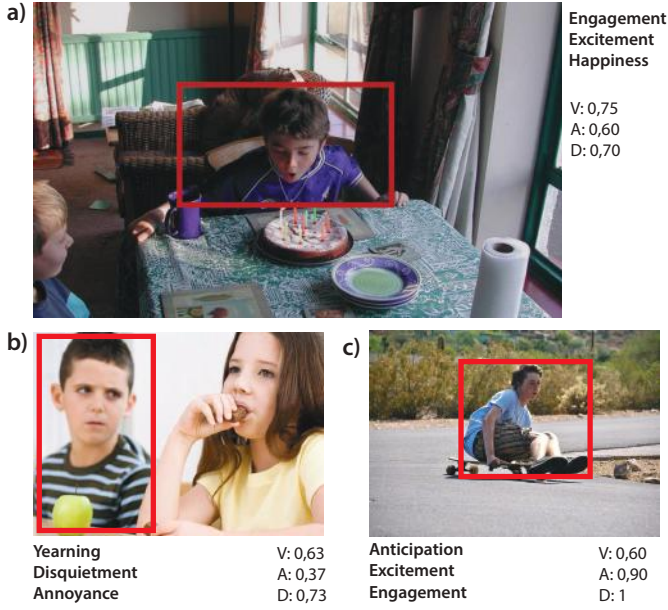


Fig. 2: Sample images in the EMOTIC dataset along with their annotations.

Our first contribution is the creation and publication of the **EMOTIC** (from EMOTions In Context) Dataset. The EMOTIC database is a collection of images of people annotated according to their apparent emotional states. Images are spontaneous and unconstrained, showing people doing different things in different environments. Fig. 2 shows some examples of images in the EMOTIC database along with their corresponding annotations. As shown, annotations combine 2 different types of emotion representation: Discrete Emotion Categories and 3 Continuous Emotion Dimensions *Valence*, *Arousal*, and *Dominance* [7]. The EMOTIC dataset is now publicly available for download at the EMOTIC website<sup>1</sup>. Details of the dataset construction process and dataset statistics can be found in section 3.

Our second contribution is the creation of a baseline system for the task of emotion recognition in context. In particular, we present and test a Convolutional Neural Network (CNN) model that jointly processes the window of the person and the whole image to predict the apparent emotional state of the person. Section 4 describes the CNN model and the implementation details while section 5 presents our experiments and discussion on the results. All the trained models resulting from this work are also publicly available at the EMOTIC website<sup>1</sup>.

This paper is an extension of the conference paper “Emotion Recognition in Context”, presented at the IEEE International Conference on Computer Vision and Pattern Recognition (CVPR) 2017 [8]. We present here an extended version of the EMOTIC dataset, with further statistical dataset analysis, an analysis of scene-centric algorithms on the data, and a study on the annotation consistency among different annotators. This new release of the EMOTIC database contains 44.4% more annotated people as compared to its previous smaller version. With the new extended dataset we retrained all the proposed baseline CNN models with additional loss

functions. We also present comparative analysis of two different scene context features, showing how the context is contributing to recognize emotions in the wild.

## 2 RELATED WORK

Emotion recognition has been broadly studied by the Computer Vision community. Most of the existing work has focused on the analysis of facial expression to predict emotions [9], [10]. The base of these methods is the *Facial Action Coding System* [11], which encodes the facial expression using a set of specific localized movements of the face, called *Action Units*. These facial-based approaches [9], [10] usually use facial-geometry based features or appearance features to describe the face. Afterwards, the extracted features are used to recognize *Action Units* and the basic emotions proposed by Ekman and Friesen [12]: *anger*, *disgust*, *fear*, *happiness*, *sadness*, and *surprise*. Currently, state-of-the-art systems for emotion recognition from facial expression analysis use CNNs to recognize emotions or Action Units [13].

In terms of emotion representation, some recent works based on facial expression [14] use the continuous dimensions of the *VAD Emotional State Model* [7]. The VAD model describes emotions using 3 numerical dimensions: *Valence* (V), that measures how positive or pleasant an emotion is, ranging from *negative* to *positive*; *Arousal* (A), that measures the agitation level of the person, ranging from *non-active / in calm* to *agitated / ready to act*; and *Dominance* (D) that measures the level of control a person feels of the situation, ranging from *submissive / non-control* to *dominant / in-control*. On the other hand, Du et al. [15] proposed a set of 21 facial emotion categories, defined as different combinations of the basic emotions, like ‘happily surprised’ or ‘happily disgusted’. With this categorization the authors can give a fine-grained detail about the expressed emotion.

Although the research in emotion recognition from a computer vision perspective is mainly focused in the analysis of the face, there are some works that also consider other additional visual cues or multimodal approaches. For instance, in [16] the location of shoulders is used as additional information to the face features to recognize basic emotions. More generally, Schindler et al. [17] used the body pose to recognize 6 basic emotions, performing experiments on a small dataset of non-spontaneous poses acquired under controlled conditions. Mou et al. [18] presented a system of affect analysis in still images of groups of people, recognizing group-level arousal and valence from combining face, body and contextual information.

Emotion Recognition in Scene Context and Image Sentiment Analysis are different problems that share some characteristics. Emotion Recognition aims to identify the emotions of a person depicted in an image. Image Sentiment Analysis consists of predicting what a person will feel when observing a picture. This picture does not necessarily contain a person. When an image contains a person, there can be a difference between the emotions experienced by the person in the image and the emotions felt by observers of the image. For example, in the image of Figure 2.b, we see a kid who seems to be annoyed for having an apple instead of chocolate and another who seems happy to have chocolate. However, as observers, we might not have any

1. <http://sunai.uoc.edu/emotic/>

of those sentiments when looking at the photo. Instead, we might think the situation is not fair and feel disapproval. Also, if we see an image of an athlete that has lost a match, we can recognize the athlete feels sad. However, an observer of the image may feel happy if the observer is a fan of the team that won the match.

## 2.1 Emotion Recognition Datasets

Most of the existing datasets for emotion recognition using computer vision are centered in facial expression analysis. For example, the GENKI database [19] contains frontal face images of a single person with different illumination, geographic, personal and ethnic settings. Images in this dataset are labelled as *smiling* or *non-smiling*. Another common facial expression analysis dataset is the ICML Face-Expression Recognition dataset [20], that contains 28,000 images annotated with 6 basic emotions and a neutral category. On the other hand, the UCDSEE dataset [21] has a set of 9 emotion expressions acted by 4 persons. The lab setting is strictly kept the same in order to focus mainly on the facial expression of the person.

The dynamic body movement is also an essential source for estimating emotion. Studies such as [22], [23] establish the relationship between affect and body posture using as ground truth the base-rate of human observers. The data consist of a spontaneous set of images acquired under a restrictive setting (people playing Wii games). The GEMEP database [24] is multi-modal (audio and video) and has 10 actors playing 18 affective states. The dataset has videos of actors showing emotions through acting. Body pose and facial expression are combined.

The Looking at People (LAP) challenges and competitions [25] involve specialized datasets containing images, sequences of images and multi-modal data. The main focus of these datasets is the complexity and variability of human body configuration which include data related to personality traits (spontaneous), gesture recognition (acted), apparent age recognition (spontaneous), cultural event recognition (spontaneous), action/interaction recognition and human pose recognition (spontaneous).

The Emotion Recognition in the Wild (EmotiW) challenges [26] host 3 databases: (1) The AFEW database [27] focuses on emotion recognition from video frames taken from movies and TV shows, where the actions are annotated with attributes like name, age of actor, age of character, pose, gender, expression of person, the overall clip expression and the basic 6 emotions and a neutral category; (2) The SFEW, which is a subset of AFEW database containing images of face-frames annotated specifically with the 6 basic emotions and a neutral category; and (3) the HAPPEI database [28], which addresses the problem of group level emotion estimation. Thus, [28] offers a first attempt to use context for the problem of predicting happiness in groups of people.

Finally, the COCO dataset has been recently annotated with object attributes [29], including some emotion categories for people, such as *happy* and *curious*. These attributes show some overlap with the categories that we define in this paper. However, COCO attributes are not intended to be exhaustive for emotion recognition, and not all the people in the dataset are annotated with affect attributes.

## 3 EMOTIC DATASET

The EMOTIC dataset is a collection of images of people in unconstrained environments annotated according to their apparent emotional states. The dataset contains 23,571 images and 34,320 annotated people. Some of the images were manually collected from the Internet by Google search engine. For that we used a combination of queries containing various places, social environments, different activities and a variety of keywords on emotional states. The rest of images belong to 2 public benchmark datasets: COCO [30] and Ade20k [31]. Overall, the images show a wide diversity of contexts, containing people in different places, social settings, and doing different activities.

Fig. 2 shows three examples of annotated images in the EMOTIC dataset. Images were annotated using Amazon Mechanical Turk (AMT). Annotators were asked to label each image according to what they think people in the images are feeling. Notice that we have the capacity of making reasonable guesses about other people's emotional state due to our capacity of being empathetic, putting ourselves into another's situation, and also because of our common sense knowledge and our ability for reasoning about visual information. For example, in Fig. 2.b, the person is performing an activity that requires *Anticipation* to adapt to the trajectory. Since he is doing a thrilling activity, he seems *excited* about it and he is *engaged* or focused in this activity. In Fig. 2.c, the kid feels a strong desire (*yearning*) for eating the chocolate instead of the apple. Because of his situation we can interpret his facial expression as *disquietness* and *annoyance*. Notice that images are also annotated according to the continuous dimensions *Valence*, *Arousal*, and *Dominance*. We describe the emotion annotation modalities of EMOTIC dataset and the annotation process in sections 3.1 and 3.2, respectively.

After the first round of annotations (1 annotator per image), we divided the images into three sets: Training (70%), Validation (10%), and Testing (20%) maintaining a similar affective category distribution across the different sets. After that, Validation and Testing were annotated by 4 and 2 extra annotators respectively. As a consequence, images in the Validation set are annotated by a total of 5 annotators, while images in the Testing set are annotated by 3 annotators (these numbers can slightly vary for some images since we removed noisy annotations).

We used the annotations from the Validation to study the consistency of the annotations across different annotators. This study is shown in section 3.3. The data statistics and algorithmic analysis on the EMOTIC dataset are detailed in sections 3.4 and 3.5 respectively.

### 3.1 Emotion representation

The EMOTIC dataset combines two different types of emotion representation:

**Continuous Dimensions:** images are annotated according to the VAD model [7], which represents emotions by a combination of 3 continuous dimensions: Valence, Arousal and Dominance. In our representation each dimension takes an integer value that lies in the range [1 – 10]. Fig. 4 shows examples of people annotated by different values of the given dimension.





Fig. 3: Examples of annotated people in EMOTIC dataset for each of the 26 emotion categories (Table 1). The person in the red bounding box is annotated by the corresponding category.



Fig. 4: Examples of annotated images in EMOTIC dataset for each of the 3 continuous dimensions Valence, Arousal & Dominance. The person in the red bounding box has the corresponding value of the given dimension.

**Emotion Categories:** in addition to VAD we also established a list of 26 emotion categories that represent various state of emotions. The list of the 26 emotional categories and their corresponding definitions can be found in Table 1. Also, Fig. 3 shows (per category) examples of people showing different emotional categories.

The list of emotion categories has been created as follows. We manually collected an affective vocabulary from dictionaries and books on psychology [32], [33], [34], [35]. This vocabulary consists of a list of approximately 400 words representing a wide variety of emotional states. After a careful study of the definitions and the similarities amongst these definitions, we formed cluster of words with similar meanings. The clusters were formalized into 26 categories such that they were distinguishable in a single image of a person with her context. We created the final list of 26 emotion categories taking into account the *Visual Separability* criterion: words that have a close meaning could not be visually separable. For instance, *Anger* is defined by

1. <b>Affection:</b> fond feelings; love; tenderness
2. <b>Anger:</b> intense displeasure or rage; furious; resentful
3. <b>Annoyance:</b> bothered by something or someone; irritated; impatient; frustrated
4. <b>Anticipation:</b> state of looking forward; hoping on or getting prepared for possible future events
5. <b>Aversion:</b> feeling disgust, dislike, repulsion; feeling hate
6. <b>Confidence:</b> feeling of being certain; conviction that an outcome will be favorable; encouraged; proud
7. <b>Disapproval:</b> feeling that something is wrong or reprehensible; contempt; hostile
8. <b>Disconnection:</b> feeling not interested in the main event of the surrounding; indifferent; bored; distracted
9. <b>Disquietment:</b> nervous; worried; upset; anxious; tense; pressured; alarmed
10. <b>Doubt/Confusion:</b> difficulty to understand or decide; thinking about different options
11. <b>Embarrassment:</b> feeling ashamed or guilty
12. <b>Engagement:</b> paying attention to something; absorbed into something; curious; interested
13. <b>Esteem:</b> feelings of favourable opinion or judgement; respect; admiration; gratefulness
14. <b>Excitement:</b> feeling enthusiasm; stimulated; energetic
15. <b>Fatigue:</b> weariness; tiredness; sleepy
16. <b>Fear:</b> feeling suspicious or afraid of danger, threat, evil or pain; horror
17. <b>Happiness:</b> feeling delighted; feeling enjoyment or amusement
18. <b>Pain:</b> physical suffering
19. <b>Peace:</b> well being and relaxed; no worry; having positive thoughts or sensations; satisfied
20. <b>Pleasure:</b> feeling of delight in the senses
21. <b>Sadness:</b> feeling unhappy, sorrow, disappointed, or discouraged
22. <b>Sensitivity:</b> feeling of being physically or emotionally wounded; feeling delicate or vulnerable
23. <b>Suffering:</b> psychological or emotional pain; distressed; anguished
24. <b>Surprise:</b> sudden discovery of something unexpected
25. <b>Sympathy:</b> state of sharing others emotions, goals or troubles; supportive; compassionate
26. <b>Yearning:</b> strong desire to have something; jealous; envious; lust

TABLE 1: Proposed emotion categories with definitions.

the words *rage*, *furious* and *resentful*. These affective states are different, but it is not always possible to separate them

**Emotion Category**

“Consider each emotion category separately and, if it is applicable to the person in the given context, select that emotion category”

**Continuous Dimension**

“Consider each emotion dimension separately, observe what level is applicable to the person in the given context, and select that level”

TABLE 2: Instruction summary for each HIT

visually in a single image. Thus, our list of affective categories can be seen as a first level of a hierarchy, where each category has associated subcategories.

Notice that the final list of affective categories also includes the 6 basic emotions (categories 2, 5, 16, 17, 21, 24), but we used the more general term *Aversion* for the category *Disgust*. Thus, the category *Aversion* also includes the subcategories *dislike*, *repulsion*, and *hate* apart from *disgust*.

### 3.2 Collecting Annotations

We used Amazon Mechanical Turk (AMT) crowd-sourcing platform to collect the annotations of the EMOTIC dataset. We designed two Human Intelligence Tasks (HITs), one for each of the 2 formats of emotion representation. The two annotation interfaces are shown in Fig. 5. Each annotator is shown a person-in-context enclosed in a red bounding-box along with the annotation format next to it. Fig. 5.a shows the interface for discrete category annotation while Fig. 5.b displays the interface for continuous dimension annotation. Notice that, in the last box of the continuous dimension interface, we also ask AMT workers to annotate the gender and estimate the age (range) of the person enclosed in red bounding-box. The designing of the annotation interface has two main focuses: *i*) the task is easy to understand and *ii*) the interface fits the HIT in one screen which avoids scrolling.

To make sure annotators understand the task, we showed them how to annotate the images step-wise, by explaining two examples in detail. Also, instructions and examples were attached at the bottom on each page as a quick reference to the annotator. Finally, a summary of the detailed instructions was shown at the top of each page (Table 2).

We adopted two strategies to avoid noisy annotations in the EMOTIC dataset. First, we conduct a qualification task to annotator candidates. This qualification task has two parts: *(i)* an Emotional Quotient HIT (based on standard EQ task [36]) and *(ii)* 2 sample image annotation tasks - one for each of our 2 emotion representations (discrete categories and continuous dimensions). For the sample annotations, we had a set of acceptable labels. The responses of the annotator candidates to this qualification task were evaluated and those who responded satisfactorily were allowed to annotate the images from the EMOTIC dataset. The second strategy to avoid noisy annotations was to insert, randomly, 2 control images in every annotation batch of 20 images; the correct assortment of labels for the control images was known beforehand. Annotators selecting incorrect labels on these control images were not allowed to annotate further and their annotations were discarded.

**(a) Discrete Categories' annotations**

Back (Image 1 of 20) Go to Next Image

- ☐ Peace (well being and relaxed/no worry/positive sensation/satisfied)
- ☐ Affection (fond feelings/tenderness/love/compassion)
- ☐ Expectation (state of anticipating/hoping on something or someone)
- ☐ Esteem (favorable opinion or judgment/gratefulness/admiration/respect)
- ☐ Confidence (feeling of being certain/proud/encouraged/optimistic)
- ☐ Engagement (occupied/absorbed/interested/paying attention to something)
- ☐ Pleasure (feeling of delight in the senses)
- ☐ Happiness (feeling delighted/enjoyment/amusement)
- ☐ Excitement (pleasant and excited state/stimulated/energetic/enthusiastic)
- ☐ Surprise (sudden discovery of something unexpected)
- ☐ Suffering (distressed/perturbed/anguished)
- ☐ Disapproval (think that something is wrong or reprehensible/contempt/hostile)
- ☐ Yearning (strong desire to have something/jealous/envious)
- ☐ Fatigue (weariness/tiredness/sleepy)
- ☐ Pain (physical suffering)
- ☐ Doubt/Confusion (difficulty to understand or decide/sceptical/lost)
- ☐ Fear (feeling afraid of danger/evil/pain/horror)
- ☐ Vulnerability (feeling of being physically or emotionally wounded)
- ☐ Disquietment (unpleasant restlessness/tense/worried/upset/stressed)
- ☐ Annoyance (bothered/irritated/impatient/troubled/frustrated)
- ☐ Anger (intense displeasure or rage/furious/resentful)
- ☐ Disgust (feeling dislike or repulsion/feeling hateful)
- ☐ Sadness (feeling unhappy/grief/disappointed/discouraged)
- ☐ Disconnection (not participating/indifferent/bored/distracted)
- ☐ Embarrassment (feeling ashamed or guilty)

**(b) Continuous Dimensions' annotations**

Back (Image 1 of 20) Go to Next Image

**Valence: Negative vs. Positive**

Negative (unpleasant) [sliders] Positive (pleasant)

**Arousal (awakeness): Calm vs. Ready to act**

Calm [sliders] Ready to act (active)

**Dominance: Dominated vs. In control**

Dominated (no control) [sliders] In control

**Gender and age of the person in the yellow box**

☐ Male ☐ Female

☐ Kid (0-12) ☐ Teenager (13-20) ☐ Adult (more than 20)

Fig. 5: AMT interface designs (a) For Discrete Categories' annotations & (b) For Continuous Dimensions' annotations

### 3.3 Agreement Level Among Different Annotators

Since emotion perception is a subjective task, different people can perceive different emotions after seeing the same image. For example in both Fig. 6.a and 6.b, the person in the red box seems to feel *Affection*, *Happiness* and *Pleasure* and the annotators have annotated with these categories with consistency. However, not everyone has selected all these emotions. Also, we see that annotators do not agree in the emotions *Excitement* and *Engagement*. We consider, however, that these categories are reasonable in this situation. Another example is that of Roger Federer hitting a tennis ball in Fig. 6.c. He is seen predicting the ball (or *Anticipating*) and clearly looks *Engaged* in the activity. He also seems *Confident* in getting the ball.

After these observations we conducted different quantitative analysis on the annotation agreement. We focused first on analyzing the agreement level in the category annotation. Given a category assigned to a person in an image, we consider as an agreement measure the number of annotators agreeing for that particular category. Accordingly, we calculated, for each category and for each annotation in the validation set, the agreement amongst the annotators and sorted those values across categories. Fig. 7 shows the distribution on the percentage of annotators agreeing for an annotated category across the validation set.

We also computed the agreement between all the annotators for a given person using *Fleiss' Kappa* ( $\kappa$ ). Fleiss' Kappa is a common measure to evaluate the agreement level among a fixed number of annotators when assigning categories to data. In our case, given a person to annotate, there is a subset of 26 categories. If we have  $N$  annotators



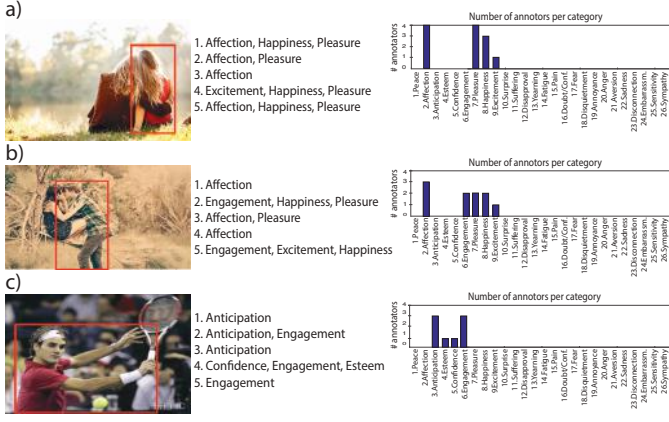


Fig. 6: Annotations of five different annotators for 3 images in EMOTIC.




Fig. 7: Representation of agreement between multiple annotators. Categories sorted in decreasing order according to the average number of annotators who agreed for that category.

per image, that means that each of the 26 categories can be selected by  $n$  annotators, where  $0 \leq n \leq N$ . Given an image we compute the Fleiss' Kappa per each emotion category first, and then the general agreement level on this image is computed as the average of these Fleiss' Kappa values across the different emotion categories. We obtained that more than 50% of the images have  $\kappa > 0.30$ . Fig. 8.a shows the distribution of kappa values across the validation set for all the annotated people in the validation set, sorted in decreasing order. Random annotations or total disagreement produces  $\kappa \sim 0$ , however for our case,  $\kappa \sim 0.3$  (on average) suggesting significant agreement level even though the task of emotion recognition is subjective.

For continuous dimensions, the agreement is measured by the standard deviation (SD) of the different annotations. The average SD across the Validation set is 1.04, 1.57 and 1.84 for Valence, Arousal and Dominance respectively - indicating that Dominance has higher ( $\pm 1.84$ ) dispersion than the other dimensions. It reflects that annotators disagree more often for Dominance than for the other dimensions which is understandable since Dominance is more difficult to interpret than Valence or Arousal [7]. As a summary, Fig. 8.b shows the standard deviations of all the images in the




Fig. 8: (a) Kappa values (sorted) and (b) Standard deviation (sorted), for each annotated person in validation set




Fig. 9: Dataset Statistics. (a) Number of people annotated for each emotion category; (b), (c) & (d) Number of people annotated for every value of the three continuous dimensions viz. Valence, Arousal & Dominance

validation set for all the 3 dimensions, sorted in decreasing order.

### 3.4 Dataset Statistics

EMOTIC dataset contains 34,320 annotated people, where 66% of them are males and 34% of them are females. There are 10% children, 7% teenagers and 83% adults amongst them.

Fig. 9.a shows the number of annotated people for each of the 26 emotion categories, sorted by decreasing order. Notice that the data is unbalanced, which makes the dataset particularly challenging. An interesting observation is that there are more examples for categories associated to positive emotions, like *Happiness* or *Pleasure*, than for categories associated with negative emotions, like *Pain* or *Embarrassment*. The category with most examples is *Engagement*. This is because in most of the images people are doing something or are involved in some activity, showing some degree of engagement. Figs. 9.b, 9.c and 9.d show the number of annotated people for each value of the 3 continuous dimensions. In this case we also observe unbalanced data but fairly distributed across the 3 dimensions which is good for modelling.




Fig. 10: Co-variance between 26 emotion categories. Each row represents the occurrence probability of every other category given the category of that particular row.

Fig. 10 shows the co-occurrence rates of any two categories. Every value in the matrix ( $r, c$ ) ( $r$  represents the row category and  $c$  column category) is a co-occurrence probability (in %) of category  $r$  if the annotation also contains the category  $c$ , that is,  $P(r|c)$ . We observe, for instance, that when a person is labelled with the category *Annoyance*, then there is 46.05% probability that this person is also annotated by the category *Anger*. This means that when a person seems to be feeling *Annoyance* it is likely (by 46.05%) that this person might also be feeling *Anger*. We also used a K-Means clustering on the category annotations to find groups of categories that occur frequently. We found, for example, that these category groups are common in the EMOTIC annotations:  $\{\text{Anticipation, Engagement, Confidence}\}$ ,  $\{\text{Affection, Happiness, Pleasure}\}$ ,  $\{\text{Doubt/Confusion, Disapproval, Annoyance}\}$ ,  $\{\text{Yearning, Annoyance, Disquietment}\}$ .

Fig. 11 shows the distribution of each continuous dimension across the different emotion categories. For every plot, categories are arranged in increasing order of their average values of the given dimension (calculated for all the instances containing that particular category). Thus, we observe from Fig. 11.a that emotion categories like *Suffering*, *Annoyance*, *Pain* correlate with low Valence values (feeling less positive) in average whereas emotion categories like *Pleasure*, *Happiness*, *Affection* correlate with higher Valence values (feeling more positive). Also interesting is to note that a category like *Disconnection* lies in the mid-range of Valence value which makes sense. When we observe Fig. 11.b, it is easy to understand that emotional categories like *Disconnection*, *Fatigue*, *Sadness* show low Arousal values and we see high activeness for emotion categories like *Anticipation*, *Confidence*, *Excitement*. Finally, Fig. 11.c shows that people are not in control when they show emotion categories like *Suffering*, *Pain*, *Sadness* whereas when the Dominance is high, emotion categories like *Esteem*, *Excitement*, *Confidence* occur more often.

An important remark about the EMOTIC dataset is that there are people whose faces are not visible. More than 25% of the people in EMOTIC have their faces partially occluded or with very low resolution, so we can not rely on facial expression analysis for recognizing their emotional state.

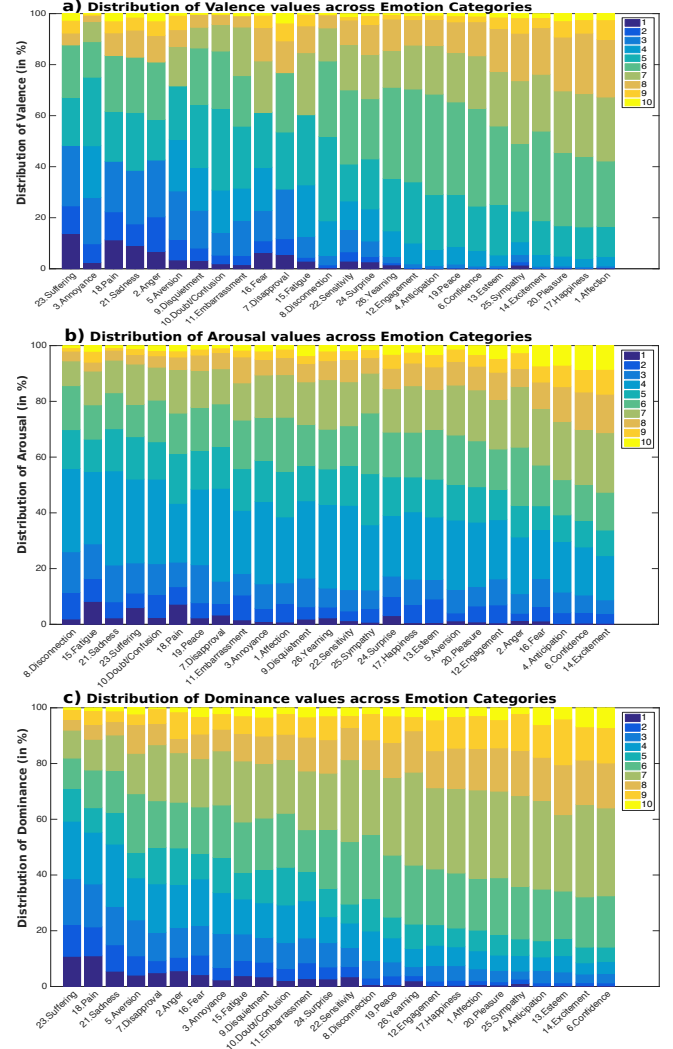


Fig. 11: Distribution of continuous dimension values across emotion categories. Average value of a dimension is calculated for every category and then plotted in increasing order for every distribution.

### 3.5 Algorithmic Scene Context Analysis

This section illustrates how current scene-centric systems can be used to extract contextual information that can be potentially useful for emotion recognition. In particular, we illustrate this idea with a CNN trained on Places dataset [37] and with the Sentibanks Adjective-Noun Pair (ANP) detectors [38], [39], a Visual Sentiment Ontology for Image Sentiment Analysis. As a reference, Fig. 12 shows Places and ANP outputs for sample images of the EMOTIC dataset.

We used AlexNet Places CNN [37] to predict the scene category and scene attributes for the images in EMOTIC. This information helps to divide the analysis into place category and place attribute. We observed that the distribution of emotions varies significantly among different place categories. For example, we found that people in the 'ski\_slope' frequently experience *Anticipation* or *Excitement*, which are associated to the activities that usually happen in this place category. Comparing sport-related and working-environment related images, we find that people in sport-

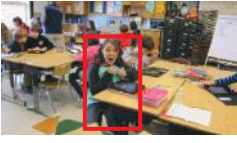

	Places CNN output	Sentibanks NAP (score). Top 8.
	<b>Place Category:</b> kindergarden_classroom, classroom <b>Attributes:</b> no_horizon, enclosed_area, man-made, working, cloth, wood, socializing, plastic, congregating.	early_childhood (0.203) early_education (0.087) early_learning (0.047) elementary_schools (0.045) elementary_education (0.041) creative_kids (0.040) final_exam (0.024) young_child (0.019)
	<b>Place Category:</b> landfill <b>Attributes:</b> natural_light, open_area, dirt, sunny, no_horizon, rugged scene, dry, foliage, trees.	long_range (0.024) outdoor_adventure (0.022) outdoor_education (0.020) hard_work (0.016) healthy_lifestyle (0.007) environmental_portrait (0.007) active_volcano (0.007) big_bear (0.007)

Fig. 12: Illustration of 2 current scene-centric methods for extracting contextual features from the scene: AlexNet Places CNN outputs (place categories and attributes) and Sentibanks ANP outputs for three example images of the EMOTIC dataset.

related images usually show *Excitement*, *Anticipation* and *Confidence*, however they show *Sadness* or *Annoyance* less frequently. Interestingly, *Sadness* and *Annoyance* appear with higher frequency in working environments. We also observe interesting patterns when correlating continuous dimensions with place attributes and categories. For instance, places where people usually show high Dominance are sport-related places and sport-related attributes. On the contrary, low Dominance is shown in ‘jail\_cell’ or attributes like ‘enclosed\_area’ or ‘working’, where the freedom of movement is restricted. In Fig. 12, the predictions by Places CNN describe the scene in general, like in the top image there is a girl sitting in a ‘kindergarden\_classroom’ (places category) which usually is situated in enclosed areas with ‘no\_horizon’ (attributes).

We also find interesting patterns when we compute the correlation between detected ANPs and emotions labelled in the image. For example, in images with people labelled with *Affection*, the most frequent ANP is ‘young\_couple’, while in images with people labelled with *Excitement* we found frequently the ANPs ‘last\_game’ and ‘playing\_field’. Also, we observe a high correlation between images with *Peace* and ANP like ‘old\_couple’ and ‘domestic\_scenes’, and between *Happiness* and the ANPs ‘outdoor\_wedding’, ‘outdoor\_activities’, ‘happy\_family’ or ‘happy\_couple’.

Overall, these observations suggest that some common sense knowledge patterns related with emotions and context could be potentially extracted, automatically, from the data.

#### 4 CNN MODEL FOR EMOTION RECOGNITION IN SCENE CONTEXT

We propose a baseline CNN model for the problem of emotion recognition in context. The pipeline of the model is shown in Fig. 13 and it is divided in three modules: *body feature extraction*, *image (context) feature extraction* and *fusion network*. The first module takes the whole image as input and generates scene-related features. The second module takes the visible body of the person and generates body-related features. Finally, the third module combines these features to do a fine-grained regression of the two types of emotion representations (section 3.1).

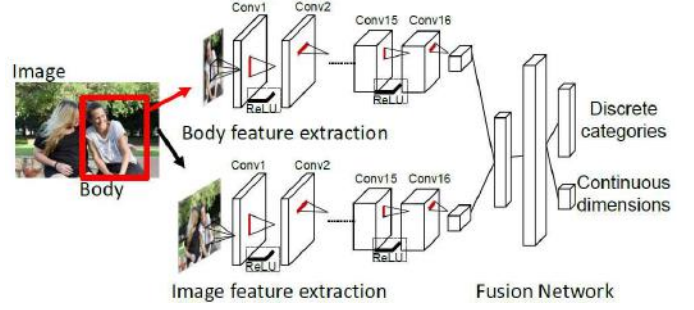


Fig. 13: Proposed end-to-end model for emotion recognition in context. The model consists of two feature extraction modules and a fusion network for jointly estimating the discrete categories and the continuous dimensions.

The body feature extraction module takes the visible part of the body of the target person as input and generates body-related features. These features include important cues like face and head aspects and pose or body appearance. In order to capture these aspects, this module is pre-trained with ImageNet [40], which is an object centric dataset that includes the category *person*.

The image feature extraction module takes the whole image as input and generates scene-context features. These contextual features can be interpreted as an encoding of the scene category, its attributes and objects present in the scene, or the dynamics between other people present in the scene. To capture these aspects, we pre-train this module with the scene-centric Places dataset [37].

The fusion module combines features of the two feature extraction modules and estimates the discrete emotion categories and the continuous emotion dimensions.

Both feature extraction modules are based on the one-dimensional filter CNN proposed in [41]. These CNN networks provide competitive performance while the number of parameters is low. Each network consists of 16 convolutional layers with 1-dimensional kernels alternating between horizontal and vertical orientations, effectively modeling 8 layers using 2-dimensional kernels. Then, to maintain the location of different parts of the image, we use a global average pooling layer to reduce the features of the last convolutional layer. To avoid internal-covariant-shift we add a batch normalizing layer [42] after each convolutional layer and rectifier linear units to speed up the training.

The fusion network module consists of two fully connected (FC) layers. The first FC layer is used to reduce the dimensionality of the features to 256 and then, a second fully connected layer is used to learn independent representations for each task [43]. The output of this second FC layer branches off into 2 separate representations, one with 26 units representing the discrete emotion categories, and second with 3 units representing the 3 continuous dimensions (section 3.1).

##### 4.1 Loss Function and Training Setup

We define the loss function as a weighted combination of two separate losses. A prediction  $\hat{y}$  is composed by the prediction of each of the 26 discrete categories and the 3 continuous dimensions,  $\hat{y} = (\hat{y}^{disc}, \hat{y}^{cont})$ . In particular,



$\hat{y}^{disc} = (\hat{y}_1^{disc}, \dots, \hat{y}_{26}^{disc})$  and  $\hat{y}^{cont} = (\hat{y}_1^{cont}, \hat{y}_2^{cont}, \hat{y}_3^{cont})$ . Given a prediction  $\hat{y}$ , the loss in this prediction is defined by  $L = \lambda_{disc}L_{disc} + \lambda_{cont}L_{cont}$ , where  $L_{disc}$  and  $L_{cont}$  represent the loss corresponding to learning the discrete categories and the continuous dimensions respectively. The parameters  $\lambda_{(disc, cont)}$  weight the contribution of each loss and are set empirically using the validation set.

**Criterion for Discrete categories ( $L_{disc}$ ):** The discrete category estimation is a multilabel problem with an inherent class imbalance issue, as the number of training examples is not the same for each class (see Fig 9.a).

In our experiments, we use a weighted Euclidean loss for the discrete categories. Empirically, we found the Euclidean loss to be more effective than using Kullback–Leibler divergence or a multi-class multi-classification hinge loss. More precisely, given a prediction  $\hat{y}^{disc}$ , the weighted Euclidean loss is defined as follows

$$L_{2disc}(\hat{y}^{disc}) = \sum_{i=1}^{26} w_i (\hat{y}_i^{disc} - y_i^{disc})^2 \quad (1)$$

where  $\hat{y}_i^{disc}$  is the prediction for the  $i$ -th category and  $y_i^{disc}$  is the ground-truth label. The parameter  $w_i$  is the weight assigned to each category. Weight values are defined as  $w_i = \frac{1}{\ln(c+p_i)}$ , where  $p_i$  is the probability of the  $i$ -th category and  $c$  is a parameter to control the range of valid values for  $w_i$ . Using this weighting scheme the values of  $w_i$  are bounded as the number of instances of a category approach to 0. This is particularly relevant in our case as we set the weights based on the occurrence of each category for each batch. Experimentally, we obtained better results using this approach compared to setting the global weights based on the entire dataset.

**Criterion for Continuous dimensions ( $L_{cont}$ ):** We model the estimation of the continuous dimensions as a regression problem. Due to multiple annotators annotating the data based on subjective evaluation, we compare the performance when using two different robust losses: (1) a margin Euclidean loss  $L_{2cont}$ , and (2) the Smooth  $L_1$   $SL_{1cont}$ . The former defines a margin of error ( $v_k$ ) when computing the loss for which the error is not considered. The margin Euclidean loss for continuous dimension is defined as:

$$L_{2cont}(\hat{y}^{cont}) = \sum_{k=1}^3 v_k (\hat{y}_k^{cont} - y_k^{cont})^2, \quad (2)$$

where  $\hat{y}_k^{cont}$  and  $y_k^{cont}$  are the prediction and the ground-truth for the  $k$ -th dimension, respectively, and  $v_k \in \{0, 1\}$  is a binary weight to represent the error margin.  $v_k = 0$  if  $|\hat{y}_k^{cont} - y_k^{cont}| < \theta$ . Otherwise,  $v_k = 1$ . If the predictions are within the error margin, i.e. error is smaller than  $\theta$ , then these predictions do not contribute to update the weights of the network.

The Smooth  $L_1$  loss refers to the absolute error using the squared error if the error is less than a threshold (set to 1 in our experiments). This loss has been widely used for object detection [44] and, in our experiments, has been shown to be less sensitive to outliers. Precisely, the Smooth  $L_1$  loss is defined as follows

$$SL_{1cont}(\hat{y}^{cont}) = \sum_{k=1}^3 v_k \begin{cases} 0.5x^2, & \text{if } |x_k| < 1 \\ |x_k| - 0.5, & \text{otherwise} \end{cases} \quad (3)$$

where  $x_k = (\hat{y}_k^{cont} - y_k^{cont})$ , and  $v_k$  is a weight assigned to each of the continuous dimensions and it is set to 1 in our experiments.

We train our recognition system end-to-end, learning the parameters jointly using stochastic gradient descent with momentum. The first two modules are initialized using pre-trained models from Places [37] and Imagenet [45] while the fusion network is trained from scratch. The batch size is set to 52 - twice the size of the discrete emotion categories. We found empirically after testing multiple batch sizes (including multiples of 26 like 26, 52, 78, 108) that batch-size of 52 gives the best performance (on the validation set).

## 5 EXPERIMENTS

We trained four different instances of our CNN model, which are the combination of two different input types and the two different continuous loss functions described in section 4.1. The input types are body (i.e., upper branch in Fig. 13), denoted by **B**, and body plus image (i.e., both branches shown in Fig. 13), denoted by **B+I**. The continuous loss types are denoted in the experiments by  $L_2$  for Euclidean loss (equation 2) and  $SL_1$  for the Smooth  $L_1$  (equation 3).

Results for discrete categories in the form of Average Precision per category (the higher, the better) are summarized in Table 3. Notice that the **B+I** model outperforms the **B** model in all categories except 1. The combination of body and image features (**B+I**( $SL_1$ ) model) is better than the **B** model.

Results for continuous dimensions in the form of Average Absolute Error per dimension,  $AAE$  (the lower, the better) are summarized in Table 4. In this case, all the models provide similar results where differences are not significant.

Fig. 14 shows the summary of the results obtained per each instance in the testing set. Specifically, Fig. 14.a shows Jaccard coefficient ( $JC$ ) for all the samples in the test set. The  $JC$  coefficient is computed as follows: per each category we use as threshold for the detection of the category the value where  $Precision = Recall$ . Then, the  $JC$  coefficient is computed as the number of categories detected that are also present in the ground truth (number of categories in the intersection of detections and ground truth) divided by the total number of categories that are in the ground truth or detected (union over detected categories and categories in the ground truth). The higher this  $JC$  is the better, with a maximum value of 1, where the detected categories and the ground truth categories are exactly the same. In the graphic, examples are sorted in decreasing order of the  $JC$  coefficient. Notice that these results also support that the **B+I** model outperforms the **B** model.

For the case of continuous dimensions, Fig. 14.b shows the Average Absolute Error ( $AAE$ ) obtained per each sample in the testing set. Samples are sorted by increasing order (best performances on the left). Consistent with the results shown in Table 4, we do not observe a significant difference among the different models.

Finally, Fig. 15 shows qualitative predictions for the best **B** and **B+I** models. These examples were randomly selected among samples with high  $JC$  in **B+I** (a-b) and samples with

Emotion Categories	CNN Inputs and $L_{cont}$ type			
	$B(L_2)$	$B(SL_1)$	$B+I(L_2)$	$B+I(SL_1)$
1. Affection	21.80	16.55	21.16	27.85
2. Anger	06.45	04.67	06.45	09.49
3. Annoyance	07.82	05.54	11.18	14.06
4. Anticipation	58.61	56.61	58.61	58.64
5. Aversion	05.08	03.64	06.45	07.48
6. Confidence	73.79	72.57	77.97	78.35
7. Disapproval	07.63	05.50	11.00	14.97
8. Disconnection	20.78	16.12	20.37	21.32
9. Disquietment	14.32	13.99	15.54	16.89
10. Doubt/Confusion	29.19	28.35	28.15	29.63
11. Embarrassment	02.38	02.15	02.44	03.18
12. Engagement	84.00	84.59	86.24	87.53
13. Esteem	18.36	19.48	17.35	17.73
14. Excitement	73.73	71.80	76.96	77.16
15. Fatigue	07.85	06.55	08.87	09.70
16. Fear	12.85	12.94	12.34	14.14
17. Happiness	58.71	51.56	60.69	58.26
18. Pain	03.65	02.71	04.42	08.94
19. Peace	17.85	17.09	19.43	21.56
20. Pleasure	42.58	40.98	42.12	45.46
21. Sadness	08.13	06.19	10.36	19.66
22. Sensitivity	04.23	03.60	04.82	09.28
23. Suffering	04.90	04.38	07.65	18.84
24. Surprise	17.20	17.03	16.42	18.81
25. Sympathy	10.66	09.35	11.44	14.71
26. Yearning	07.82	07.40	08.34	08.34
Mean	23.86	22.36	24.88	27.38

TABLE 3: Average Precision (AP) obtained on test set per category. Results for models where the input is just the body  $B$ , and models where the input are both the body and the whole image  $B+I$ . The type of  $L_{cont}$  used is indicated in parenthesis ( $L_2$  refers to equation 2 and  $SL_1$  refers to equation 3).

Continuous Dimensions	CNN Inputs and $L_{cont}$ type			
	$B(L_2)$	$B(SL_1)$	$B+I(L_2)$	$B+I(SL_1)$
Valence	0.0537	0.0545	0.0546	0.0528
Arousal	0.0600	0.0630	0.0648	0.0611
Dominance	0.0570	0.0567	0.0573	0.0579
Mean	0.0569	0.0581	0.0589	0.0573

TABLE 4: Average Absolute Error (AAE) obtained on test set per each continuous dimension. Results for models where the input is just the body  $B$ , and models where the input are both the body and the whole image  $B+I$ . The type of  $L_{cont}$  used is indicated in parenthesis ( $L_2$  refers to equation 2 and  $SL_1$  refers to equation 3).

low  $JC$  in  $B+I$  (g-h). Incorrect category recognition is indicated in red. As shown, in general,  $B+I$  model outperforms  $B$ , although there are some exceptions, like Fig. 15.c.

### 5.1 Context Features Comparison

The goal of this section is to compare different context features for the problem of emotion recognition in context. A key aspect for incorporating the context in an emotion recognition model is to be able to obtain information from the context that is actually relevant for emotion recognition. Since the context information extraction is a scene-centric task, the information extracted from the context should be based in a scene-centric feature extraction system. That is why our baseline model uses a Places CNN for the context feature extraction module. However, recent works in sentiment analysis (detecting the emotion of a person when




Fig. 14: Results per each sample (Test Set, sorted): (a) Jaccard Coefficient ( $JC$ ) of the recognized discrete categories (b) Average Absolute Error ( $AAE$ ) in the estimation of the three continuous dimensions.

he/she observes an image) also provide a system for scene feature extraction that can be used for encoding the relevant contextual information for emotion recognition.

To compute body features, denoted by  $B_f$ , we fine tune an AlexNet ImageNet CNN with EMOTIC database, and use the average pooling of the last convolutional layer as features. For the context (image), we compare two different feature types, which are denoted by  $I_f$  and  $I_S$ .  $I_f$  are obtained by fine tuning an AlexNet Places CNN with EMOTIC database, and taking the average pooling of the last convolutional layer as features (similar to  $B_f$ ), while  $I_S$  is a feature vector composed of the sentiment scores for the ANP detectors from the implementation of [39].

To fairly compare the contribution of the different context features, we train Logistic Regressors for the following features and combination of features: (1)  $B_f$ , (2)  $B_f+I_f$ , and (3)  $B_f+I_S$ . For the discrete categories we obtain mean APs  $AP = 23.00$ ,  $AP = 27.70$ , and  $AP = 29.45$ , respectively. For the continuous dimensions, we obtain AAE 0.0704, 0.0643, and 0.0713 respectively. We observe that, for the discrete categories, both  $I_f$  and  $I_S$  contribute relevant information to the emotion recognition in context. Interestingly,  $I_S$  performs better than  $I_f$ , even though these features have not been trained using EMOTIC. However, these features are smartly designed for sentiment analysis, which is a problem closely related to extracting relevant contextual information for emotion recognition, and are trained with a large dataset of images.

## 6 CONCLUSIONS

In this paper we pointed out the importance of considering the person scene context in the problem of automatic emotion recognition in the wild. We presented the EMOTIC database, a dataset of 23,571 natural unconstrained images with 34,320 people labeled according to their apparent emotions. The images in the dataset are annotated using two different emotion representations: 26 discrete categories, and the 3 continuous dimensions *Valence*, *Arousal* and *Dominance*. We described in depth the annotation process and analyzed the annotation consistency of different annotators. We also provided different statistics and algorithmic analysis on the data, showing the characteristics of the EMOTIC database. In addition, we proposed a baseline

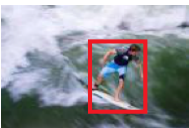







	Ground Truth	B (L2)		B+H (SL1)	
a)		Anticipation V: 0,57 Confidence A: 0,83 Engagement D: 0,67	Anticipation V: 0,61 Confidence A: 0,61 Engagement D: 0,67	Anticipation V: 0,62 Confidence A: 0,70 Engagement D: 0,66	
			Happiness Surprise Sympathy JC: 0,57		JC: 1,00
b)		Anticipation V: 0,50 Confidence A: 0,63 Engagement D: 0,67	Anticipation V: 0,50 Confidence A: 0,54 Engagement D: 0,64	Anticipation V: 0,62 Confidence A: 0,56 Engagement D: 0,61	
			Esteem Excitement Happiness Peace JC: 0,71		JC: 1,00
c)		Anticipation V: 0,60 Confidence A: 0,33 Engagement D: 0,63	Anticipation V: 0,59 Confidence A: 0,52 Engagement D: 0,61	Anticipation V: 0,63 Confidence A: 0,56 Engagement D: 0,63	
			Disquietment Doubt/Confusion Excitement Happiness, Surprise JC: 0,75		JC: 0,71
d)		Anticipation V: 0,60 Confidence A: 0,53 Excitement D: 0,73	Anticipation V: 0,59 Aversion A: 0,52 Engagement D: 0,58	Anticipation V: 0,64 Confidence A: 0,54 Excitement D: 0,62	
			Engagement Happiness Peace JC: 0,38		JC: 0,71
e)		Affection V: 0,67 Anticipation A: 0,43 Engagement D: 0,83	Anticipation V: 0,59 Confidence A: 0,52 Doubt/Confusion D: 0,61	Anticipation V: 0,62 Confidence A: 0,52 Disconnection D: 0,62	
			Engagement Esteem Happiness Pain Pleasure JC: 0,30		JC: 0,40
f)		Affection V: 0,53 Anticipation A: 0,70 Disquietment D: 0,63	Anticipation V: 0,60 Engagement A: 0,56 Doubt/Confusion D: 0,63	Affection V: 0,62 Anticipation A: 0,58 Confidence D: 0,63	
			Happiness Peace Pleasure JC: 0,22		JC: 0,40
g)		Annoyance V: 0,40 Engagement A: 0,33 Excitement D: 0,63	Affection, Anticipation, Aversion V: 0,59 Confidence A: 0,52 Disapproval D: 0,61	Affection V: 0,62 Anticipation A: 0,50 Disconnection D: 0,59	
			Doubt/Confusion Disapproval Doubt/Confusion Embarrassment, Engagement Esteem, Fatigue, Happiness, Peace, Pleasure, Sympathy JC: 0,17		JC: 0,23
h)		Anger V: 0,50 Annoyance A: 0,33 Aversion D: 0,67	Anticipation V: 0,60 Confidence A: 0,50 Disconnection D: 0,63	Affection V: 0,64 Anticipation A: 0,54 Disquietment D: 0,62	
			Doubt/Confusion Engagement Happiness Pain JC: 0,00		JC: 0,08

Fig. 15: Ground truth and results on images randomly selected with different  $JC$  scores.

CNN model for emotion recognition in scene context that combines the information of the person (body bounding box) with the scene context information (whole image). We also compare two different feature types for encoding the contextual information. Our results show the relevance of using contextual information to recognize emotions and, in conjunction with the EMOTIC dataset, motivate further research in this direction. All the data and trained models are publicly available for the research community in the website of the project.

## ACKNOWLEDGMENT

This work has been partially supported by the *Ministerio de Economía, Industria y Competitividad (Spain)*, under the Grants Ref. TIN2015-66951-C2-2-R and RTI2018-095232-B-C22, and by Innovation and Universities (FEDER funds). The authors also thank NVIDIA for their generous hardware donations.

## REFERENCES

- [1] D. Borth, R. Ji, T. Chen, T. Breuel, and S.-F. Chang, "Large-scale visual sentiment ontology and detectors using adjective noun pairs," in *Proceedings of the 21st ACM international conference on Multimedia*. ACM, 2013, pp. 223–232.
- [2] H. Aviezer, R. R. Hassin, J. Ryan, C. Grady, J. Susskind, A. Anderson, M. Moscovitch, and S. Bentin, "Angry, disgusted, or afraid? studies on the malleability of emotion perception," *Psychological science*, vol. 19, no. 7, pp. 724–732, 2008.
- [3] R. Righart and B. De Gelder, "Rapid influence of emotional scenes on encoding of facial expressions: an erp study," *Social cognitive and affective neuroscience*, vol. 3, no. 3, pp. 270–278, 2008.
- [4] T. Masuda, P. C. Ellsworth, B. Mesquita, J. Leu, S. Tanida, and E. Van de Veerdonk, "Placing the face in context: cultural differences in the perception of facial emotion," *Journal of personality and social psychology*, vol. 94, no. 3, p. 365, 2008.
- [5] L. F. Barrett, B. Mesquita, and M. Gendron, "Context in emotion perception," *Current Directions in Psychological Science*, vol. 20, no. 5, pp. 286–290, 2011.
- [6] L. F. Barrett, *How emotions are made: The secret life of the brain*. Houghton Mifflin Harcourt, 2017.
- [7] A. Mehrabian, "Framework for a comprehensive description and measurement of emotional states," *Genetic, social, and general psychology monographs*, 1995.
- [8] R. Kosti, J. M. Alvarez, A. Recasens, and A. Lapedriza, "Emotion recognition in context," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [9] M. Pantic and L. J. Rothkrantz, "Expert system for automatic analysis of facial expressions," *Image and Vision Computing*, vol. 18, no. 11, pp. 881–905, 2000.
- [10] Z. Li, J.-i. Imai, and M. Kaneko, "Facial-component-based bag of words and phog descriptor for facial expression recognition," in *SMC*, 2009, pp. 1353–1358.
- [11] E. Friesen and P. Ekman, "Facial action coding system: a technique for the measurement of facial movement," *Palo Alto*, 1978.
- [12] P. Ekman and W. V. Friesen, "Constants across cultures in the face and emotion," *Journal of personality and social psychology*, vol. 17, no. 2, p. 124, 1971.
- [13] C. F. Benitez-Quiroz, R. Srinivasan, and A. M. Martinez, "Emotionet: An accurate, real-time algorithm for the automatic annotation of a million facial expressions in the wild," in *Proceedings of IEEE International Conference on Computer Vision & Pattern Recognition (CVPR16)*, Las Vegas, NV, USA, 2016.
- [14] M. Soleymani, S. Asghari-Esfeden, Y. Fu, and M. Pantic, "Analysis of eeg signals and facial expressions for continuous emotion detection," *IEEE Transactions on Affective Computing*, vol. 7, no. 1, pp. 17–28, 2016.
- [15] S. Du, Y. Tao, and A. M. Martinez, "Compound facial expressions of emotion," *Proceedings of the National Academy of Sciences*, vol. 111, no. 15, pp. E1454–E1462, 2014.
- [16] M. A. Nicolaou, H. Gunes, and M. Pantic, "Continuous prediction of spontaneous affect from multiple cues and modalities in valence-arousal space," *IEEE Transactions on Affective Computing*, vol. 2, no. 2, pp. 92–105, 2011.
- [17] K. Schindler, L. Van Gool, and B. de Gelder, "Recognizing emotions expressed by body pose: A biologically inspired neural model," *Neural networks*, vol. 21, no. 9, pp. 1238–1246, 2008.
- [18] W. Mou, O. Celiktutan, and H. Gunes, "Group-level arousal and valence recognition in static images: Face, body and context," in *Automatic Face and Gesture Recognition (FG)*, 2015 11th IEEE International Conference and Workshops on, vol. 5. IEEE, 2015, pp. 1–6.
- [19] "GENKI database," [http://mplab.ucsd.edu/wordpress/?page\\_id=398](http://mplab.ucsd.edu/wordpress/?page_id=398), accessed: 2017-04-12.
- [20] "ICML face expression recognition dataset," <https://goo.gl/nn9w4R>, accessed: 2017-04-12.
- [21] J. L. Tracy, R. W. Robins, and R. A. Schriber, "Development of a facs-verified set of basic and self-conscious emotion expressions," *Emotion*, vol. 9, no. 4, p. 554, 2009.
- [22] A. Kleinsmith and N. Bianchi-Berthouze, "Recognizing affective dimensions from body posture," in *Proceedings of the 2Nd International Conference on Affective Computing and Intelligent Interaction*, ser. ACII '07. Berlin, Heidelberg:



- Springer-Verlag, 2007, pp. 48–58. [Online]. Available: [http://dx.doi.org/10.1007/978-3-540-74889-2\\_5](http://dx.doi.org/10.1007/978-3-540-74889-2_5)
- [23] A. Kleinsmith, N. Bianchi-Berthouze, and A. Steed, “Automatic recognition of non-acted affective postures,” *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 41, no. 4, pp. 1027–1038, Aug 2011.
- [24] T. Bänziger, H. Pirker, and K. Scherer, “Gemep-geneva multimodal emotion portrayals: A corpus for the study of multimodal emotional expressions,” in *Proceedings of LREC*, vol. 6, 2006, pp. 15–019.
- [25] S. Escalera, X. Baró, H. J. Escalante, and I. Guyon, “Chalearn looking at people: Events and resources,” *CoRR*, vol. abs/1701.02664, 2017. [Online]. Available: <http://arxiv.org/abs/1701.02664>
- [26] A. Dhall, R. Goecke, J. Joshi, J. Hoey, and T. Gedeon, “EmotiW 2016: Video and group-level emotion recognition challenges,” in *Proceedings of the 18th ACM International Conference on Multimodal Interaction*, ser. ICMI 2016. New York, NY, USA: ACM, 2016, pp. 427–432. [Online]. Available: <http://doi.acm.org/10.1145/2993148.2997638>
- [27] A. Dhall *et al.*, “Collecting large, richly annotated facial-expression databases from movies,” 2012.
- [28] A. Dhall, J. Joshi, I. Radwan, and R. Goecke, “Finding happiest moments in a social context,” in *Asian Conference on Computer Vision*. Springer, 2012, pp. 613–626.
- [29] G. Patterson and J. Hays, “Coco attributes: Attributes for people, animals, and objects,” in *European Conference on Computer Vision*. Springer, 2016, pp. 85–100.
- [30] T. Lin, M. Maire, S. J. Belongie, L. D. Bourdev, R. B. Girshick, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, “Microsoft COCO: common objects in context,” *CoRR*, vol. abs/1405.0312, 2014. [Online]. Available: <http://arxiv.org/abs/1405.0312>
- [31] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba, “Semantic understanding of scenes through ade20k dataset,” 2016. [Online]. Available: <https://arxiv.org/pdf/1608.05442>
- [32] “Oxford english dictionary,” <http://http://www.oed.com>, accessed: 2017-06-09.
- [33] “Merriam-webster online english dictionary,” <https://www.merriam-webster.com>, accessed: 2017-06-09.
- [34] E. G. Fernández-Abascal, B. García, M. Jiménez, M. Martín, and F. Domínguez, *Psicología de la emoción*. Editorial Universitaria Ramón Areces, 2010.
- [35] R. W. Picard and R. Picard, *Affective computing*. MIT press Cambridge, 1997, vol. 252.
- [36] Y. Groen, A. B. M. Fuermaier, A. E. Den Heijer, O. Tucha, and M. Althaus, “The empathy and systemizing quotient: The psychometric properties of the dutch version and a review of the cross-cultural stability,” *Journal of Autism and Developmental Disorders*, vol. 45, no. 9, pp. 2848–2864, 2015. [Online]. Available: <http://dx.doi.org/10.1007/s10803-015-2448-z>
- [37] B. Zhou, A. Khosla, A. Lapedriza, A. Torralba, and A. Oliva, “Places: An image database for deep scene understanding,” *CoRR*, vol. abs/1610.02055, 2015.
- [38] D. Borth, R. Ji, T. Chen, T. Breuel, and S.-F. Chang, “Large-scale visual sentiment ontology and detectors using adjective noun pairs,” in *Proceedings of the 21st ACM international conference on Multimedia*. ACM, 2013, pp. 223–232.
- [39] T. Chen, D. Borth, T. Darrell, and S.-F. Chang, “Deepsentibank: Visual sentiment concept classification with deep convolutional neural networks,” *arXiv preprint arXiv:1410.8586*, 2014.
- [40] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [41] J. Alvarez and L. Petersson, “Decomposeme: Simplifying convnets for end-to-end learning,” *CoRR*, vol. abs/1606.05426, 2016.
- [42] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International Conference on Machine Learning*, 2015, pp. 448–456.
- [43] R. Caruana, *A Dozen Tricks with Multitask Learning*, 2012, pp. 163–189.
- [44] R. Girshick, “Fast r-cnn,” in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 1440–1448.
- [45] J. Deng, W. Dong, R. Socher, L. Li, K. Li, and L. Fei-Fei, “Imagenet: A large-scale hierarchical image database,” in *CVPR*, 2009.




**Ronak Kosti** is pursuing his PhD at Universitat Oberta de Catalunya, Spain advised by Prof. Agata Lapedriza. He is working with Scene Understanding and Artificial Intelligence (SUNAI) group in computer vision, specifically in the area of affective computing. He obtained his Masters in Machine Intelligence from DA-IICT (Dhirubhai Ambani Institute of Information and Communication Technology) in 2014. His master's research was based on depth estimation from single image using Artificial Neural Networks.



supervision of Prof. Yann LeCun.

**Jose M. Alvarez** is a senior research scientist at NVIDIA. Previously, he was senior deep learning researcher at Toyota Research Institute, US; prior to that he was with Data61, CSIRO, Australia (formerly NICTA) as a researcher. He obtained his Ph.D. in 2010 from Autonomous University of Barcelona under the supervision of Prof. Antonio Lopez and Prof. Theo Gevers. Previous to CSIRO, he worked as a postdoctoral researcher at the Courant Institute of Mathematical Science at New York University under the



**Adria Recasens** is pursuing his PhD in computer vision at the Computer Science and Artificial Intelligence Laboratory (CSAIL) of the Massachusetts Institute of Technology advised by Professor Antonio Torralba. His research interests range on various topics in computer vision and machine learning. He is focusing most of his research on automatic gaze-following. He received a Telecommunications Engineer's Degree and a Mathematics Licentiate Degree from the Universitat Politècnica de Catalunya.



**Agata Lapedriza** is an Associate Professor at the Universitat Oberta de Catalunya. She received her MS degree in Mathematics at the Universitat de Barcelona and her Ph.D. degree in Computer Science at the Computer Vision Center, at the Universitat Autònoma de Barcelona. She was working as a visiting researcher in the Computer Science and Artificial Intelligence Lab, at the Massachusetts Institute of Technology (MIT), from 2012 until 2015. Currently she is a visiting researcher at the MIT Medialab, Affective

Computing group from September 2017. Her research interests are related to image understanding, scene recognition and characterization, and affective computing.