

Affective Processes: stochastic modelling of temporal context for emotion and facial expression recognition

Enrique Sanchez¹ Mani Kumar Tellamekala² Michel Valstar² Georgios Tzimiropoulos^{1,3}

¹ Samsung AI Center, Cambridge, UK

² University of Nottingham, Nottingham, UK

³ Queen Mary University London, London, UK

e.lozano@samsung.com {psxmkt, michel.valstar}@nottingham.ac.uk g.tzimiropoulos@qmul.ac.uk

Abstract

Temporal context is key to the recognition of expressions of emotion. Existing methods, that rely on recurrent or self-attention models to enforce temporal consistency, work on the feature level, ignoring the task-specific temporal dependencies, and fail to model context uncertainty. To alleviate these issues, we build upon the framework of Neural Processes to propose a method for apparent emotion recognition with three key novel components: (a) probabilistic contextual representation with a global latent variable model; (b) temporal context modelling using task-specific predictions in addition to features; and (c) smart temporal context selection. We validate our approach on four databases, two for Valence and Arousal estimation (SEWA and AffWild2), and two for Action Unit intensity estimation (DISFA and BP4D). Results show a consistent improvement over a series of strong baselines as well as over state-of-the-art methods.

1. Introduction

In this paper, we address the problem of facial behaviour recognition from video, in particular, the problem of recognising apparent emotions in terms of Valence and Arousal [49], and facial expressions in terms of Action Unit intensity [9]. This is a longstanding problem in video recognition which has been extensively studied by the computer vision community [24, 62, 39, 19, 31, 26, 65, 72, 75, 69, 53, 51, 52, 37, 44]. Nevertheless, even recent methods [25, 31, 75] struggle to achieve high accuracy on the most difficult datasets including SEWA [32], AffWild2 [27], BP4D [73] and DISFA [40]. Hence, there is still great progress to be made towards solving this challenging problem. In this paper, we show that effective temporal context modelling is a key feature for significantly advancing the state-of-the-art.

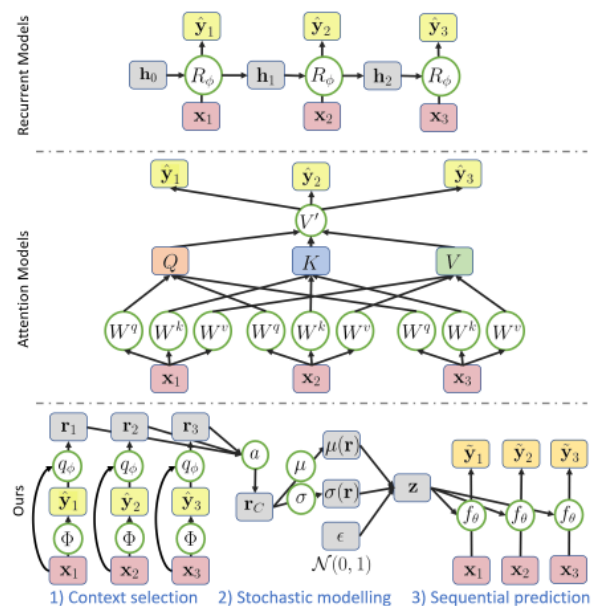


Figure 1. Different ways of modelling temporal context in the literature of emotion recognition: Recurrent Models (**top**), Attention-based Models (**middle**), and our proposed approach based on Neural Processes (**bottom**). Our method proposes a probabilistic modelling of temporal context through a global latent variable $\mathbf{z} \sim \mathcal{N}(\mu, \sigma)$, and includes a novel method for automatic context selection. The global latent variable allows to sample temporal functions $f_{\theta}(\cdot, \mathbf{z})$ that are consistent with the captured context.

But what makes emotion (and facial expression) recognition such a difficult problem? There are many, often inter-related, reasons for this including: (a) annotation of emotions is often both subjective and laborious making it hard to annotate it consistently [54]; (b) there exist only small- and medium-size emotion video datasets and, due to point (a) above, the annotations for these datasets are often noisy; (c) emotions are subtle, acting as unobserved, latent variables that only partially explain facial appearance over time [2]. They often require temporal as well as multi-modal context

in order to be robustly recognised. From these challenges, this paper focuses on solving problem (c). In particular, it proposes a completely unexplored perspective for effective incorporation of temporal context into emotion recognition.

Previous work in temporal modelling of emotions has primarily focused on modelling facial expression dynamics, i.e., the way that facial expressions evolve over time for recognition purposes. A typical deep learning pipeline for this has been a Convolution Neural Network (CNN) followed by a Recurrent Neural Network (RNN), usually an LSTM or GRU [30, 26, 24, 69]. Although these dynamics can be particularly useful in recognising specific facial expressions (e.g. [61, 62, 19, 39, 68]), we believe that their importance has been over-emphasised for the problem of emotion recognition.

We posit that temporal context is more important than facial dynamics for emotion recognition. The cues for inferring a person’s apparent emotion from their facial behaviour are often sparsely and non-regularly distributed over a temporal window of interest and collecting such distributed contextual cues is critical to inferring emotions robustly. In addition, due to the person-specific variability of facial expressions but, most importantly, due to their subjective annotation, context must be modelled in a stochastic manner.

To address the aforementioned challenges, in this work, and for the first time to the best of our knowledge, we build upon the framework of Neural Processes [14, 15] to propose a model for emotion recognition with 3 key components: (1) stochastic contextual representation with a *global* latent variable model; (2) task-aware temporal context modelling not only by using features but also task-specific predictions; and (3) effective temporal context selection.

Fig. 1 depicts an overview of the working flow our method (bottom), RNN-based methods (top), and methods based on self-attention [64] (middle). Note methods for context modelling based on self-attention do not satisfy any of (1)-(3): They are deterministic and model long-term dependencies not globally, but by using pair-wise feature similarities. Moreover, they do not use task-aware context modelling and do not perform context selection.

Overall, **we make** the following **3 contributions**:

1. We propose Affective Processes: the very first model for emotion recognition with three key properties: (a) global stochastic contextual representation; (b) task-aware temporal context modelling; and (c) temporal context selection. We conduct a large number of ablation studies illustrating the contribution of each of our model’s key features and components.
2. We show that our model is more effective in modelling temporal context not only than CNNs+RNNs but also than a strong baseline based on self-attention. Our model outperforms both baselines by large margin.

3. We validated our approach on the most difficult databases for emotion recognition in terms of Valence and Arousal estimation, namely SEWA [32] and AfFWild2 [27]. We further show that our approach can be effective for the problem of Action Unit intensity estimation on DISFA [40] and BP4D [73] datasets. On all datasets used, we show a consistent and often significant improvement over state-of-the-art methods.

2. Related Work

This section reviews related work from temporal apparent emotion recognition, Action Unit intensity estimation and neural modelling of Stochastic Processes.

Temporal Modelling in Emotion Recognition: Most existing methods model the temporal dynamics of continuous emotions using deterministic approaches such as Time-Delay Neural Networks [41], RNNs, LSTMs and GRUs [30, 26, 24, 69, 7, 57], multi-head attention models [18], 3D Convolutions [74, 33], 3D ConvLSTMs [17], and temporal-hourglass CNNs [8]. While these deterministic models are capable of effectively learning the temporal dynamics, they do not take the inherent stochastic nature of the continuous emotion labels into account. Deviating from this trend, recently, [45] applied variational RNNs [4] to the valence estimation problem to add stochasticity to the temporal model of emotion recognition. Similarly, in [46], emotion is treated as a dynamical latent system state which is estimated using Gaussian processes coupled with Kalman filters, while Deep Extended Kalman filters are used in [47]. These probabilistic methods have better latent structure-discovery capabilities than the deterministic temporal models, however, they focus only on the temporal dynamics of emotions, ignoring the temporal context. In contrast, our method primarily focuses on directly learning the global temporal context in the input sequence using stochastic processes with data-driven ‘priors’.

AU Intensity Estimation: In order to model the subjective-variability of the contextual factors in AU intensity estimation, several works advocate for probabilistic models [65, 52, 5, 51, 11, 12, 53]. Particularly, Gaussian Processes based models [11, 12] and variational latent variable models [10, 37, 66] have been extensively used for this purpose. Other probabilistic approaches to facial affect modelling include Conditional Ordinal Random Fields (CORF) [65, 52, 51] and Markov Random Fields (MRF) [53]. To learn the task-specific temporal context, context-aware feature fusion and label fusion modules are proposed in [75] to model AU intensity temporal dynamics. Although we also propose to learn the task-specific temporal context, our context learning method is completely different to that of [75].

Stochastic Process Modelling [70, 67, 59, 60]: Our method primarily builds on the recently proposed stochastic regression frameworks of Neural Processes (NPs). NPs [14,

15, 21, 16, 56, 36] are a family of predictive stochastic processes that aim to learn distributions over functions by adapting their priors to the data using context observations. Application of NPs in Computer Vision has been limited to basic image completion and super-resolution tasks [21] as *proof-of-concept* experiments. To our knowledge, we are the first to show how NPs can be used to solve challenging real-world regression tasks like that of emotion recognition.

3. Method

3.1. Problem Formulation and Background

Temporal Regression: Given as input signal the sequence \mathbf{x}_t representing, for example, a sequence of images, and the corresponding output labels $\mathbf{y}_t \in \mathbb{R}^N$, we wish to learn a function from \mathcal{X} to label space \mathcal{Y} as $f : \mathcal{X} \rightarrow \mathcal{Y}$.

Stochastic Process (SP): Unlike previous work in deterministic regression, we are not tasked with estimating a single function f but rather a distribution of functions \mathcal{P} , i.e. a *Stochastic Process (SP)*, from which we can sample $f \sim \mathcal{P}$. Each observed sequence in a training set is considered a realisation of the SP.

Neural Process (NP): NP uses an encoder-decoder architecture to learn how to sample functions $f \sim \mathcal{P}$ (i.e. from an SP) *conditioned* on a set of *Context Points*. Specifically, a given sequence is split into two parts: the context \mathcal{C} for which it is assumed that the ground truth labels are available $\{\mathbf{x}_c, \mathbf{y}_c\}_{c \in \mathcal{C}}$ and the target \mathcal{T} points, for which only the input $\{\mathbf{x}_t\}_{t \in \mathcal{T}}$ is available. Given a function f drawn from an SP conditioned on the context \mathcal{C} , the goal is to estimate $f(\mathbf{x}_t)$ for each $\mathbf{x}_t \in \mathcal{T}$. Later on, we propose a way to deal with the unavailability of context ground truth labels $\{\mathbf{y}_c\}_{c \in \mathcal{C}}$ for the practical scenario of emotion recognition.

The NP Encoder is an MLP q_ϕ that computes an individual feature representation $\mathbf{r}_c = q_\phi(\mathbf{x}_c, \mathbf{y}_c) \in \mathbb{R}^d$ ($d = 128$ in our work) for each input context pair, followed by a permutation-invariant aggregation operation. We choose simple summation for aggregation to compute a global deterministic context representation $\mathbf{r}_C = \frac{1}{C} \sum_c \mathbf{r}_c \in \mathbb{R}^d$, with C the number of context points. To account for the context stochasticity, \mathbf{r}_C is used to parameterise a Gaussian distribution from which a global latent variable \mathbf{z} can be drawn. To this end, two linear layers are used to compute the mean $\mu_C(\mathbf{r}_C)$ and standard deviation $\sigma_C(\mathbf{r}_C)$. The global latent variable $\mathbf{z} \sim \mathcal{N}(\mu_C, \sigma_C) \in \mathbb{R}^d$ represents a different realisation of the SP conditioned on the context.

The NP Decoder is an MLP f_θ that samples functions $f \sim \mathcal{P}$ through the latent variable \mathbf{z} . For a given \mathbf{z} , function samples are generated at all target points as $\mathbf{y}_t = f_\theta(\mathbf{x}_t, \mathbf{z})$. Assuming some i.i.d. Gaussian observation noise with learnable variance [20], the output of f_θ can be that of a mean and standard deviation of a Gaussian distribution: $\mathbf{y}_t \sim \mathcal{N}(f_\theta^\mu(\mathbf{x}_t, \mathbf{z}), f_\theta^\sigma(\mathbf{x}_t, \mathbf{z}))$.

3.2. Affective Processes

Let us assume that we are given a sequence of images $\{I_t\}_{t=1}^T$, and that we want to estimate facial affect, in terms of Valence and Arousal, on a per-frame basis $\mathbf{y}_t \in \mathbb{R}^2$ ¹. The framework of NPs cannot be applied to emotion recognition as is. Actually, application of NPs in Computer Vision has been limited to basic image completion and super-resolution tasks [21] as *proof-of-concept* experiments. One of our main contributions is to show how such a framework can be applied for challenging real-world temporal regression tasks like the one of emotion recognition.

3.2.1 Input Features and Pseudo-labels

The original NP formulation would directly operate on the sequence of images $\{I_t\}$. Furthermore, it would require ground truth labels for the context points \mathbf{y}_c both at training and test time. While one can assume the ground-truth labels to be available for training, their need at test time would limit the application of NPs to real-world problems. To alleviate these issues, we firstly propose to train a backbone CNN, comprising a feature extractor $\Phi(\cdot)$ and a classifier \mathbf{W} , in a standard supervised manner, to provide *independent* per-frame estimates of Valence and Arousal $\hat{\mathbf{y}}_t \in \mathbb{R}^2$.

The purpose of this network is two-fold: Firstly, the feature extractor $\Phi(\cdot)$ is used to provide a per-frame feature representation $\mathbf{x}_t = \Phi(I_t)$. Then, we can use the sequence of features \mathbf{x}_t as input to the NP instead of the image sequence. Secondly, we propose to use the output of the classifier $\hat{\mathbf{y}}_t$ as *pseudo-ground truth* labels to be fed as input to the NP. These labels are noisy, and lack temporal consistency, yet they still provide some information regarding the true labels. In this paper, we choose as our backbone an architecture inspired by that of [44, 58, 71]. We refer the reader to the Supplementary Material for further details.

Frozen Backbone: It is important to remark that $\Phi(\cdot)$ and \mathbf{W} are kept frozen during the training of our AP. While one could opt for finetuning the backbone while training the AP, the change in $\Phi(\cdot)$ and \mathbf{W} would lead to changing both the input space and the pseudolabels, leading to poorer performance. When minimising the AP objective (see below), \mathbf{W} is no longer learned through the error of the values $\hat{\mathbf{y}}$, and thus it can lead to produce pseudolabels that lie out of the space \mathcal{Y} , i.e. that are not valid. Recall that NPs rely on estimating $f(\mathbf{x}_t)$ for the target points, provided some observed values $\mathbf{y}_c = f(\mathbf{x}_c)$ for the context. In this paper the values \mathbf{y}_c are replaced by the predicted values $\hat{\mathbf{y}}_c$ provided by \mathbf{W} . If these values change over the course of training, and become non plausible values of f , the whole NP framework becomes invalid. We experimentally tried fine-tuning the backbone, and observed a poor performance.

¹For our experiments, we will also consider the task of Action Unit intensity estimation.

Overall, to our knowledge, we are the first to show in NP literature that pseudo-ground truth generated from an independent network can be effectively used for defining the labels for the context points used as input to the NP.

3.2.2 Context Selection

Our first and obvious option for context selection is uniform random sampling from the sequence, which would imply the strong assumption that the errors of the backbone are consistent. Considering that the backbone works in a per-frame basis, it is likely that some of the input features are more indicative of the context than others. In this paper, we propose to delve into the properties of the NPs to investigate whether we can infer which frames are most optimal for context representation.

Given that the number of context points can be any number between 1 and the length of the sequence, the choice of $\mathbf{r}_C = \mathbf{r}_c$ for a single context point c also results in a valid latent function. Building on this concept and the fact that $\sigma_C(\mathbf{r}_C)$ represents the uncertainty of the latent functions, we propose to select the context by first estimating the set of latent functions that each of the individual representations \mathbf{r}_c would propose, and select those with lowest uncertainty. That is, we propose to first estimate \mathbf{r}_c and their corresponding $\sigma_C(\mathbf{r}_c)$ for all the frames in a given sequence, and select the C with lowest uncertainty. The context set \mathcal{C} is then defined as the C frames for which $\sigma_c(\mathbf{r}_c)$ is the lowest. Note that this step is used to select the context points, only. Once context points are selected, they are passed through the aggregation function to compute \mathbf{r}_C in the standard way. This adds a negligible extra complexity to the framework (the values of \mathbf{r}_c do not need to be computed again). Finally, we emphasize that we apply context selection only at test time.

3.2.3 The Complete Architecture

The complete architecture is shown in Fig. 2. We use as input the features \mathbf{x}_t extracted from the pre-trained backbone, as well as the corresponding predictions. The encoder q_ϕ is a three-layer MLP that maps each of the input pairs $(\mathbf{x}_c, \hat{\mathbf{y}}_c)$ into a latent representation $\mathbf{r}_c \in \mathbb{R}^{128}$. The input pseudolabels $\hat{\mathbf{y}}_t$ are firstly upsampled to match the dimension of the feature representation. Once \mathbf{r}_C has been computed from the average of the selected context points, it is forwarded to a *latent encoder*, which comprises a common linear layer and two independent linear layers, to compute μ_C and σ_C , respectively. The decoder f_θ is another three-layer MLP that receives the individual features \mathbf{x}_t and the latent representation \mathbf{z} , and produces two outputs $\tilde{\mathbf{y}}_t^\mu$ and $\tilde{\mathbf{y}}_t^\sigma$, parameterising the output distribution. During training, $\mathbf{z} \sim \mathcal{N}(\mu_C, \sigma_C)$, and $\tilde{\mathbf{y}}_t \sim \mathcal{N}(\tilde{\mathbf{y}}_t^\mu, \tilde{\mathbf{y}}_t^\sigma)$. At test time, we use the predictive mean and set $\mathbf{z} = \mu_C$ and $\tilde{\mathbf{y}}_t = \tilde{\mathbf{y}}_t^\mu$.

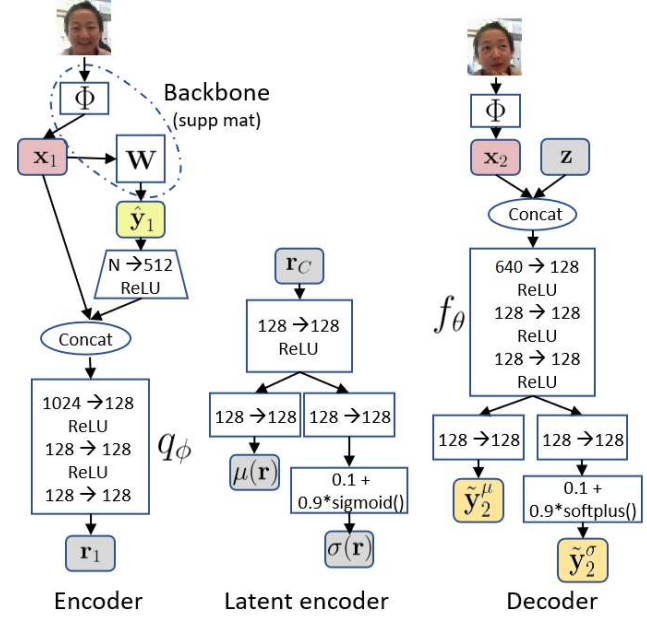


Figure 2. Architecture of AP. The encoder takes as input individual pairs $(\mathbf{x}_c, \hat{\mathbf{y}}_c)$, to produce the corresponding feature representations \mathbf{r}_c . The representations \mathbf{r}_c for the context are summarised into \mathbf{r}_C , and passed through a latent encoder that outputs the latent distribution, defined as a Gaussian parameterised by $\mu(\mathbf{r})$ and $\sigma(\mathbf{r})$. The decoder takes a feature representation \mathbf{x}_t from the target set, as well as a sample from the latent distribution, and produces an output Gaussian distribution parameterised by the outputs $(\tilde{\mathbf{y}}_t^\mu, \tilde{\mathbf{y}}_t^\sigma)$ of the decoder. The numbers represent the dimensions of the linear layers.

3.2.4 Training Losses

Learning in our model comprises learning the parameters of the encoder q_ϕ and the decoder f_θ . For a given set of annotated sequences of varying length, a random subset of frames are used as context, and the learning is the task of reconstructing the labels for the target points. In the NPs, the goal is to maximise the evidence lower-bound (ELBO) of $\log p(\mathbf{y}_T | \mathbf{x}_T, \mathbf{x}_C, \mathbf{y}_C)$. Because Valence and Arousal take values in $[-1, 1]$, we also consider that the labels \mathbf{y} are normally distributed around the origin, and propose a regularisation term for the output distribution as $\mathcal{L}_{reg} = D_{KL}(\mathcal{N}(\sum \tilde{\mathbf{y}}_t^\mu, \sum \tilde{\mathbf{y}}_t^\sigma) \| \mathcal{N}(0, 1))$.

The full objective to minimise is:

$$\mathcal{L} = \mathcal{L}_{nll} + \lambda_{kl} \mathcal{L}_{kl} + \lambda_{reg} \mathcal{L}_{reg}, \quad (1)$$

with $\mathcal{L}_{nll} = -\mathbb{E}_{q_\phi(\mathbf{z}|\mathcal{C})}[\log p(\mathbf{y}_T | \mathbf{x}_T, \mathbf{z})]$ and $\mathcal{L}_{kl} = D_{KL}(q_\phi(\mathbf{z}|\mathcal{T}) \| q_\phi(\mathbf{z}|\mathcal{C}))$. The optimisation can be efficiently done through the reparametrization trick [23]. In this paper λ_{kl} and λ_{reg} are set to 1.

3.3. Comparison with other methods

It is important to remark that our model comes with the following benefits: a) it allows to automatically infer the global context; this adds a new key to existing Neural Processes where context is manually given, b) it allows to model the uncertainty on the estimated context, and c) it allows to draw coherent samples; note that while the decoder works on a per-frame basis, the fact that it is conditioned on a global latent variable whose value can be drawn from a latent distribution, makes all of the samples be dependent on it [6].

Vs. RNN: The task of temporal regression could be approached by using a parametric form for f , by means of an RNN. However, this assumes that the variables \mathbf{y}_t are drawn from Markov chains, which in practice disregard long-term dependencies needed for temporal context modelling. Furthermore, it can be shown [15] that the joint distribution $p(\mathbf{y}_T)$ is exchangeable, i.e. it is permutation invariant. This offers a new insight into temporal modelling that differentiates from RNNs that condition a current observation on the previous hidden state, i.e. where order matters. We argue that, provided a given static image, the corresponding labels should not be position-dependent. If we shuffle a given sequence and pick up a frame at random, the corresponding annotation is not expected to change. Our approach seeks an orderless modelling of sequences, by looking at them as a whole through the global latent variable \mathbf{z} .

Vs. Self-attention: Self-attention models [64] can also model long-term dependencies using pair-wise feature similarities. On the contrary, we choose NPs to estimate f by sampling from a global latent random variable. This allows for a more direct and holistic modelling of the temporal context. Furthermore, unlike self-attention models, in NPs both input signal \mathbf{x}_c , and labels \mathbf{y}_c are used to learn context.

4. Experimental Settings

4.1. Datasets and Performance Metrics

We validate our approach for emotion recognition on **SEWA** [32] and **Aff-Wild2** [24, 27, 28, 26, 29, 25, 72]. We also validate our approach for the task of Action Unit intensity estimation on **BP4D** [73, 63] and **DISFA** [40].

SEWA [32] is a large video dataset for affect estimation in-the-wild. It comprises over 2,000 minutes of video data annotated with valence and arousal values. It contains subjects from 6 different cultures. We follow the evaluation protocol of [32, 31], and apply a 8:1:1 dataset partition².

AffWild2 [25] is a large-scale video dataset divided in training, validation, and test partitions. Because the labels for the latter are not made publicly available, we use the training and validation only for our evaluation purposes.

²The same partitions were kindly provided by the data owners.

BP4D [73, 63] contains videos collected from 41 subjects, performing 8 different tasks. It is the main corpus of FERA 2015 challenge [63]. We use the official train (168 videos), validation (160 videos), and test (159 videos) partitions, which contain annotations for 5 different AUs.

DISFA [40] contains 27 videos of 27 different subjects, with 4,844 frames each. It contains annotations for 12 AUs. We follow a 3-fold cross-validation protocol with 18 subjects used for training and 9 subjects used for testing.

Performance Metrics: For Valence and Arousal, we follow the AVEC Challenge series [50] and report the Concordance Correlation Coefficient [34]. For Action Units, we follow the ranking criteria used in FERA challenges [63], and report the Intra Class Correlation (ICC [55]). A full definition of these metrics can be found in the Supp. Material.

4.2. Implementation details

All experiments are developed in PyTorch [48]. First, we pre-train the backbone architecture Φ and \mathbf{W} on a per-frame basis for each corresponding dataset (see Supplementary Material). We then use as input the features \mathbf{x}_t extracted from the pre-trained backbone, as well as the corresponding predictions. The dimension of input features \mathbf{x}_t is 512.

Valence and Arousal: To train the AP for Valence and Arousal, we randomly sample sequences between 35 and 70 frames from the corresponding training set. The range 35 – 70 can be seen as an augmentation parameter: using a shorter range yielded worse results due to lack of generalisation, whereas no improvement was observed when using a longer range of sequences. The number of context points is randomly chosen from 3 to the length of the given sequence. During training, the target labels are the ground-truth annotations. For the context labels, we use the ground-truth annotations or the predictions from \mathbf{W} with probability 0.5. The batch size is 16 videos, and the learning rate and weight decay are set to 0.00025 and 0.0001, respectively. The AP is trained for 25 epochs, each of 1000 iterations, using Adam [22] with $(\beta_1, \beta_2) = (0.9, 0.999)$. We use a Cosine Annealing strategy for the learning rate [38]. At test time, we use a fixed sequence length of 70 frames.

Action Units: We used the same setting as for Valence and Arousal, with the following differences: the batch size is set to 6 videos, and the learning rate and weight decay are set to 0.0001 and 0.0005, respectively.

5. Ablation Studies

We use the validation partition of SEWA to perform our ablation studies. We study the effect of our proposed context selection method, the contribution of the pseudo labels, the effect of the different losses used for the AP training, and the effect of the latent variable. Finally, we explore a different alternative to the AP configuration [35]. We study in Sec. 6 the complexity of our approach.

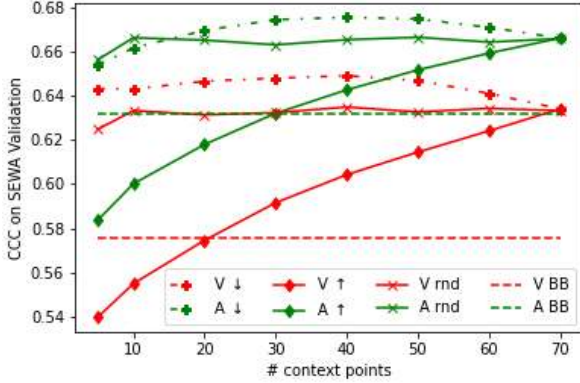


Figure 3. Accuracy on SEWA validation w.r.t. the number of context points and context selection method. The results in red correspond to the accuracy obtained for Valence (V), and those in green for Arousal (A). ↓ (dashed and dotted lines) represents the results obtained by selecting the context frames according to the lowest values of $\|\sigma_c\|$, and ↑ (solid lines with diamond markers) represents the results obtained by selecting the frames according to the largest values of $\|\sigma_c\|$ and, *rnd* (solid lines with cross markers) represents the results obtained by choosing the frames randomly. The dashed horizontal lines represent the results of the backbone.

Effect of context selection: Firstly, we investigate the effect of the number of selected context points and the selection method proposed in Sec. 3.2. For a given trained model, we study the effect of choosing the context points (a) randomly, (b) according to the lowest values of $\|\sigma_c\|$, and (c) according to the largest values of $\|\sigma_c\|$. Fig. 3 shows that the best results are attained when choosing 40 context points according to the lowest value of $\|\sigma_c\|$. We fix for the remaining experiments the number of context points to 40.

Impact of task-aware context modelling: One of the contributions of our method is that of task-aware context modelling in the form of *pseudo-ground truth labels*. We conducted an experiment where the AP model is trained without the predictions provided by **W**. The architecture, training, losses, learning rate and weight decay remain the same as for the main AP. The results, reported in Table 1 as $AP_{w/o \mathbf{W}}$, show that the performance drops significantly without the aid of the noisy task-specific predictions.

Effect of global latent random variable: We firstly investigate the effect of the latent variable on regressing coherent functions. Recall that our proposed approach consists of a probabilistic regression of coherent functions, i.e. it produces a set of functional proposals. An alternative to this approach is to use a non-probabilistic approach where the latent representation does not parameterise a latent random distribution, but is rather *deterministic*. The architecture of a deterministic AP is similar to that shown in Fig. 2,

Model	Valence	Arousal	Mean
Backbone	0.576	0.632	0.604
AP_{nll}	0.645	0.676	0.660
AP_{nll+kl}	0.649	0.669	0.659
$AP_{nll+reg}$	0.649	0.676	0.662
$AP_{nll+reg+kl}$	0.643	0.628	0.635
Deterministic AP	0.650	0.576	0.613
AP + Det.	0.649	0.658	0.652
$AP_{w/o \mathbf{W}}$	0.600	0.454	0.527
AP + Det. + Att	0.662	0.672	0.667

Table 1. Ablation studies on the validation set of SEWA.

although with a single-branch latent encoder and no sampling. With a deterministic AP, one dispenses with the coherent sampling i.e. the dependence among the target outputs is lost, and assumed to be i.i.d. The learning burden is put only into capturing the observation variance \tilde{y}^σ to account for the output uncertainty (i.e. it is probabilistic in the output, but not in the latent representation). The results of the deterministic AP are shown in Table 1 (Deterministic AP). We observe that the performance degrades substantially, in particular regarding the Arousal predictions. This illustrates the importance of the latent variable to model the underlying variability for the joint prediction of Valence and Arousal. Finally, we also investigated the dimensionality of **z**, observing little effect when having more dimensions.

Effect of losses: The training of the APs is carried out by minimising the objective shown in Eqn. 1. Recent findings [13] suggest that training an NP using only a maximum likelihood approach (i.e. by minimising the negative log-likelihood \mathcal{L}_{nll}) can improve performance and avoid the practical issues of using amortised variational inference (i.e. $\mathcal{L}_{nll} + \mathcal{L}_{kl}$). We studied the effect of each case along with our proposed loss. The results are shown in Table 1. We refer with AP_{nll} to the maximum likelihood approach, with AP_{nll+kl} to the one that minimises $\mathcal{L}_{nll} + \mathcal{L}_{kl}$, with $AP_{nll+reg}$ to the one that minimises $\mathcal{L}_{nll} + \mathcal{L}_{reg}$, and with $AP_{nll+reg+kl}$ to the one that minimises the full loss in Eqn. 1. The numerical results do not shed light into whether one is better than the other. We observe that a combination of the maximum likelihood approach with our proposed regularisation slightly improves the numerical results, although we acknowledge these not to be significant enough. If we use the three losses together we observe a performance degradation. These results align with recent findings that advocate for a further investigation of which of the losses is a better objective for Neural Processes [13].

Different network configurations: We also study different network configurations. We first evaluate the perfor-

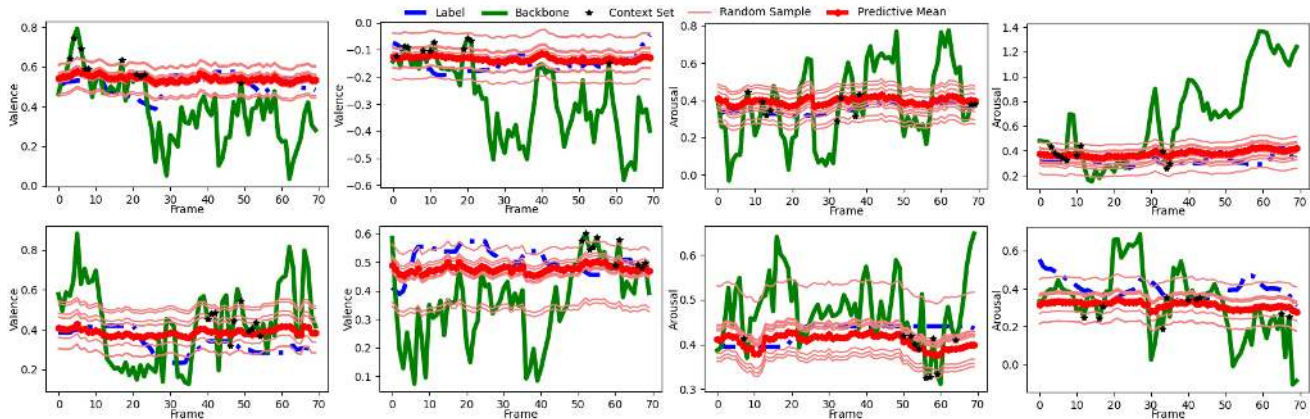


Figure 4. Examples of modelling Valence and Arousal on SEWA dataset, when using as context only 10 points, selected according to the lowest values of $\|\sigma_c\|$. Blue line represents the ground-truth labels. The green line represents the predictions from the backbone. The black dots represent the frames selected as context using σ_c . The light red curves represent different realisations of $f(\cdot, \mathbf{z})$ for 10 different (sampled) \mathbf{z} , and the red curve represents the process returned by the predictive mean μ_C . The variation in the functions corresponds to the global latent uncertainty given the selected context.

mance of using a combined latent and deterministic layer denoted as AP + Det.. The performance of the AP + Det. is on par with that of the single AP. We then explore the use of Attention in the AP + Det. where self-attention is used to replace the aggregation of the individual representations (attention is permutation-invariant). Considering the trade-off between complexity and accuracy, we opt for keeping the simplest architecture, described in Fig. 2.

Qualitative results: We show a conceptual demonstration of the AP capabilities for both context selection and functional proposals in Fig. 4. The blue lines represent the ground-truth annotations for different segments of 70 frames. The green lines correspond to the individual predictions \hat{y}_t given by the backbone \mathbf{W} . The black dots represent the frames selected as context according to the lowest values of $\|\sigma_c\|$. For the sake of clarity, we used only 10 points to select context. The light red curves represent different realisations of the functions $f_\theta(\cdot, \mathbf{z})$ given by different samples of \mathbf{z} , with the function corresponding to μ_C represented in dark red. We can see that the proposed functions are able to better approximate the sequential predictions from the noisy individual representations, leading to much more stable predictions. We can observe that the model tends to select frames that correspond to lower error, indicating how the individual latent representations can represent their confidence in the context.

6. Comparison with state-of-the-art

We compare the results of our models w.r.t. the most recent works in emotion and Action Unit intensity estimation. The bulk of our results are shown in Tables 2 and 3 for Valence and Arousal, and Tables 4 and 5 for Action Units (MSE is also reported in the Supplementary Material).

In-house baselines: We trained a set of baselines based on GRU [3] and Attention [64], using, for each case, the same backbone as the one used in AP to provide the input features. We also optimised for each task and model the sequence length (herein denoted by L). For Valence and Arousal, we used a 4-layer Convolutional GRU with 128-d hidden layers, and $L = 30$. For Action Units, we trained a 2-layer Bidirectional GRU with 256-d hidden states, with $L = 60$ frames. For Attention, we used, for both tasks, an architecture based on the Transformer encoder: we used a 2-layer MultiHead Self-Attention encoder with 32 learnable heads, each of 512 dimensions. The sequence lengths are $L = 30$ and $L = 40$ frames for Valence and Arousal and Action Units, respectively. The results in Tables 2-5 indicate that our AP outperforms both recurrent and attention-based models. We also experimented with using as input both the features and the labels, as well as finetuning the backbone, observing a performance degradation.

Different backbones: We validate that our AP can be used with other backbone choices. On SEWA, we first train a ResNet-50, yielding a CCC of 0.567, and then train our AP on top yielding a CCC of 0.655. On AffWild-2, we train our AP using as backbone the publicly available model of [33]: the backbone of [33] reports a CCC score on the AffWild-2 test set of 0.43, whereas our AP yields a CCC score of 0.47.

Analysis: On SEWA, we compare against Mitekova et al. [42] and Kossaifi et al. [31]. Notably, our backbone is already on par with the state of the art results of [31]. We can observe that our method based on temporal context offers the largest improvement. The same behaviour is observed

Model	Valence	Arousal	Mean
Mitekova et al. [42]	0.44	0.39	0.42
Kossafi et al. [31]	0.75	0.52	0.64
Ours backbone	0.69	0.61	0.65
ConvGRU [3] [†]	0.72	0.62	0.67
Self-Attn [64] [†]	0.70	0.61	0.65
Ours AP	0.75	0.64	0.69

Table 2. Comparison against state of the art methods on the SEWA test partition. [†] denotes in-house implementation.

Model	Valence	Arousal	Mean
Ours Backbone	0.398	0.445	0.422
ConvGRU [3] [†]	0.398	0.503	0.450
Self-Attn [64] [†]	0.419	0.505	0.462
Ours AP	0.438	0.498	0.468

Table 3. Comparison of different methods on the validation set of AffWild2. [†] denotes in-house implementation.

AU	1	2	4	5	6	9	12	15	17	20	25	26	Avg.
VGP-AE [10]	0.48	0.47	0.62	0.19	0.50	0.42	0.80	0.19	0.36	0.15	0.84	0.53	0.46
2DC [37]	0.70	0.55	0.69	0.05	0.59	0.57	0.88	0.32	0.10	0.08	0.90	0.50	0.50
HR [44]	0.56	0.52	0.75	0.42	0.51	0.55	0.82	0.55	0.37	0.21	0.93	0.62	0.57
Ours backbone	0.25	0.15	0.76	0.77	0.50	0.62	0.82	0.44	0.58	0.22	0.93	0.69	0.56
BiGRU [3] [†]	0.23	0.13	0.77	0.70	0.53	0.64	0.82	0.42	0.58	0.25	0.92	0.69	0.56
Self-Attn [64] [†]	0.30	0.17	0.78	0.76	0.51	0.64	0.82	0.46	0.59	0.17	0.92	0.68	0.57
Ours AP	0.35	0.19	0.78	0.73	0.52	0.65	0.81	0.49	0.61	0.28	0.92	0.67	0.58

Table 4. Comparison against state of the art methods on the DISFA database (in ICC values). [†] denotes in-house evaluation

AU	6	10	12	14	17	Avg.
CDL [1]	0.69	0.73	0.83	0.50	0.37	0.62
ISIR [43]	0.79	0.80	0.86	0.71	0.44	0.72
HR [44]	0.82	0.82	0.80	0.71	0.50	0.73
Ours backbone	0.76	0.74	0.82	0.62	0.39	0.67
BiGRU [3] [†]	0.78	0.76	0.83	0.62	0.50	0.70
Self-Attn [64] [†]	0.80	0.78	0.79	0.60	0.42	0.68
Ours AP	0.82	0.80	0.86	0.69	0.51	0.74

Table 5. Results on the test partition of BP4D dataset (in ICC values). [†] denotes in-house evaluation

on AffWild2, where our AP significantly improves the Valence CCC. On DISFA, our backbone is again on par with the state of the art methods VGP-AE [10] and 2DC [37]. However, our AP brings the largest improvement setting a new state-of-the-art ICC of 0.58, which improves over the Heatmap Regression method of Ntinou et al. [44]. Finally, we compare our method against the winner and runner-up methods of FERA 2015, ISIR [43] and CDL [1], respectively, as well as against the Heatmap Regression method of [44]. Our method sets up a new state-of-the-art result with an ICC score of 0.74. We include a set of qualitative results in the Supplementary Material.

Complexity: Our AP adds very little complexity to the pre-trained backbone. The configuration shown in Fig. 2 comprises only $\sim 330K$ parameters, and requires only ~ 20 MFLOPs for a given sequence of $L = 70$ frames³. In com-

³We measured the complexity using the Thop library for PyTorch <https://github.com/Lyken17/pytorch-OpCounter>

parison, the ConvGRU model has $\sim 789K$ parameters, and requires ~ 758 MFLOPs for $L = 30$ frames. The PyTorch built-in 2-layer BiGRU has $\sim 1.2M$ parameters, and needs ~ 71 MFLOPs for $L = 60$ frames. The Multi-Head attention model has $\sim 4.2M$ parameters, and requires ~ 125 MFLOPs for $L = 30$ frames.

7. Conclusions

We proposed a novel temporal context learning approach to facial affect recognition, by modelling the stochastic nature of affect labels. To model uncertainty in the temporal context of the affective signals, our proposed method, called Affective Processes, builds on Neural Processes to learn a distribution of functions. Our method addresses a key limitation of the NPs — that of requiring ground-truth labels at inference time — by proposing the use of pseudo-labels provided by a pre-trained backbone and by using a novel method for automatic context selection. We believe that the key elements of our method pave the way to apply NPs to large-scale supervised learning problems which are inherently stochastic and for which deterministic temporal models like RNNs and self-attention could fall short.

Acknowledgements

The work by Michel Valstar was part-funded by the National Institute for Health Research. The views represented are the views of the authors alone and do not necessarily represent the views of the Department of Health in England, NHS, or the National Institute for Health Research.

References

- [1] Tadas Baltrušaitis, Marwa Mahmoud, and Peter Robinson. Cross-dataset learning and person-specific normalisation for automatic action unit detection. In *IEEE Conf. and Workshops on Auto. Face and Gesture Recog.*, volume 6, pages 1–6. IEEE, 2015.
- [2] Lisa Feldman Barrett, Ralph Adolphs, Stacy Marsella, Aleix M Martinez, and Seth D Pollak. Emotional expressions reconsidered: Challenges to inferring emotion from human facial movements. *Psychological science in the public interest*, 20(1):1–68, 2019.
- [3] Kyunghyun Cho, Bart Van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014.
- [4] Junyoung Chung, Kyle Kastner, Laurent Dinh, Kratarth Goel, Aaron C Courville, and Yoshua Bengio. A recurrent latent variable model for sequential data. In *Adv. Neural Inform. Process. Syst.*, pages 2980–2988, 2015.
- [5] Arnaud Dapogny, Kevin Bailly, and Severine Dubuisson. Pairwise conditional random forests for facial expression recognition. In *Int. Conf. Comput. Vis.*, pages 3783–3791, 2015.
- [6] Bruno De Finetti. La prévision: ses lois logiques, ses sources subjectives. *Annales de l’institut Henri Poincaré*, 7(1):1–68, 1937.
- [7] Didan Deng, Zhaokang Chen, and Bertram E Shi. Multitask emotion recognition with incomplete labels. In *IEEE Conf. and Workshops on Auto. Face and Gesture Recog.*, pages 828–835, 2020.
- [8] Zhengyin Du, Suowei Wu, Di Huang, Weixin Li, and Yunhong Wang. Spatio-temporal encoder-decoder fully convolutional network for video-based dimensional emotion recognition. *IEEE Trans. on Affect. Comput.*, 2019.
- [9] Paul Ekman. Facial action coding system. 1977.
- [10] Stefanos Eleftheriadis, Ognjen Rudovic, Marc Peter Deisenroth, and Maja Pantic. Variational gaussian process auto-encoder for ordinal prediction of facial action units. In *ACCV*, pages 154–170. Springer, 2016.
- [11] Stefanos Eleftheriadis, Ognjen Rudovic, Marc Peter Deisenroth, and Maja Pantic. Gaussian process domain experts for modeling of facial affect. *IEEE Trans. Image Process.*, 26(10):4697–4711, 2017.
- [12] Stefanos Eleftheriadis, Ognjen Rudovic, and Maja Pantic. Discriminative shared gaussian processes for multiview and view-invariant facial expression recognition. *IEEE Trans. Image Process.*, 24(1):189–204, 2014.
- [13] Andrew YK Foong, Wessel P Bruinsma, Jonathan Gordon, Yann Dubois, James Requeima, and Richard E Turner. Meta-learning stationary stochastic process prediction with convolutional neural processes. In *Adv. Neural Inform. Process. Syst.*, 2020.
- [14] Marta Garnelo, Dan Rosenbaum, Christopher Maddison, Tiago Ramalho, David Saxton, Murray Shanahan, Yee Whye Teh, Danilo Rezende, and SM Ali Eslami. Conditional neural processes. In *International Conference on Machine Learning*, pages 1704–1713, 2018.
- [15] Marta Garnelo, Jonathan Schwarz, Dan Rosenbaum, Fabio Viola, Danilo J. Rezende, S. M. Ali Eslami, and Yee Whye Teh. Neural processes. In *ICML Workshop on Theoretical Foundations and Applications of Deep Generative Models*, 2018.
- [16] Jonathan Gordon, Wessel P Bruinsma, Andrew YK Foong, James Requeima, Yann Dubois, and Richard E Turner. Convolutional conditional neural processes. In *International Conference on Learning Representations*, 2020.
- [17] Jian Huang, Ya Li, Jianhua Tao, Zheng Lian, and Jiangyan Yi. End-to-end continuous emotion recognition from video using 3d convlstm networks. In *ICASSP*, pages 6837–6841. IEEE, 2018.
- [18] Jian Huang, Jianhua Tao, Bin Liu, Zhen Lian, and Mingyue Niu. Efficient modeling of long temporal contexts for continuous emotion recognition. In *Int. Conf. Affect. Comput. and Intel. Inter.*, pages 185–191. IEEE, 2019.
- [19] Shashank Jaiswal and Michel Valstar. Deep learning the dynamic appearance and shape of facial action units. In *IEEE Winter Conf. on Appl. of Comput. Vis.*, pages 1–8. IEEE, 2016.
- [20] Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Adv. Neural Inform. Process. Syst.*, pages 5574–5584. Curran Associates, Inc., 2017.
- [21] Hyunjik Kim, Andriy Mnih, Jonathan Schwarz, Marta Garnelo, Ali Eslami, Dan Rosenbaum, Oriol Vinyals, and Yee Whye Teh. Attentive neural processes. In *International Conference on Learning Representations*, 2018.
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [23] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [24] Dimitrios Kollias, Mihalis A Nicolaou, Irene Kotsia, Guoying Zhao, and Stefanos Zafeiriou. Recognition of affect in the wild using deep neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 1972–1979. IEEE, 2017.
- [25] Dimitrios Kollias, Attila Schulc, Elnar Hajiyev, and Stefanos Zafeiriou. Analysing affective behavior in the first abaw 2020 competition. *arXiv preprint arXiv:2001.11409*, 2020.
- [26] Dimitrios Kollias, Panagiotis Tzirakis, Mihalis A Nicolaou, Athanasios Papaioannou, Guoying Zhao, Björn Schuller, Irene Kotsia, and Stefanos Zafeiriou. Deep affect prediction in-the-wild: Aff-wild database and challenge, deep architectures, and beyond. *Int. J. Comput. Vis.*, pages 1–23, 2019.
- [27] Dimitrios Kollias and Stefanos Zafeiriou. Aff-wild2: Extending the aff-wild database for affect recognition. *arXiv preprint arXiv:1811.07770*, 2018.
- [28] Dimitrios Kollias and Stefanos Zafeiriou. A multi-task learning & generation framework: Valence-arousal, action units & primary expressions. *arXiv preprint arXiv:1811.07771*, 2018.

- [29] Dimitrios Kollias and Stefanos Zafeiriou. Expression, affect, action unit recognition: Aff-wild2, multi-task learning and arcface. *arXiv preprint arXiv:1910.04855*, 2019.
- [30] Dimitrios Kollias and Stefanos P Zafeiriou. Exploiting multi-cnn features in cnn-rnn based dimensional emotion recognition on the omg in-the-wild dataset. *IEEE Trans. on Affect. Comput.*, 2020.
- [31] Jean Kossaifi, Antoine Toisoul, Adrian Bulat, Yannis Panagakis, Timothy M. Hospedales, and Maja Pantic. Factorized higher-order cnns with an application to spatio-temporal emotion estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2020.
- [32] Jean Kossaifi, Robert Walecki, Yannis Panagakis, Jie Shen, Maximilian Schmitt, Fabien Ringeval, Jing Han, Vedhas Pandit, Antoine Toisoul, Bjoern W Schuller, et al. Sewa db: A rich database for audio-visual emotion and sentiment research in the wild. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019.
- [33] F Kuhnke, L Rumberg, and J Ostermann. Two-stream aural-visual affect analysis in the wild. In *IEEE Conf. and Workshops on Auto. Face and Gesture Recog.*, pages 366–371, 2020.
- [34] I Lawrence and Kuei Lin. A concordance correlation coefficient to evaluate reproducibility. *Biometrics*, pages 255–268, 1989.
- [35] Tuan Anh Le, Hyunjik Kim, Marta Garnelo, Dan Rosenbaum, Jonathan Schwarz, and Yee Whye Teh. Empirical evaluation of neural process objectives. In *Adv. Neural Inform. Process. Syst. Worksh.*, 2018.
- [36] Juho Lee, Yoonho Lee, Jungtaek Kim, Eunho Yang, Sung Ju Hwang, and Yee Whye Teh. Bootstrapping neural processes. *Adv. Neural Inform. Process. Syst.*, 33, 2020.
- [37] Dieu Linh Tran, Robert Walecki, Stefanos Eleftheriadis, Bjorn Schuller, Maja Pantic, et al. Deepcoder: Semi-parametric variational autoencoders for automatic facial action coding. In *Int. Conf. Comput. Vis.*, pages 3190–3199, 2017.
- [38] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016.
- [39] Brais Martinez, Michel F Valstar, Bihan Jiang, and Maja Pantic. Automatic analysis of facial actions: A survey. *IEEE Trans. on Affect. Comput.*, 2017.
- [40] S Mohammad Mavadati, Mohammad H Mahoor, Kevin Bartlett, Philip Trinh, and Jeffrey F Cohn. Disfa: A spontaneous facial action intensity database. *IEEE Trans. on Affect. Comput.*, 4(2):151–160, 2013.
- [41] Hongying Meng, Nadia Bianchi-Berthouze, Yangdong Deng, Jinkuang Cheng, and John P Cosmas. Time-delay neural network for continuous emotional dimension prediction from facial expression sequences. *IEEE transactions on cybernetics*, 46(4):916–929, 2015.
- [42] Anna Mitenkova, Jean Kossaifi, Yannis Panagakis, and Maja Pantic. Valence and arousal estimation in-the-wild with tensor methods. In *IEEE Conf. and Workshops on Auto. Face and Gesture Recog.*, pages 1–7. IEEE, 2019.
- [43] Jeremie Nicolle, Kevin Bailly, and Mohamed Chetouani. Facial action unit intensity prediction via hard multi-task metric learning for kernel regression. In *IEEE Conf. and Workshops on Auto. Face and Gesture Recog.*, volume 6, pages 1–6. IEEE, 2015.
- [44] I. Ntinou, E. Sanchez, A. Bulat, M. Valstar, and G. Tzimiropoulos. A transfer learning approach to heatmap regression for action unit intensity estimation. *IEEE Transactions on Affective Computing*, pages 1–1, 2021.
- [45] Desmond Ong, Zhengxuan Wu, Zhi-Xuan Tan, Marianne Reddan, Isabella Kahhale, Alison Mattek, and Jamil Zaki. Modeling emotion in complex stories: the stanford emotional narratives dataset. *IEEE Trans. on Affect. Comput.*, 2019.
- [46] Meshia Cédric Oveneke, Isabel Gonzalez, Valentin Enescu, Dongmei Jiang, and Hichem Sahli. Leveraging the bayesian filtering paradigm for vision-based facial affective state estimation. *IEEE Trans. on Affect. Comput.*, 9(4):463–477, 2017.
- [47] Meshia Cédric Oveneke, Yong Zhao, Ercheng Pei, Abel Diaz Berenguer, Dongmei Jiang, and Hichem Sahli. Leveraging the deep learning paradigm for continuous affect estimation from facial expressions. *IEEE Trans. on Affect. Comput.*, 2019.
- [48] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. In *Autodiff workshop - NeurIPS*, 2017.
- [49] Jonathan Posner, James A Russell, and Bradley S Peterson. The circumplex model of affect: An integrative approach to affective neuroscience, cognitive development, and psychopathology. *Development and psychopathology*, 17(3):715, 2005.
- [50] Fabien Ringeval, Björn Schuller, Michel Valstar, Nicholas Cummins, Roddy Cowie, Leili Tavabi, Maximilian Schmitt, Sina Alisamir, Shahin Amiriparian, Eva-Maria Messner, et al. Avec 2019 workshop and challenge: state-of-mind, detecting depression with ai, and cross-cultural affect recognition. In *Proceedings of the 9th International on Audio/Visual Emotion Challenge and Workshop*, pages 3–12, 2019.
- [51] Ognjen Rudovic, Vladimir Pavlovic, and Maja Pantic. Kernel conditional ordinal random fields for temporal segmentation of facial action units. In *Eur. Conf. Comput. Vis.*, pages 260–269. Springer, 2012.
- [52] Ognjen Rudovic, Vladimir Pavlovic, and Maja Pantic. Multi-output laplacian dynamic ordinal regression for facial expression recognition and intensity estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 2634–2641. IEEE, 2012.
- [53] Georgia Sandbach, Stefanos Zafeiriou, and Maja Pantic. Markov random field structures for facial action unit intensity estimation. In *Int. Conf. Comput. Vis. Worksh.*, pages 738–745, 2013.
- [54] Vidhyasaharan Sethu, Emily Mower Provost, Julien Epps, Carlos Busso, Nicholas Cummins, and Shrikanth Narayanan. The ambiguous world of emotion representation. *arXiv preprint arXiv:1909.00360*, 2019.

- [55] P.E. Shrout and J.L. Fleiss. Intraclass correlations: uses in assessing rater reliability. *Psychological bulletin*, 1979.
- [56] Gautam Singh, Jaesik Yoon, Youngsung Son, and Sungjin Ahn. Sequential neural processes. In *Adv. Neural Inform. Process. Syst.*, pages 10254–10264, 2019.
- [57] Mani Kumar Tellamekala and Michel Valstar. Temporally coherent visual representations for dimensional affect recognition. In *2019 8th International Conference on Affective Computing and Intelligent Interaction (ACII)*, pages 1–7. IEEE, 2019.
- [58] Antoine Toisoul, Jean Kossaifi, Adrian Bulat, Georgios Tzimiropoulos, and Maja Pantic. Estimation of continuous valence and arousal levels from faces in naturalistic conditions. *Nature Machine Intelligence*, 2021.
- [59] Martin Trapp, Robert Peharz, Franz Pernkopf, and Carl Edward Rasmussen. Deep structured mixtures of gaussian processes. In *International Conference on Artificial Intelligence and Statistics*, pages 2251–2261. PMLR, 2020.
- [60] Volker Tresp. Mixtures of gaussian processes. In *Adv. Neural Inform. Process. Syst.*, pages 654–660, 2001.
- [61] Michel Valstar and Maja Pantic. Fully automatic facial action unit detection and temporal analysis. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.*, pages 149–149. IEEE, 2006.
- [62] Michel François Valstar. Timing is everything: A spatio-temporal approach to the analysis of facial actions. *PhD Dissertation*, 2008.
- [63] Michel F Valstar, Timur Almaev, Jeffrey M Girard, Gary McKeown, Marc Mehu, Lijun Yin, Maja Pantic, and Jeffrey F Cohn. Fera 2015-second facial expression recognition and analysis challenge. In *IEEE Conf. and Workshops on Auto. Face and Gesture Recog.*, volume 6, pages 1–8. IEEE, 2015.
- [64] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Adv. Neural Inform. Process. Syst.*, pages 5998–6008, 2017.
- [65] Robert Walecki, Vladimir Pavlovic, Björn Schuller, Maja Pantic, et al. Deep structured learning for facial action unit intensity estimation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3405–3414, 2017.
- [66] Robert Walecki, Ognjen Rudovic, Vladimir Pavlovic, and Maja Pantic. Variable-state latent conditional random fields for facial expression recognition and action unit detection. In *IEEE Conf. and Workshops on Auto. Face and Gesture Recog.*, volume 1, pages 1–8. IEEE, 2015.
- [67] Ke Wang, Geoff Pleiss, Jacob Gardner, Stephen Tyree, Kilian Q Weinberger, and Andrew Gordon Wilson. Exact gaussian processes on a million data points. In *Adv. Neural Inform. Process. Syst.*, pages 14648–14659, 2019.
- [68] Shangfei Wang, Zhuangqiang Zheng, Shi Yin, Jiajia Yang, and Qiang Ji. A novel dynamic model capturing spatial and temporal patterns for facial expression analysis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 2019.
- [69] Shu-Hui Wang and Chiou-Ting Hsu. Ast-net: An attribute-based siamese temporal network for real-time emotion recognition. In *Brit. Mach. Vis. Conf.*, 2017.
- [70] Christopher KI Williams and Carl Edward Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA, 2006.
- [71] Jing Yang, Adrian Bulat, and Georgios Tzimiropoulos. Fan-face: a simple orthogonal improvement to deep face recognition. In *AAAI Conference on Artificial Intelligence*, 2020.
- [72] Stefanos Zafeiriou, Dimitrios Kollias, Mihalis A Nicolaou, Athanasios Papaioannou, Guoying Zhao, and Irene Kotzia. Aff-wild: Valence and arousal ‘in-the-wild’ challenge. In *IEEE Conf. Comput. Vis. Pattern Recog. Worksh.* IEEE, 2017.
- [73] Xing Zhang, Lijun Yin, Jeffrey F Cohn, Shaun Canavan, Michael Reale, Andy Horowitz, Peng Liu, and Jeffrey M Girard. Bp4d-spontaneous: a high-resolution spontaneous 3d dynamic facial expression database. *Image and Vision Computing*, 32(10):692–706, 2014.
- [74] Y Zhang, R Huang, J Zeng, and S Shan. M3f: Multi-modal continuous valence-arousal estimation in the wild. In *IEEE Conf. and Workshops on Auto. Face and Gesture Recog.*, pages 617–621.
- [75] Yong Zhang, Haiyong Jiang, Baoyuan Wu, Yanbo Fan, and Qiang Ji. Context-aware feature and label fusion for facial action unit intensity estimation with partially labeled data. In *Int. Conf. Comput. Vis.*, pages 733–742, 2019.