MORPHSET: AUGMENTING CATEGORICAL EMOTION DATASETS WITH DIMENSIONAL AFFECT LABELS USING FACE MORPHING

Vassilios Vonikakis

Amazon Web Services Singapore Neo Yuan Rong Dexter, Stefan Winkler

National University of Singapore School of Computing

ABSTRACT

Emotion recognition and understanding is a vital component in human-machine interaction. Dimensional models of affect such as those using valence and arousal have advantages over traditional categorical ones due to the complexity of emotional states in humans. However, dimensional emotion annotations are difficult and expensive to collect, therefore they are not as prevalent in the affective computing community. To address these issues, we propose a method to generate synthetic images from existing categorical emotion datasets using face morphing as well as dimensional labels in the circumplex space with full control over the resulting sample distribution, while achieving augmentation factors of at least 20x or more.

1. INTRODUCTION

Classification of basic prototypical high-intensity facial expressions is an extensively researched topic. Inspired initially by the seminal work of Ekman [1], it has made significant strides in recent years [2, 3]. However, such approaches have limited applicability in real life, where people rarely exhibit high-intensity prototypical expressions; low-key, non-prototypical expressions are much more common in every-day situations. Consequently, researchers have started to explore alternative approaches, such as intensity of facial action units [4,5], compound expressions [6], or dimensional models of facial affect [7–9]. Yet these alternatives have not received much attention in the computer vision community compared to categorical models.

One major problem that impedes the widespread use of dimensional models is the limited availability of datasets. This stems from the difficulty of collecting large sets of images across many subjects and expressions. It is even more difficult to acquire reliable emotion annotations for supervised learning. Continuous dimensional emotion labels such as Valence and Arousal are difficult for laymen users to assess and assign, and hiring experienced annotators to label a large corpus of images is prohibitively expensive and time consuming. Since even experienced annotators may disagree on these labels, multiple annotations per image are required, which further increases the cost and complexity of the task. Yet there are no guarantees that the full range of possible expressions and intensities will be covered, resulting in imbalanced datasets. Consequently, large, *balanced* emotion datasets, with highquality annotations, covering a wide range of expression variations and expression intensities of many different subjects, are in short supply.

Our approach attempts to address this need. We propose a fast and cost-effective augmentation framework to create balanced, annotated image datasets, appropriate for training Facial Expression Analysis (FEA) systems for dimensional affect. The framework uses high-quality facial morphings to transform typical categorical datasets (usually 7 expressions per subject) into dimensional ones, with an augmentation factor of at least 20x or more. More importantly, it produces *multiple annotated expressions per subject*, balanced across the Valence-Arousal (VA) space. The resulting synthesized facial images look realistic and visually convincing. We also demonstrate that they can be used very effectively for training and testing real-world FEA systems.

Although morphing, as a means of deriving an extended set of facial expressions, is a widely used tool in psychology [10, 11], it has found limited adoption in the computer vision community. Traditional work on expression synthesis usually incorporates manipulation of the facial geometry and texture mapping in images or videos [12, 13]. Other approaches include the use of 3D meshes adopted from RGB-D space [14] or the adjustment of Action Units from the Facial Action Coding System (FACS) [15]. More recently, researchers have employed Generative Adversarial Networks (GANs) for this purpose [16]. Various conditional GAN variations have been used to generate novel expressions while preserving identity and other facial details [17-19]. While most of them also take the categorical approach, a few models have been proposed based on action units [20] and continuous emotion dimensions [21]. These GANs generally require

Part of this work was conducted while V. Vonikakis and S. Winkler were with the Advanced Digital Sciences Center (ADSC), University of Illinois at Urbana-Champaign, Singapore. The contribution from V. Vonikakis was made by him prior to joining AWS. Contact: winkler@comp.nus.edu.sg

a large dataset to start with, with no guarantees that generated faces will not exhibit unnatural artifacts, and the difficulty of creating proper annotations remains. In our approach on the other hand, we have full control over the pipeline, resulting in deterministic outputs both in terms of synthetic images and dimensional emotion labels.

Our main contributions can be summarized as follows:

- A new dataset augmentation framework that can transform a typical categorical facial expression dataset into a *balanced* augmented one.
- The framework can generate hundreds of different expressions *per subject* with full user control over their distribution.
- The augmented dataset comes with automatically generated, highly consistent Valence/Arousal annotations of continuous dimensional affect.

The code for the proposed augmentation framework is available at https://github.com/dexterdley/MorphSet.

2. DATASET GENERATION

We assume a 2-dimensional polar affective space, similar to the Valence-Arousal (VA) space of the circumplex model [7], with Neutral located at the center. The typical 7 facial expressions, which are usually included in categorical emotion datasets, Neutral (Ne), Happy (Ha), Surprised (Su), Afraid (Af), Angry (An), Disgusted (Di) and Sad (Sa), can be mapped to points with specific coordinates in the polar AV space. Apart from these 7 points however, there is a lot of empty space on the remaining AV plane. These missing facial expressions comprise: (a) different expression variations e.g. Delighted, Excited, Upset etc., located at different angles in the AV polar space, and (b) different intensity variations of the expressions, e.g. slightly happy, moderately happy, extremely happy etc., spanning the area from the center (Neutral) outwards toward the periphery of the AV space. The basic premise of MorphSet is that many of these expression variations can be synthesized by high-quality morphings between images of categorical expressions.

Let \mathbf{F}_i^E denote the face image of subject i with facial expression E. For categorical datasets, usually $E \in \{\text{Ne,Ha,Su,Af,An,Di,Sa}\}$. Let θ^E denote the specific angle of each expression in the polar AV space, estimated from emotion studies [7, 23]. Let $I_i^E \in [0, 1]$ denote the intensity of expression E of subject i. Zero expression intensity $I^E = 0$ coincides with Neutral (by definition $I^{\text{Ne}} = 0$), while $I^E = 1$ represents the highest possible intensity. Let $M_p\left(\mathbf{F}_i^{source}, \mathbf{F}_i^{target}, r\right)$ be a morphing function,

Let $M_p(\mathbf{F}_i^{source}, \mathbf{F}_i^{target}, r)$ be a morphing function, based on p facial landmarks, that returns a new face image, which is the result of morphing \mathbf{F}_i^{source} towards \mathbf{F}_i^{target} with a ratio $r \in [0, 1]$; when r = 0 the morphed image is identical to \mathbf{F}_i^{source} , and when r = 1 it is identical to \mathbf{F}_i^{target} . Any contemporary morphing approach can be used for this, such as Delaunay triangulation followed by local warping of 68 facial landmarks from Dlib [24] face recognition system.

Our augmentation framework is based on 2 types of morphings. In order to synthesize new expression variations, Apex to Apex morphing (1) is used, between the given apex expressions of the categorical dataset:

ł

Apex to Apex
$$\begin{cases} \mathbf{F}_{i}^{A_{1}rA_{2}} = M_{p} \left(\mathbf{F}_{i}^{A_{1}}, \mathbf{F}_{i}^{A_{2}}, r \right) \\ I_{i}^{A_{1}rA_{2}} = (1-r)I_{i}^{A_{1}} + rI_{i}^{A_{2}} \\ \theta^{A_{1}rA_{2}} = (1-r)\theta^{A_{1}} + r\theta^{A_{2}} \end{cases}$$
(1)

where A, A_1 and A_2 are apex expressions from the parent dataset. In order to synthesize new intensity variations, Neutral to Apex morphing (2) is used, between the NE image and a given (or interpolated) apex image:

Neutral to Apex
$$\begin{cases} \mathbf{F}_{i}^{rA} = M_{p} \left(\mathbf{F}_{i}^{\text{Ne}}, \mathbf{F}_{i}^{A}, r \right) \\ I_{i}^{rA} = rI_{i}^{A} \\ \theta^{rA} = \theta^{A} \end{cases}$$
(2)

Fig. 1 shows an example of these 2 types of morphing. Once new interpolated apex expressions are generated by equation (1), 'neutral to interpolated apex' morphings can further be generated by applying equation (2) on them.

For every given or generated face image \mathbf{F}_i^E , with I_i^E and θ^E , the Valence and Arousal annotations can be computed as $V_i^E = I_i^E \cos(\theta^E)$ and $A_i^E = I_i^E \sin(\theta^E)$.

Fig. 2 illustrates the proposed augmentation framework for a typical categorical dataset with the 7 prototypical expressions. We start from 10° , the approximate location of Happy in VA space, and proceed in increments of 15° steps in order to span the whole range up to 205° , where Sad is approximately located. The proposed template is bounded within $[10^{\circ}, 205^{\circ}]$ only because negative Arousal expressions (Sleepy, Tired, Bored, Calm, etc.) are absent from the typical categorical emotion datasets. We use an angle increment of 15° and an intensity increment of 0.1, because they strike a good balance between expression granularity, augmentation factor and symmetry between the positions of the given prototypical expressions in the AV space.

Based on the above selected granularity, for a typical categorical dataset of 7 facial images per subject, the proposed augmentation framework can generate 134 new images, reaching a total of 134+7=141 facial images per subject. This translates to an augmentation factor of 20x, or 40x with simple image mirroring. Doubling the angular and radial granularity to 7.5° and 0.05, respectively, results in an augmentation factor of 80x, or 160x with image mirroring.

We build an example augmented dimensional dataset from a combination of 3 categorical datasets, which have been extensively used in psychology: the Radboud [22], Karolinska [25], and Warsaw [26] datasets. We select them



Fig. 1. Examples of the 2 types of face morphings utilized in the proposed augmentation framework, using images from the Radboud dataset [22]. In this example, all images are synthesized out of 4 given images from the original dataset (outlined in black). Top: Neutral to Apex (Happy) morphing. Bottom: Apex1 (Sad) to Apex2 (Disgusted) morphing.



Fig. 2. Dataset augmentation framework based on face morphing. Intensity of expression is represented with size and color saturation. Outlined shapes indicate apex expressions.

because they have superior image quality, their facial expressions were guided by FACS experts, and they are accompanied by validation studies that provide the *perceived* intensity of expression I_i^E .

Using the proposed MorphSet augmentation framework, all faces are aligned with respect to the subjects' eyes, and synthetic images are generated for each subject. The resulting augmented dataset comprises more than 55,000 images, which – with finer granularity and mirroring – can reach over 450,000. More importantly though, each image comes

with continuous Valence, Arousal, and Intensity annotations, which can be used to train dimensional FEA systems.

As the proposed augmented dataset contains only frontal images, an FEA system trained on it would not be invariant to different head poses. There are various solutions to this, which are discussed in detail in [27]. In fact, we have successfully used such an approach to build a robust in-the-wild facial expression analysis system [28].

3. COMPARISON & DISCUSSION

In this section, we evaluate the benefits of MorphSet as a dataset for facial emotion recognition. We also describe the AffectNet [8] and Aff-Wild [9] databases commonly used for emotion recognition.

AffectNet [8] is a dataset comprising of roughly 450,000 images of in-the-wild-facial expressions collected from the Internet using affect keywords. These images are manually annotated with both categorical and dimensional labels. Despite its size, the training samples of AffectNet are noisy and highly imbalanced, causing many learning algorithms to perform poorly on the minority classes. Furthermore, human annotator agreement is just over 60%, suggesting that the dataset suffers from noisy/incorrect annotations.

Aff-Wild [9] is an in-the-wild video dataset, consisting of 298 Youtube videos displaying reactions of 200 subjects. The annotations for valence and arousal were collected continuously via joystick. For direct comparison with AffectNet and MorphSet, we extract individual frames according to the affect annotations from the time-stamped video sequences.

Table 1 compares these two databases with MorphSet, which is unique in that it provides a balanced distribution of facial expressions for each subject.

To compare baseline FEA results, we train a ResNet-18 model [29] on each dataset. We add two neurons after the



Fig. 3. Randomly sampled training images around the same valence and arousal annotations from AffectNet, Aff-Wild, and MorphSet (top to bottom rows). Circles indicate wrong Arousal/Valence annotations (e.g. positive instead of negative), while squares indicate correct overall emotion but wrong intensity (e.g. extremely afraid instead of slightly afraid).

Table 1. Comparison to other dimensional datasets.					
	AffectNet	Aff-Wild	MorphSet		
Unique subjects	≈450,000	200	167		
Expressions	1	N/A	≈342		
per subject	1	(video)			
Total	$\sim 150,000$	1,224,100	55,000 up to		
images	~450,000		≈450,000		
Annotators	12	8	N/A		

final fully connected layer to predict valence and arousal dimensions and minimize the L2 loss. Input images are resized to 224x224 with standard augmentations (e.g. affine transformations). We evaluate on the validation set of AffectNet and on 20% of randomly selected, unseen identities of Aff-Wild and MorphSet.

Table 2 shows the Root Mean Square error (RMSE) and Concordance Correlation Coefficient (CCC) for each dataset. Although the results are not directly comparable because of the difference in test sets, they provide useful insights. The performance for MorphSet is significantly better than Affect-Net and Aff-Wild, likely due to the frontal and highly controlled conditions of the images as well as higher consistency of the VA annotations.

The fact that wild datasets are often noisier and less controlled in terms of facial expressions than MorphSet is further illustrated by Fig. 3, where images with specific VA labels are randomly sampled from each of the three datasets. AffectNet and Aff-Wild samples show significant fluctuations and outliers in facial expressions (indicated with circles and squares in Fig. 3), whereas the facial expressions from MorphSet are much more consistent across different subjects.

Tuble 1 . Rebitter 10 Subellife rebuild for each autubet.	Table 2. ResNet-18	baseline results for each dataset.
------------------------------------------------------------------	--------------------	------------------------------------

	AffectNet	Aff-Wild	MorphSet
RMSE Valence	0.427	0.407	0.157
RMSE Arousal	0.390	0.266	0.155
CCC Valence	0.533	0.186	0.915
CCC Arousal	0.418	0.174	0.821

4. CONCLUSIONS

We presented MorphSet, a dataset augmentation framework that can transform a typical categorical dataset of facial expressions into a balanced augmented one using image morphing. We use the framework to generate hundreds of different expressions per subject with full user control over their distribution. The augmented dataset comes with automatically generated, highly consistent dimensional annotations suitable for supervised learning of continuous affect.

5. REFERENCES

- P. Ekman and W. V. Friesen, "Constants across cultures in the face and emotion," *Journal of Personality and Social Psychol*ogy, vol. 17, no. 2, pp. 124–129, 1971.
- [2] E. Sariyanidi, H. Gunes, and A. Cavallaro, "Automatic analysis of facial affect: A survey of registration, representation, and recognition," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 37, no. 6, pp. 1113–1133, June 2015.
- [3] S. Li and W. Deng, "Deep facial expression recognition: A survey," *IEEE Trans. Affective Computing*, March 2020.
- [4] O. Rudovic, V. Pavlovic, and M. Pantic, "Context-sensitive dynamic ordinal regression for intensity estimation of facial action units," *IEEE Trans. Pattern Analysis and Machine Intelli*gence, vol. 37, no. 5, pp. 944–958, 2015.
- [5] F. Benitez-Quiroz, R. Srinivasan, and A. M. Martinez, "Discriminant functional learning of color features for the recognition of facial action units and their intensities," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 41, no. 12, pp. 2835–2845, Dec. 2019.
- [6] S. Du, Y. Tao, and A. M. Martinez, "Compound facial expressions of emotion," *Proc. National Academy of Sciences*, vol. 111, no. 15, pp. E1454–E1462, 2014.
- [7] J. A. Russell, "A circumplex model of affect," *Journal of Personality and Social Psychology*, vol. 39, no. 6, p. 1161–1178, 1980.
- [8] A. Mollahosseini, B. Hasani, and M. H. Mahoor, "AffectNet: A database for facial expression, valence, and arousal computing in the wild," *IEEE Trans. Affective Computing*, vol. 10, no. 1, pp. 18–31, Jan. 2019.
- [9] S. Zafeiriou, D. Kollias, M. A. Nicolaou, A. Papaioannou, G. Zhao, and I. Kotsia, "Aff-Wild: Valence and arousal 'in-thewild' challenge," in *Proc. IEEE Computer Vision and Pattern Recognition Workshops (CVPRW)*, July 2017, pp. 1980–1987.
- [10] R. P. C. Kessels, B. Montagne, A. W. Hendriks, D. I. Perrett, and E. H. F. Haan, "Assessment of perception of morphed facial expressions using the emotion recognition task: Normative data from healthy participants aged 8-75," *Journal of Neuropsychology*, vol. 8, no. 1, pp. 75–93, 2014.
- [11] L. J. Wells, S. M. Gillespie, and P. Rotshtein, "Identification of emotional facial expressions: Effects of expression, intensity, and sex on eye gaze," *PLOS One*, vol. 11, no. 12, pp. 1–20, Dec. 2016.
- [12] B.-J. Theobald, I. Matthews, M. Mangini, J. R. Spies, T. R. Brick, J. F. Cohn, and S. M. Boker, "Mapping and manipulating facial expression," *Language and Speech*, vol. 52, no. 2-3, pp. 369–386, 2009.
- [13] D. Kollias, S. Cheng, E. Ververas, I. Kotsia, and S. Zafeiriou, "Deep neural network augmentation: Generating faces for affect analysis," *International Journal of Computer Vision*, vol. 128, pp. 1455–1484, May 2020.
- [14] J. Thies, M. Zollhöfer, M. Nießner, L. Valgaerts, M. Stamminger, and C. Theobalt, "Real-time expression transfer for facial reenactment," *ACM Trans. Graphics*, vol. 34, no. 6, pp. 183:1–183:14, Oct. 2015.

- [15] H. Yu, O. Garrod, R. Jack, and P. Schyns, "Realistic facial animation generation based on facial expression mapping," in *Proc. SPIE 5th International Conference on Graphic and Image Processing (ICGIP)*, vol. 9069, 2014.
- [16] J. Han, Z. Zhang, N. Cummins, and B. W. Schuller, "Adversarial training in affective computing and sentiment analysis: Recent advances and perspectives," *IEEE Computational Intelligence Magazine*, vol. 14, no. 2, pp. 68–81, 2019.
- [17] H. Ding, K. Sricharan, and R. Chellappa, "ExprGAN: Facial expression editing with controllable expression intensity," in *Proc. 32nd AAAI Conference on Artificial Intelligence*, New Orleans, LA, Feb. 2018.
- [18] X. Wang, W. Li, G. Mu, D. Huang, and Y. Wang, "Facial expression synthesis by U-Net conditional generative adversarial networks," in *Proc. ACM International Conference on Multimedia Retrieval (ICMR)*, Yokohama, Japan, 2018, pp. 283– 290.
- [19] E. Ververas and S. Zafeiriou, "SliderGAN: Synthesizing expressive face images by sliding 3D blendshape parameters," *International Journal of Computer Vision*, vol. 128, pp. 2629–2650, Nov. 2020.
- [20] A. Pumarola, A. Agudo, A. M. Martinez, A. Sanfeliu, and F. Moreno-Noguer, "GANimation: Anatomically-aware facial animation from a single image," in *Proc. European Conf. Computer Vision (ECCV)*, 2018, pp. 835–851.
- [21] V. Vielzeuf, C. Kervadec, S. Pateux, and F. Jurie, "The many moods of emotion," *CoRR*, vol. abs/1810.13197, 2018.
- [22] O. Langner, R. Dotsch, G. Bijlstra, D. H. J. Wigboldus, S. T. Hawk, and A. van Knippenberg, "Presentation and validation of the Radboud faces database," *Cognition and Emotion*, vol. 24, no. 8, pp. 1377–1388, 2010.
- [23] G. Paltoglou and M. Thelwall, "Seeing stars of valence and arousal in blog posts," *IEEE Trans. Affective Computing*, vol. 4, no. 1, pp. 116–123, Jan 2013.
- [24] D. E. King, "Dlib-ml: A machine learning toolkit," *Journal of Machine Learning Research*, vol. 10, pp. 1755–1758, 2009.
- [25] E. Goeleven, R. D. Raedt, L. Leyman, and B. Verschuere, "The Karolinska directed emotional faces: A validation study," *Cognition and Emotion*, vol. 22, no. 6, pp. 1094–1118, 2008.
- [26] M. Olszanowski, G. Pochwatko, K. Kuklinski, M. Scibor-Rylski, P. Lewinski, and R. K. Ohme, "Warsaw set of emotional facial expression pictures: A validation study of facial display photographs," *Frontiers in Psychology*, vol. 5, p. 1516, 2015.
- [27] V. Vonikakis and S. Winkler, "Identity-invariant facial landmark frontalization for facial expression analysis," in *Proc. IEEE Int'l Conf. Image Processing (ICIP)*, Abu Dhabi, Oct. 2020.
- [28] V. Vonikakis and S. Winkler "Efficient facial expression analysis for dimensional affect recognition using geometric features," *arXiv eprint*, vol. 2106.07817, June 2021.
- [29] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Computer Vision* and Pattern Recognition (CVPR), June 2016, pp. 770–778.