

Face expression understanding by geometrical characterization of deep human face representation

Adrien RAISON^{1,2}, Théo BIARDEAU¹, Pascal BOURDON^{1,2}, David HELBERT^{1,2}

¹University of Poitiers, XLIM Laboratory, ASALI department, CNRS 7252

²I3M Common Laboratory CNRS-Siemens, University and Hospital of Poitiers
Poitiers, France

Abstract

Emotion understanding is a key to have a better understanding of the human nature. In this contribution we propose an end-to-end pipeline that takes color images as inputs and produces a semantic graph that encodes numerically what are facial emotions. This approach leverages low-level geometric details as face representation which are numerical representations of facial muscle activation patterns to build this emotional understanding. It shows that our method recovers social expectations of what characterizes facial emotions.

Introduction

Emotions are one of the main human communication channels. They are resultant of many complex human internal processes that themselves have been triggered by external stimuli. Understanding how human emotions work is a continuously investigated research topic and results favor both academic (e.g human psychological understanding improvement) and industrial (e.g advertising targeting) developments. In this study we propose a new framework allowing understanding visually facial emotions; in which manner they are recognizable from others; in other words, what is defining them intrinsically from a numerical point of view. Our approach relies on embedding human faces as graphs where nodes represent both semantic and positional landmarks. The spatial distribution of those landmarks is highly related to facial emotion revealing procedures. We firstly classify those graphs with a semantic-aware Graph Neural Network (GNN) model. We then provide a qualitative study that explains accurate internal decision processes the classifier has designed for it, which in this context, is a proxy for emotions numerical definition framework.

Related work

Emotion characterization is foremost an emotion recognition dependent problem. Emotion recognition problems have been deeply investigated according to different angles such as natural language processing paradigm [7]. Other studies have been conducted with EEG signals [2] that allow us to have a deep notion of what is emotion from a neurological point of view. Other studies had taken interest in human emotions characterization namely [8, 9] and a graph-based study has been done by [4]. Understanding facial emotions from a visual viewpoint, on a side, deals with facial alignment. Facial alignment is at the core interest of computer vision community. It consists of specifying facial landmarks position regarding specific semantics. Some relevant solutions have been proposed by the computer vision community

leading to a wide deployment in industrial applications although it remains an important research topic still investigated. There are non-linear correlated phenomena between facial landmarks groundtruth distribution and facial emotion revealing process. In this study we propose an end-to-end framework that first, encodes human faces present in color images as graphs; then explain deep face representation classification yielding numerical definition of what are human facial emotions. Many explaining methods suited for GNN have been proposed in the literature. XGNN [21] is a model-level approach that generates iteratively explaining graphs through a reinforcement learning procedure. GNNExplainer [20] is a mask generator model based on mutual information optimization. It starts with randomly initialized node and node features mask jointly optimized, with mutual information, against the class label of the assessed graph. LRP-GNN [16] adopts a walk-based approach, introducing the node anteriority and apply the original Layer-wise Relevance Propagation (LRP) [1] propagation rule. EiX-GNN [10] is a model-agnostic approach that is ordering sub-graphs by determining the asymptotic behavior of a random walk that is performed over the explained graphs. Nonetheless, explaining is often a context-dependent task. In order to get rid of any contextualized expertise to assess these methods in terms of relevance, objectives metrics [19] has been designed and have helped us to choose our explaining methods. Despite all these considerations, it turns out that our pipeline allows recovering social expectations of what are facial emotion recognition.

Problem formulation

Emotion is a human multi-factorial response to environmental interaction. For human-to-human communication, emotions help to design exchanged messages. Expressing emotions can be made by different channels. One of them is to put face configuration in a specific spatial configuration that is socially recognizable. From this, revealing an emotion is a transfer process between an initial face shape to another specific one. The resulting pose can be defined as being what is the expressed emotion. For instance in 2-person conversation, expressing happiness is done most of the time by smiling and uprising cheeks of one of the involved people. Noticing this change helps to visually perceive happiness for the conversation partner, the emotional message is thus received if both individuals share the same facial emotion revealing process. Under computer-aided approach, understanding and characterizing human face emotions can be addressed by computationally understand the underlying conditioned facial geometry and especially given these precise configurations. We present here our proposed approach providing relevant results

even on a wide range of social contexts.

Our approach

Face shape deformation is due to a joint facial muscle activation, some patterns occur and some of them are basis components for facial emotion revealing. Catching such patterns and characterizing them computationally are the key to provide a computational understanding of facial emotions. Encoding such patterns require to define them physically and transcript it in computing machines. Faces are complex biological structure but we can represent facial muscle groups as semantic positional landmarks. To cover this issue, we use state-of-art face alignment method that has been especially designed for such semantic mapping. The landmark spatiality variation is due to the adjacent facial muscles that squeeze or stretch conjointly in some face neighborhoods. Acquiring this relational knowledge, in particular when emotions are expressed, can be done by representing faces as a graph where each node is a facial landmark and where adjacency is represented by both intra-muscles (local) and inter muscles (global) scale interactions. We will describe our image-based facial graph-embedding process suited for various contexts; introducing our emotional supervised classification problem and its interpretability aspect. This last will let us highlight what are key insights to understand facial emotions at a visual scope.

Feature conception and labeling

These raw materials are the inputs of our preprocessing step : image-based facial landmark detection and emotion recognition. For the sake of clarity, we denote D as being our images dataset. The dataset set size is denoted by $|D| \in \mathbb{N}$ and we see an image of size $k, l \in \mathbb{N}$ as a triplet $\mathbf{X} \in \mathbb{M}_{k,l}(\{0, 255\} \cap \mathbb{N})^3$.

Facial landmark inference From raw color images we determine graph features to learn thank to a face alignment method. The development of such method is at a core interest of computer vision community and detecting facial landmarks is a well-studied problem. Basically, assuming that we have an optimal image-based facial landmark detector \mathbf{fa}^* that determined carefully $p \in \mathbb{N}$ three-dimensional facial landmarks, i.e :

$$\begin{aligned} \mathbf{fa}^* : D &\rightarrow \mathbb{R}^{3 \times p} \\ \mathbf{X} &\mapsto (\mathbf{x}_i)_{i \in \{1, \dots, p\}} \end{aligned}$$

Since each instance \mathbf{X} of D has its own recording context, the inference $\mathbf{fa}^*(\mathbf{X})$ lean on significant different spatially distributed subspace of \mathbb{R}^3 notably in terms of space-barycentric position or overall scaling ratio. Without a standardization procedure, this will lead to an irrelevant statistical learning because intrinsecal emotion signal is only a in between relative landmarks interaction. To gain statistical stability, we apply a standardization process relying on finding optimal rigid affine transformations. We will consider here the scale, rotation and translation transformation since this is these three affine transformations that overcome the wide-context image recording issue. This standardization puts $\mathbf{fa}^*(D)$ in a common, sustainable and statistically efficient dataset representation. This framework is defined regarding a mean face that stands as a regressed and standardized objective. This mean face is emotionless, synthetic, and has been designed by averaging

spatially p landmarks of neutral aligned faces. It also means that neutral emotion will be our baseline to characterize other facial emotions. We denote this mean face has $(\mathbf{m}_i)_{i \in \{1, \dots, p\}} \in \mathbb{R}^{3 \times p}$. The standardization process is formalized as follows; for all $(\mathbf{x}_i)_i \in \mathbf{fa}^*(D)$:

$$(s_{\mathbf{X}}^*, \mathbf{R}_{\mathbf{X}}^*, \mathbf{t}_{\mathbf{X}}^*) = \arg \min_{(s, \mathbf{R}, \mathbf{t}^T) \in \Omega} \sum_{i=1}^p \|\mathbf{m}_i^T - s\mathbf{R}\mathbf{x}_i^T - \mathbf{t}^T\|_2 \quad (1)$$

with $\Omega = \mathbb{R}^+ \times SO(3) \times \mathbb{R}^3$.

Note that those rigid transformations follow the data distribution prior. Indeed, due to the various contexts present in D , head poses are various as well. Considered geometric transformations must not alter the global biological structure of human faces which is, without abusive considerations, assumed to be constant across any individuals. And since affine transformations preserve notably collinearity, parallelism, ratio of length and barycenters, it conserves the global physical faces aspect and does not introduce any semantic alterations or non-linear effects. This problem (1) has a closed-form solution where proof is given in [5]. Due to face global face shape variation over individuals in D , non-linear variation (i.e individual singularity) is not taken into account but can be seen as noisy features that may not affect meaningfully the learning phase. We denote the standardized dataset as D^* , it means:

$$D^* = \{s_{\mathbf{X}}^* \mathbf{R}_{\mathbf{X}}^* \mathbf{X}^T - \mathbf{t}_{\mathbf{X}}^{*T} | \mathbf{X} \in \mathbf{fa}^*(D)\} \quad (2)$$

Then we embed this new data standardized representation as a graph according to the following scheme. From original definition, graphs are couple (V, E) where V is the set of nodes and $E \subset V \times V$ is the set of edges. The set E is the adjacency representation of the considered graph. There is a matrix formulation of edges representation that is fully equivalent and allowing performing easier algebraic operation. The adjacency matrix $\mathbf{A} \in \mathbb{M}_{|V|}(\{0, 1\})$ as its entry $a_{i,j}$ defined by $a_{i,j} = 1$ iff $(i, j) \in E$, 0 otherwise. In the context of deep representation of signals evolving over the graph, graphs are rather seen as (\mathbf{L}, \mathbf{A}) such that \mathbf{L} is a column vector of size $|V|$ valued in \mathcal{H} , a d -dimensional Hilbert space ($d \in \mathbb{N}$). It thus means that each node in V as an \mathcal{H} -valued row vector of size d representing the signal evolving on this node. The adjacency matrix \mathbf{A} is used to describe the domain structure.

We map bijectively D^* to $D_G^* = \{G_{\hat{\mathbf{X}}} | \hat{\mathbf{X}} \in D^*\}$ where for each $\hat{\mathbf{X}} \in D^*$, $G_{\hat{\mathbf{X}}} = (\hat{\mathbf{X}}, \mathbf{A}) \in D^* \times \{\mathbf{A}\}$. Note that \mathbf{A} does not depend on $\hat{\mathbf{X}}$ because the topology of induced graphs (i.e, the face muscles interaction adjacency) is a common feature shared across all humans being, indeed as mentioned before giving a social context and an emotion, peoples express it according to the same joint muscle activation patterns. The design of \mathbf{A} has been driven by the muscle momentum connexity assumption. The muscle momentum connexity assumption is based upon the fact that in every face localities, physically connected muscles act together as a group (i.e given a muscle, when it is activating, adjacent muscles are more likely to be driven by this activation and acting in turn). And so, we assume that this assumption holds to reveal facial emotion.

We designed facial muscles interactivity according to a hand-crafted strategy built upon local and global semantic information. Local adjacency is conceived according to the muscle momentum

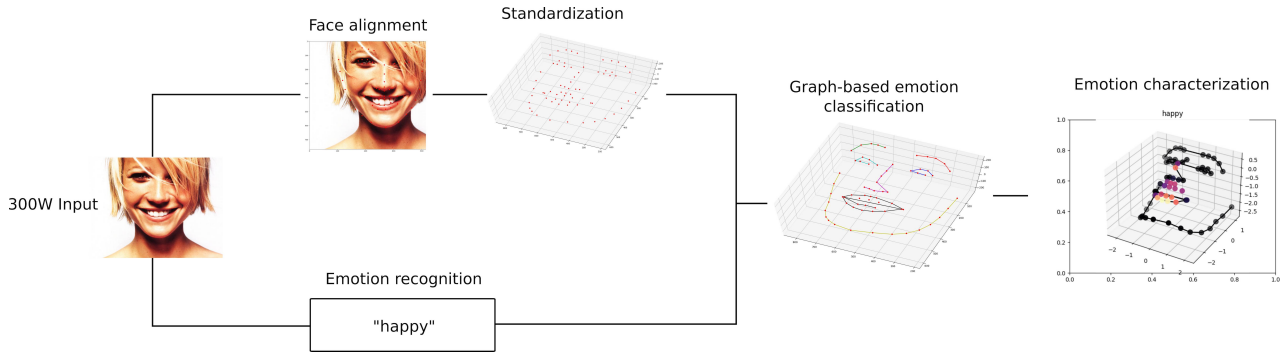


Figure 1. Our method relies on three compounds: a labeling module that is used to build our graph-fashion face emotional datasets; once this assembling is done, a deep model will classify our graph by parcellating in human-intractable high-dimensional space standardized graph vertices spatial positions; a conversion is then operated to put this complex representation into a human readable that provide an understanding of human facial emotions thanks to the ScoreCAM GNN explainer.

Semantic features	Node index
Left eyebrow	17,18,19,20,21
Right eyebrow	22,23,24,25,26
Left eye	36,37,38,39,40,41
Right eye	42,43,44,45,46,47
Nose	27,28,29,30,31,32,33,34,35
Thin & cheeks	0,1,2,3,4,5,6,7,8,9,10,11,...,16
Upper mouth	48,49,50,51,52,53,54,64
Lower mouth	48,60,59,58,57,56,55,54,64,65,66,67

Semantic face parts - Nodes index mapping

connexity assumption. At the local scope, landmarks, that represent the same positional semantic, are represented as a chained subgraph. Global adjacency is focused on gathering that local semantics in an prior-free representation, so as a complete adjacency structure according to the graph theory terminology. This strategy allows embedding efficiently both fine-details semantics (local scale) and higher-level semantics (global scale), including biological prior knowledge. The careful mixing of these connex components (according to the input graph) has a convenient algebraic representation. Indeed, the general adjacency matrix \mathbf{A} is a block matrix designed respectively to respect the aforementioned $\{1, \dots, p\}$ semantic parcellation. Landmarks semantic splitting mapping is described in Table 1.

Emotion recognizer inference As well as facial landmarks detection method, emotion recognition based on color images has been heavily investigated by the computer vision community. Many studies have been conducted and emotion recognition problems are often framed as supervised classification problems since nowadays many data are publicly available. Considering $C \in \mathbb{N}$ different emotions, an emotion recognizer is simply an optimal mapping er^* such that:

$$\begin{aligned} \text{er}^* : D &\rightarrow \{1, \dots, C\} \\ \mathbf{X} &\mapsto c_{\mathbf{X}} \end{aligned}$$

We then infer $\text{er}^*(D)$ to obtain labels that will be afterward used to classify emotions with respect to $D_G^* \times \text{er}^*(D)$, i.e. in a graph-based fashion rather than in an image-based setting.

Note: Inferences made through fa^* and er^* may not be

absolutely accurate in terms of, respectively, emotion and facial landmark groundtruths. The $D_G^* \times \text{er}^*(D)$ quality is highly dependent on fa^* and er^* accuracy. It has to be noticed that the wrong assignment may lead to irrelevant results and may bias the conclusion of this study.

Feature classification

In order to acquire high understanding of facial emotions and to encode efficiently the joint landmarks relative position distribution given an emotion, we leverage the powerful data representation abilities that deep classifier have. As a consequence, we used a deep model to classify emotions, especially a graph neural networks. Graph Neural Network is a general deep model introduced in [15] that is able to deal with data represented as graph. Contemporain GNN models have been proposed [6, 17]. We have used such models to perform our classification task. Now, we introduce $\theta \in \Theta$ such that Θ is an m -dimensional Euclidean space. We denote by f_θ our parametrized GNN classifier that determines for each instance the discrete conditional probability distribution given θ and the instanced graph $G_{\mathbf{X}}$ over these C classes. The optimal parameter $\theta^* \in \Theta$ exists and determines what is an optimal classifier that we denote f_{θ^*} .

Feature understanding

Once having determined θ^* , we can analyze how f_{θ^*} internal decision processes have been designed. The classifier f_θ , combine non-linearly many different scales of input representations. Diving into such tortuous mixing is hopeless in order to have a human-understandable representation of this mixing. We rather use a method that seeks to highlight in a human-affordable and precise manner this representation combining. Explaining methods do not have a formalized formulation since what they are supplying are still under investigation from psychological, philosophical or even computational aspects.

Experiment

In this section we provide our actual framework from the dataset we have used to labelizing tools as well as our implementation setup that justify the results shown below.

Dataset

For a wide deployment purpose of our method, we use publicly and highly available RGB images coming from famous computer vision datasets.

300W dataset [12, 13, 14] is widely used computer vision datasets for facial landmarks problems. This dataset is interesting since it supplies a very large contextual example including a wide range of emotions. This variety of data covers the majority of human emotions that a person using our tool can encounter in daily life. This dataset is composed of respectively 300 indoor-contextualized and 300 outdoor-contextualized instances. It is from this dataset that we have inferred three-dimensional landmarks and emotions.

Face alignment tools

As a three-dimensional landmark detector fa^* we used [3]. This model is one of the state-of-the-art 2D-to-3D landmark detector and it detects $p = 68$ facial landmarks. It relies on the Face Alignment Network. Several methods (e.g [22]) have been designed to directly find three-dimensional landmarks from raw images. But for accuracy concerns, regressing models in two-dimensional space is easier from optimization point of view than in three-dimensional ones because the searching space is in the lower dimension. So a 2D-to-3D models that just extend the last dimension based on an already accurate 2D position is preferable to a 3D regression model that is trained from scratch.

Emotional recognizer tools

As an emotion recognizer tool er^* we have chosen the state-of-the-art model [9]. It achieves 76.82% of accuracy on FER2013 dataset which is the state-of-the-art dataset for such classification task. It has been trained to detect 7 emotion types : angry, fear, disgust, happy, sad, surprise and neutral. This method is using a segmentation network to refine feature maps that are then plugged into a Deep Residual Network and a U-Net based architecture.

As far as we know, both tools achieved respectfully state-of-the-art results in their own domain. That is why we have chosen them to lead this study. We now provide our own classifier architecture that helps us to classify our graph representation regarding those $C = 7$ emotion types.

Classifier architecture and implementation setup

The architecture of f_θ (Figure 2a) is designed upon a local encoding that focuses on understanding each connex components and a global setting that combines, without any prior, these local information flows that afterward feed a linear layer allowing the weighting of each preprocessed semantics (with the concatenation of global average pooling and global maximum pooling layers) for the emotion classification. Our experimental setup is defined as follows : for our deep classifier we have used GCN modules for each semantic component encoding and a GCN for the subsequent global encoding. A linear module large of 512 neurons is used for the classification phase. We have used the Adam optimizer with a learning rate of 7×10^{-4} . From a hardware point, we have used Nvidia A100 40 Gb GPUs.

	Infidelity	Sparsity
GNNExplainer	1.67	0.23
LRP-GNN	0.79	0.14
ScoreCAM GNN	0.21	0.89

Statistical comparative study of explaining methods: left column is the averaged infidelity over the dataset 300W given each assessed explaining methods. The infidelity [19] measures the unfaithfulness of a explaining method regarding an instance and a deep model. The lowest the infidelity, the better. We have measured as well the average sparsity that provided explanations actually are regarding each explaining method. Sparsity is another objective metric that measuring the conciseness of each provided explanation. The higher the sparsity, the better. According to both objective metrics, ScoreCAM GNN has shown empirically better results.

Results

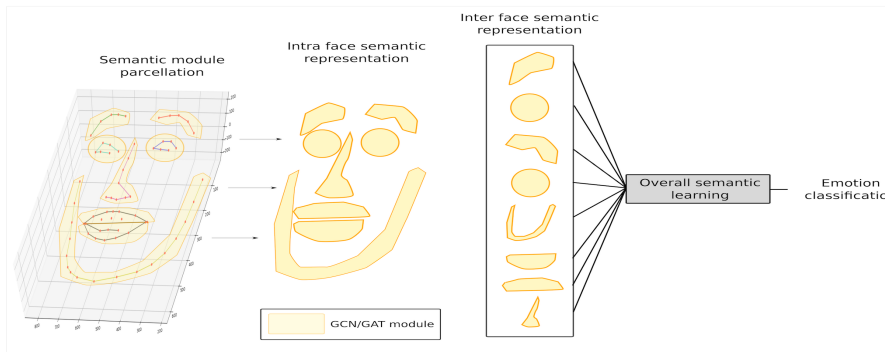
In this section, we firstly provide the performances we have reached for classifying emotions within our graph-based approach. Then, we will illustrate we are our results regarding the emotion singularization obtained through ScoreCAM GNN which has been proven to be the most relevant regarding objective assessment metrics.

Classification performances

For the classification task label distribution needs to be balanced for an efficient learning. It turns out that emotions were not uniformly distributed over instances of 300W. Consequently, we had a focus only on classes which occur at least at 10% among 300W instances. Those emotions are happy, neutral, sad and fear, it thus means that $C = 4$. We also note some emotion misattributions by er^* . Under our experimental setup, we reach 66% of accuracy on the graph-variant of 300W dataset. This result remains acceptable considering the misattribution problem and general classification model accuracy that surround 80% of accuracy. Based on these results, we now provide some qualitative emotion characterization measurements.

Relative explaining method assessment

In order to supply relevant results as possible, we have leaded a statistical comparative study among state-of-the-art explaining methods. In this study we retain GNNExplainer, the state-of-the-art method, LRP-GNN and ScoreCAM GNN that provide explanations under realistic time amount with respect to XGNN or Sub-graphX. As shown in Table 2, we have measured some explaining methods objective assessment metrics (faithfulness and sparsity) and ScoreCAM GNN seems to be the most relevant explaining method. As a consequence, we will only consider ScoreCAM GNN explanations to lead our emotional understanding study. ScoreCAM GNN [11] is an extension to non-Euclidean domain of ScoreCAM, initially introduced by [18]. ScoreCAM linearly combine the highest level of data representation and weight each of them by its own contribution weight relatively to the classification task. The outcome of ScoreCAM GNN is, given a classifier and an instance, a normalized distribution of these contributions for classifying the instance, of each element of the domain the instance leans on (i.e nodes). In our context, it provided the impact of each node regarding the classification of the considered graph according to f_{θ^*} . In other terms, it provides the importance of




(a) Graph-based classifier structure: We designed the signal propagating scheme by concentrate semantic localities between themselves. Each semantic part is seen as a connex component of the assessed graph. As long as each component encodes a partial emotion information, we gather each of them in a semantic graph with eight components that are connected in an unconstrained manner, it means that it is a complete graph. The obtained embedded data distribution is then fed to a GCN module and a linear part that helps to classifying.

each landmark which has an almost one-to-one physical mapping (facial muscle neighborhood) involved in emotion description, so a deep facial emotion understanding.

Emotion understanding

Understanding deep model behaviors may be dependent on model accuracy since model accuracy reflects the numerical understanding of the model learning tasks. Under the above considerations, some results may be a bit irrelevant but it appears experimentally that provided an explanation of f_{θ^*} behaviors are in accordance with the social baseline of what an emotion is with respect to other emotions thus what is characterizing facial emotions visually. As shown in Figure 3, we have displayed the four emotions characterization with their initial image representation. For classification task, a single instance explaining method always provided their explanation with including what is the relevant information (current class understanding) and what is not relevant (remaining classes) since classification are seamlessly a space partition problem. We firstly described what is the neutral emotion because it can be seen as a baseline for understanding and characterizing other emotions. Indeed we can see neutral emotion as a non-emotional emotion. Other emotions characterization will be then supplied in a constrastive fashion, with as a baseline, the neutral emotion that is geometrically unambiguous. For happy emotion, what we noticed under what ScoreCAM GNN revealed is that upper, lower mouth and nose are conjointly involved in the raw description of what happiness is over a human face. Analogously, sadness emotion can be seen as the symmetric of happiness and we found out that recognizing sadness from the geometrical point of view is also done by a specific mouth and nose muscle configuration that is far different from the happiness configuration. These results are in social accordance of what we can expect from the response of what characterizing facial emotions.

(b) The initial semantic node indexation that we used for our experimentation. This indexation is widely used for tackling face landmark assignment problems.



Conclusion

In this study we propose an original approach to characterize visually what are human face emotions. By using, state-of-the-art methods to recover high-level information, such as image-based emotion recognition or semantic facial landmarks positioning, we encode human faces as a graph to enforce and leverage the relational aspects of muscle activation patters involved in human emotions revealing processes. Even with uncomplete information due mainly to self-weakness of labeling methods, that stands out as being state-of-the-art method in their own field of application, we recover social expected results of what are emotions and which muscle group are involved to visually described them. Further works may include more accurate labeling methods in order to increase the overall accuracy of the graph based classifier. Other explaining methods may be used to provide emotional understandings.

References

- [1] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On Pixel-Wise Explanations for Non-Linear Classifier Decisions by Layer-Wise Relevance Propagation. *PLOS ONE*, 10(7):e0130140, July 2015. Publisher: Public Library of Science.
- [2] Danny Oude Bos. EEG-based Emotion Recognition. page 18.
- [3] Adrian Bulat and Georgios Tzimiropoulos. How Far are We from Solving the 2D & 3D Face Alignment Problem? (and a Dataset of 230,000 3D Facial Landmarks). In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1021–1030, Venice, October 2017. IEEE.
- [4] Luke Guerdan, Alex Raymond, and Hatice Gunes. Toward Affective XAI: Facial Affect Analysis for Understanding Explainable Human-AI Interactions. In *2021 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, pages 3789–3798, Montreal, BC, Canada, October 2021. IEEE.
- [5] Berthold K. P. Horn. Closed-form solution of absolute orientation using unit quaternions. *J. Opt. Soc. Am. A, JOSAA*, 4(4):629–642, April 1987. Publisher: Optica Publishing Group.
- [6] Thomas N. Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv:1609.02907 [cs, stat]*, February 2017. arXiv: 1609.02907.
- [7] Zheng Lian, Jianhua Tao, Bin Liu, Jian Huang, Zhanlei Yang, and

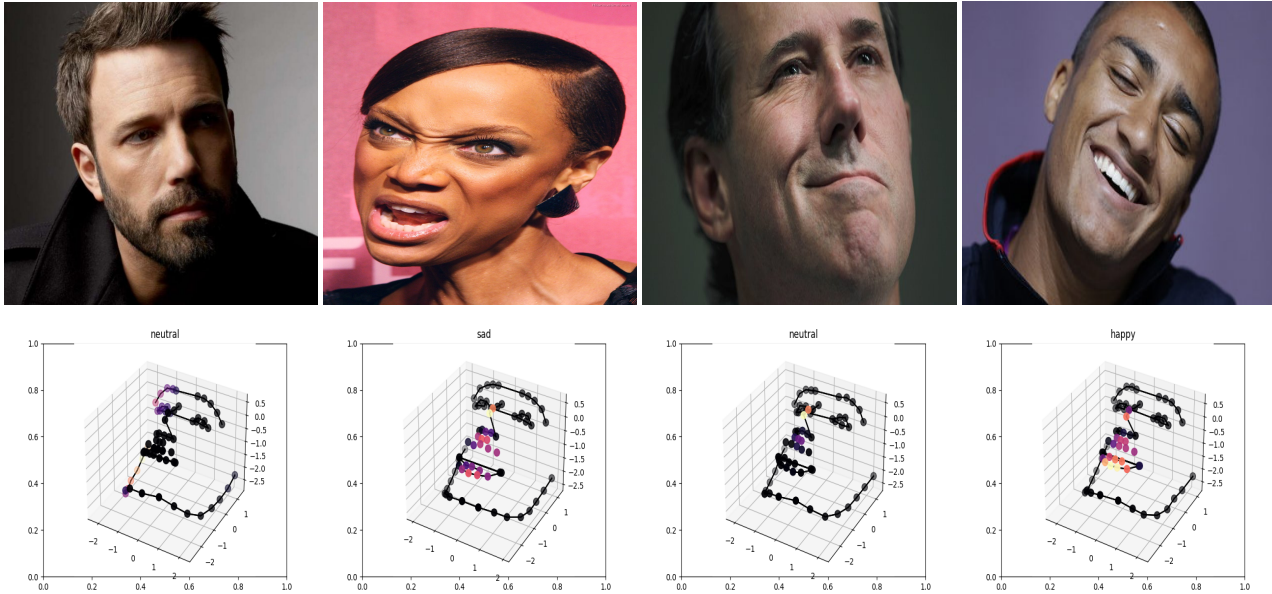


Figure 3. Emotion characterization through ScoreCAM GNN: In these four dual representation of human faces we highlight where is the numerical definition of facial emotions. Neutral emotion is the baseline emotion to characterize other ones. It shows that happiness and sadness are both highly embedded in mouth and nose motions which are socially relevant.

- Rongjun Li. Context-Dependent Domain Adversarial Neural Network for Multimodal Emotion Recognition. In *Interspeech 2020*, pages 394–398. ISCA, October 2020.
- [8] Mark A. Neerinx, Jasper van der Waa, Frank Kaptein, and Jurriaan van Diggelen. Using Perceptual and Cognitive Explanations for Enhanced Human-Agent Team Performance. In *Engineering Psychology and Cognitive Ergonomics: 15th International Conference, EPCE 2018, Held as Part of HCI International 2018, Las Vegas, NV, USA, July 15-20, 2018, Proceedings*, pages 204–214, Berlin, Heidelberg, July 2018. Springer-Verlag.
- [9] Luan Pham, The Huynh Vu, and Tuan Anh Tran. Facial Expression Recognition Using Residual Masking Network. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 4513–4519, January 2021. ISSN: 1051-4651.
- [10] Adrien Raison, Pascal Bourdon, and David Helbert. EiX-GNN : Concept-level eigencentality explainer for graph neural networks. June 2022.
- [11] Adrien Raison, Pascal Bourdon, and David Helbert. ScoreCAM GNN : une explication optimale des réseaux profonds sur graphes. In *XXVIII^{ème} Colloque Francophone de Traitement du Signal et des Images (GRETSI)*, Nancy, France, September 2022.
- [12] Christos Sagonas, Epameinondas Antonakos, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 Faces In-The-Wild Challenge: database and results. *Image and Vision Computing*, 47:3–18, March 2016.
- [13] Christos Sagonas, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. 300 Faces in-the-Wild Challenge: The First Facial Landmark Localization Challenge. In *2013 IEEE International Conference on Computer Vision Workshops*, pages 397–403, Sydney, Australia, December 2013. IEEE.
- [14] Christos Sagonas, Georgios Tzimiropoulos, Stefanos Zafeiriou, and Maja Pantic. A Semi-automatic Methodology for Facial Landmark Annotation. In *2013 IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 896–903, OR, USA, June 2013. IEEE.
- [15] F. Scarselli, M. Gori, Ah Chung Tsoi, M. Hagenbuchner, and G. Monfardini. The Graph Neural Network Model. *IEEE Trans. Neural Netw.*, 20(1):61–80, January 2009.
- [16] Thomas Schnake, Oliver Eberle, Jonas Lederer, Shinichi Nakajima, Kristof T. Schütt, Klaus-Robert Müller, and Grégoire Montavon. Higher-Order Explanations of Graph Neural Networks via Relevant Walks. *IEEE Trans. Pattern Anal. Mach. Intell.*, pages 1–1, 2021. arXiv:2006.03589 [cs, stat].
- [17] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. Graph Attention Networks. *arXiv:1710.10903 [cs, stat]*, February 2018. arXiv: 1710.10903.
- [18] Haofan Wang, Zifan Wang, Mengnan Du, Fan Yang, Zijian Zhang, Sirui Ding, Piotr Mardziel, and Xia Hu. Score-CAM: Score-Weighted Visual Explanations for Convolutional Neural Networks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 111–119, Seattle, WA, USA, June 2020. IEEE.
- [19] Chih-Kuan Yeh, Cheng-Yu Hsieh, Arun Suggala, David I Inouye, and Pradeep K Ravikumar. On the (In)Fidelity and Sensitivity of Explanations. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [20] Zhitao Ying, Dylan Bourgeois, Jiaxuan You, Marinka Zitnik, and Jure Leskovec. GNNExplainer: Generating Explanations for Graph Neural Networks. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [21] Hao Yuan, Jiliang Tang, Xia Hu, and Shuiwang Ji. XGNN: Towards Model-Level Explanations of Graph Neural Networks. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '20*, pages 430–438, New York, NY, USA, August 2020. Association for Computing Machinery.
- [22] Xiangyu Zhu, Zhen Lei, Xiaoming Liu, Hailin Shi, and Stan Z Li. Face Alignment Across Large Poses: A 3D Solution. page 10.