

# Less is More: Facial Landmarks can Recognize a Spontaneous Smile

Md. Tahrim Faroque Tushar\*<sup>1</sup>

tahrim.tushar@northsouth.edu

Yan Yang\*<sup>2,3</sup>

u6169130@anu.edu.au

Md Zakir Hossain<sup>2,3,4†</sup>

zakir.hossain@csiro.au

Sheikh Motahar Naim<sup>5</sup>

sheinaim@amazon.com

Nabeel Mohammed<sup>1</sup>

nabeel.mohammed@northsouth.edu

Shafin Rahman<sup>1†</sup>

shafin.rahman@northsouth.edu

<sup>1</sup> Department of Electrical and

Computer Engineering,

North South University, Bangladesh

<sup>2</sup> Biological Data Science Institute,

The Australian National University,

Canberra, Australia

<sup>3</sup> CSIRO Agriculture & Food,

Canberra, Australia

<sup>4</sup> CSIRO Machine Learning & Artificial

Intelligence Future Science Platform,

Canberra, Australia

<sup>5</sup> Amazon Web Services

## Abstract

Smile veracity classification is a task of interpreting social interactions. Broadly, it distinguishes between spontaneous and posed smiles. Previous approaches used hand-engineered features from facial landmarks or considered raw smile videos in an end-to-end manner to perform smile classification tasks. Feature-based methods require intervention from human experts on feature engineering and heavy pre-processing steps. On the contrary, raw smile video inputs fed into end-to-end models bring more automation to the process with the cost of considering many redundant facial features (beyond landmark locations) that are mainly irrelevant to smile veracity classification. It remains unclear to establish discriminative features from landmarks in an end-to-end manner. We present a **MeshSmileNet** framework, a transformer architecture, to address the above limitations. To eliminate redundant facial features, our landmarks input is extracted from Attention Mesh, a pre-trained landmark detector. Again, to discover discriminative features, we consider the *relativity* and *trajectory* of the landmarks. For the *relativity*, we aggregate facial landmark that conceptually formats a curve at each frame to establish local spatial features. For the *trajectory*, we estimate the movements of landmark composed features across time by self-attention mechanism, which captures pairwise dependency on the trajectory of the same landmark. This idea allows us to achieve state-of-the-art performances on UVA-NEMO, BBC, MMI Facial Expression, and SPOS datasets.

# 1 Introduction

Facial expression can reveal the emotional state of a person. One of the most common facial expressions representing happiness is the smile. More specifically, a genuine and spontaneous smile can only represent happiness. However, when the smile becomes deliberate or posed, it may signify anxiety, embarrassment, politeness, fear, affirmation and sarcasm [1]. Due to the social impact of the smile, many studies have been conducted to recognize a genuine/real/spontaneous from a posed/deliberate/fake smile [6, 23, 29, 31]. Such research aims to improve human-computer interaction by allowing computers/robots to perceive users' emotional states. It can be helpful in the fields of neuroscience by studying the influence of subtle movements of eye and lip muscles in human emotions, psychology by investigating the genuine feeling of patients with mental illness, autism, or verbally handicapped people, criminology by aiding polygraph test for police interrogation and so on. In this paper, we study the spontaneous and posed smile expression and present a **MeshSmileNet** framework that can work end-to-end and fully-automatic manner.

Unlike traditional facial expression classification tasks, smile veracity classification is more challenging because of the difficulty of capturing and interpreting micro-facial movements in spatial and time dimensions. We identify several limitations of past methods. **(a) Redundant facial features.** Existing end-to-end learning approaches [29, 31] study the problem by considering video frames as inputs. As indicated by [6], the smile veracity mostly depends on the variation of key facial landmarks. For example, the pull of the zygomatic major and orbicularis oculi muscle leads to a spontaneous smile [10]. However, video frames contain a large volume of irrelevant facial features such as identity that is trivial to the recognition performance. **(b) Poor inductive bias.** The spatial location of the same feature may vary across time dimensions. For example, the lips corner may not be in the similar location across different frames. For some methods, [29, 31], it creates difficulty in tracking feature trajectories because they assume an inductive bias that discriminative features will be located in a similar position across smile videos. **(c) Manual feature engineering.** Some methods [6] manually craft features including eye angle, cheek pull speed, and so on after tracking the facial landmark of interest over the smile video. Others [28, 29] explore Local Binary Pattern features for facial regions of interest. This process brings the intervention of human experts with strong domain knowledge and limits the feature generality to other similar tasks. **(d) Smile phase segmentation.** Manual feature engineering-based approach [6, 28, 29] usually requires the segmentation of the smile phase (onset, apex, offset) to underneath the recognition performance. It exacerbates the cumbersome recognition processes of these feature engineering-based approaches. **(e) Spatial isolation.** Manually crafted features are processed by a single-layer model [6, 6, 6, 29], where the interactions of features from different facial regions are strictly treated as linear. This strong assumption hurts the recognition performance because the complex interplay of smile features should be predominantly non-linear. Figure 1 summarizes facial feature tracking on different methods.

We present a **MeshSmileNet** framework that avoids redundant facial features and inject inductive bias of facial movements by converting each frame of the smile video into facial landmarks. We extract landmarks using Attention Mesh [15], a pre-trained model, frame by frame. Landmarks as input (instead of raw frames) make the input representation compact and enable landmark tracking across the time dimension. To yield discriminative landmark features, we explore the concept of *relativity* and *trajectory* for the smile veracity classification task. For the *relativity*, we process landmarks locally in the spatial domain with a sequence of CurveNet block [30]. It selectively groups the nearby landmarks that interact




Figure 1: Visual illustration of facial feature tracking of different methods related to smile classification. (a) [4, 5, 6, 18, 21, 26, 29] work on semi-automatic or automatic landmarks but extracts hand-engineered features. (b) [20, 80] use raw frames directly as input without any landmark location. Although they extract automatic features but need to consider redundant information. (c) Our approach contracts 3D landmark mesh to track down feature spatial and time dimensions automatically. It neither uses hand-engineered features nor redundant information of faces.

with each other. Moreover, it conceptually approximates geometrical landmark features, for example, the eye angle from those manual feature-based approaches. The spatial relativity of each landmark provides efficient features variation tracking compared to raw landmarks because of containing the nearby landmark information. For the *trajectory*, movements for the same landmark feature across time are accounted by a temporal self-attention mechanism. We place a lightweight spatial self-attention to avoid spatial isolation. The spontaneous smile usually has slow facial feature movements than the posed smile [22]. In time domains, we ground feature variation by self-attention (e.g., the speed of pulling lip corners). With this idea, our **MeshSmileNet** model extracts effective landmark features for the smile veracity classification. We perform extensive experiments and achieve state-of-the-art results on UVA-NEMO, BBC, MMI facial expression, and SPOS datasets.

Our contributions are listed below: (1) We propose the **MeshSmileNet** framework for classifying spontaneous and posed smiles based on facial landmarks. (2) We explore the concepts, relativity and trajectory, of learning end-to-end landmark features for smile classification task. (3) We perform detailed ablation studies, compare our methods with strong baselines and state-of-the-art results on four well-known smile datasets.

## 2 Related Works

**Smile Recognition.** In the facial action coding system (FACS) [4, 5], a smile triggers both the Orbicularis oculi (AU6) and the Zygomatic major (AU12) muscle by rising cheeks and lip corners. In regards to the smile duration, Dibeklioglu *et al.* [4, 5] separate it into onset (neutral to expressive), apex (stay expressive), and offset (decay to neutral) phases. Later works on smile motions are studied to distinguish spontaneous and posed smiles. Valstar *et al.* [26, 27] studies the eyebrows motion patterns and movements of head, face, and shoulders. Dibeklioglu *et al.* [4, 5] present D-Marker features that consider eyelids, cheeks, and

lip corners movements. Mandal *et al.* [18] distinguish spontaneous smiles from posed smiles by dense optical flows. Pfister *et al.* [21] and Wu *et al.* [28, 29] present a spatial-temporal texture description, CLBP-TOP, to encode eye, lip, and cheek for identifying spontaneous smile. However, all the above methods are semi-automatic that requires human intervention for feature engineering. Mandal *et al.* [21], and Yang *et al.* [51] study automatic spontaneous and posed smile recognition with convolution networks by smile videos. However, they consider redundant information from raw video frames that slowdowns the efficiency of the classification process. Our paper proposes a novel automatic framework that studies the relativity and trajectory of facial landmarks for spontaneous and posed smile recognition to address the issues of past methods.

**Relating smiles with video and point cloud classification.** Recent advancements in video and point cloud sequence classification approaches can benefit real/fake smile recognition methods. Existing transformer-based approaches for video classification focus on the study of patch tokenization, positional embedding, feedforward network activation, normalization, and efficient self-attention to scale to high volume video classification [24]. A standard approach uses a convolution network to extract spatial features of each frame and aggregate them by self-attention in the time domain. In our proposed method, we adopt such self-attention. Furthermore, to classify smile videos, one can convert a smile video to facial landmarks with 4D points and apply point cloud classification methods. Fan *et al.* [12, 13] propose PSTNet and P4Transformer with a spatial and temporal convolution to learn informative feature representation of the point cloud sequences. By a fixed spatial radius and time span, this convolution aggregates features in nearby spatial locations and time separately. In PSTNet, they progressively downsample point features in the spatial domain with the furthest point sample to perceive the entire point cloud. P4Transformer combines PSTNet with transformers to have a global interaction in the time domain for features at close spatial locations. The spatial feature aggregations of both PSTNet and P4Transformer consider all nearby landmarks. In our task, nearby landmarks are not equally important. They potentially flatten beneficial features for spontaneous and posed smile recognition. Our framework explores the relativity of nearby landmarks to help the task.

## 3 Method

### 3.1 Problem Formulation

From psychological research, we know that lips, cheeks, and eyes play important roles in a genuine smile [4, 5, 6]. For this, traditional methods find landmark locations in semi-automatic or automatic ways and track the movement of those locations using some pre-defined rules/equations/statistics (e.g., piecewise Bézier volume deformation tracker and speed, acceleration, duration, symmetries and amplitude of landmarks [4, 5, 6]). In some cases, methods [21, 28, 29] extract low-level vision features (e.g., HoG, LBP) surrounding landmark locations. We consider traditional approaches as hand-engineered feature-based methods because it requires expert intervention. Recent deep learning approaches [21, 51] make feature extraction automatic (replacing hand-engineered features) by using whole raw video frames as input. Since such approaches use CNN features without considering landmark information, they need to consider redundant facial information unrelated to smiles. Moreover, the spatial location of key landmarks might not be similar across diverse human subjects (e.g., young vs. adults), which may restrict the network from extracting necessary

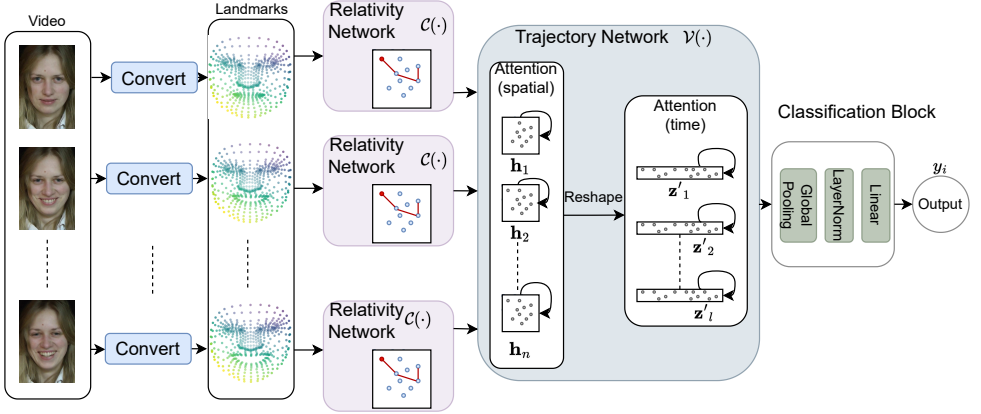


Figure 2: Our proposed architecture, **MeshSmileNet**. See supplementary material for more details.

information. Our goal is to propose a deep architecture that can track landmark locations in an end-to-end manner without using redundant information and expert hand-engineering. For this, we leverage 3D landmarks obtained from Attention Mesh [15].


Given a video  $\mathcal{X} = [\mathbf{x}_n \mid n \in 1 \dots N]$ , where  $\mathbf{x}_t$  and  $N$  are  $t$ th frame and video length, respectively, we aim to identify if  $\mathcal{X}$  is a spontaneous ( $y = 0$ ) or posed ( $y = 1$ ) smile. Here,  $y$  is the true label. From each frame,  $\mathbf{x}_n \in \mathcal{X}$ , we extract 3D facial landmarks,  $\mathbf{a}_n \in \mathbb{R}^{3 \times L}$ , where  $L$  is the total number of landmarks. Suppose,  $\mathcal{A} = [\mathbf{a}_n \mid n \in 1 \dots N]$  denotes the landmark representation of a smile video. Our goal is to train a model  $\mathcal{F}(\cdot)$  that maps  $\mathcal{A}$  to  $y$ .

### 3.2 MeshSmileNet

**Model Overview.** First, we extract  $L$  number of facial landmarks ( $\mathbf{a}_n$ ) from each frame ( $\mathbf{x}_n$ ) of a video using Attention Mesh [15]. As shown in Figure 2, all extracted landmarks are fed into our proposed network,  $\mathcal{F}(\cdot)$ . It composes of three main components. (1) **Relativity network:** We explore spatial geometry relations of landmarks at each frame ( $\mathbf{a}_n \in \mathcal{A}$ ) based on CurveNet blocks [30]. It links landmarks with similar features into a curve before aggregation. (2) **Trajectory network:** As  $\mathbf{a}_n \in \mathcal{A}$  is ordered, we track the movement of each landmark  $\mathbf{a}_{n,l}$ ,  $l$ th landmark at  $n$ th frame. This operation is based on the self-attention mechanism, aiming to measure the variation of the constructed geometry feature across time. To have interactions between the learned geometric features (avoid spatial isolation), we employ a lightweight self-attention for them on spatial domain. It also helps trajectory measurement of our purpose. (3) **Classification network:** We classify the video into label  $y$ , spontaneous smiles or posed smiles, from the relativity and trajectory based landmark features. Now, we elaborate on each network’s components.

**Relativity network.** Our relativity network  $\mathcal{C}(\cdot)$  models the local geometry of nearby landmarks. Past studies show geometry-based features such as lip pulling, which is formed by a curve from at least three landmarks, are beneficial in the spontaneous and posed smile recognition task. We borrow CurveNet blocks [30] into our framework for this purpose. For landmarks of the same frame, a CurveNet block forms  $k$  curves. It determines their starting points by top- $k$  selection method [14] that ranks projected scores of each landmark. Afterward, a curve is sequentially formulated by connecting to neighborhood landmarks with

Figure 3: Overall distinguishing face landmarks (marked as blue) for real fake smile classification while (a) only spatial, (b) only time dimensional, and (c) both spatial and time dimensional self-attention is used on UVA-NEMO dataset. Our method finds landmarks related to lips, cheeks, eyes, and face borders essential for the classification task. The bottom two rows show distinguishing face landmarks of five individual samples of real and fake smile videos, respectively. It tells that important landmarks are found in similar locations regardless of real or fake smiles.



Gumbel-Softmax liked selection methods before pooling into a unified size of landmark features. We denote these processes by:

$$\mathbf{h}_n = \mathcal{C}(\mathbf{a}_n) \quad \forall n \in [1 \cdots N],$$

where,  $h_n \in \mathbb{R}^{d \times L}$  is  $d$ -dimensional output representation of  $L$  landmarks.

**Trajectory network.** We measure the variation of the learned geometry feature by the trajectory network  $\mathcal{V}(\cdot)$ . It builds up representations that describe the movements of facial features such as the pull of the zygomatic major and orbicularis oculi muscle. Naively, it could be achieved by introducing a self-attention to all landmarks. However, this is computationally expensive. Motivated by the axial attention [16], we decompose it into self-attention in spatial and time domains. Our  $\mathcal{V}(\cdot)$  has

$$\begin{aligned} \mathbf{o}_n &= \text{Attention}(\mathbf{h}_n, \mathbf{h}_n, \mathbf{h}_n) & \forall n \in [1 \cdots N], \\ \mathbf{z}'_l &= [\mathbf{o}_{1,l} \cdots \mathbf{o}_{N,l}] & \forall l \in [1 \cdots L], \\ \mathbf{z}_l &= \text{Attention}(\mathbf{z}'_l, \mathbf{z}'_l, \mathbf{z}'_l) & \forall l \in [1 \cdots L], \end{aligned}$$

where,  $\mathbf{o}_n \in \mathbb{R}^{d \times L}$  is the output for attention on each frame,  $\mathbf{z}'_l \in \mathbb{R}^{d \times N}$  concatenates feature of  $l$ th landmark across time, and  $\mathbf{z}_l \in \mathbb{R}^{d \times N}$  is the attention output for,  $\mathbf{z}_l$ , the same landmark across time.  $\mathbf{o}_n$  facilitates better spatial local feature representation with interacting spatial landmark features, and  $\mathbf{z}_l$  tracks the movement of local facial features. Semantically, following on the above example,  $\forall l \in [1 \cdots L]$   $\mathbf{z}_l$  can model ‘the pull of the zygomatic major muscle’ or ‘the pull of the orbicularis oculi muscle’ separately. Here,  $\forall n \in [1 \cdots N]$   $\mathbf{o}_n$  brings interactions by aggregating them.

In Figure 3, we analyze the complex interplay of relativity and trajectory networks by visualizing the importance of different landmark locations by calculating their gradient with respect to the target label. Note that we average the background face, landmark position, and gradient of the entire UVA-NEMO dataset. The trajectory network,  $\mathcal{V}(\cdot)$ , calculates self-attention across spatial and time dimensions of videos. While using only spatial self-attention in Figure 3(a), we notice that lips regions (spanning from nose bottom to chin) play decisive facial features for the task. Then, only self-attention across time dimensions in Figure 3(b) also finds the landmarks around the left and right sides of the face border. By using self-attention across both spatial and time dimensions, our final recommendation in



Figure 3(c) combines individual contributions (of spatial and time) and additionally focuses on the nose and eyes regions. This finding is somewhat similar to Dibeklioglu *et al.* [8] which establishes that lips, cheek, and eyes contribute more to real and fake smile classification. Our method automatically arrives at the similar conclusion without any hand-engineering. Being data dependent nature of deep learning, our method also concentrates on the nose and face border-related landmarks.

**Classification Block.** We apply a global pooling on  $\forall l \in [1 \cdots L] \mathbf{z}_l$ , and then it is passed to distinguish between spontaneous and posed smiles. Our Classification Block has a global pooling layer, a Layer Normalization, a Linear layer, and a Sigmoid activation function.

**Objective:** Our goal is to classify spontaneous and posed smiles. We penalize deviation of network predictions  $\mathcal{F}(\mathcal{A})$  from the ground truth  $y_i$  using Binary Cross Entropy loss:

$$\mathcal{L}_{BCE} = y_i \log \mathcal{F}(\mathcal{A}) + (1 - y_i) \log(1 - \mathcal{F}(\mathcal{A}))$$

**Inference.** Given a test smile video, after extracting landmarks locations representing that video,  $\mathcal{A}^*$ , we perform a forward pass for prediction,  $\hat{y}$ , i.e.,  $\hat{y} = \mathcal{F}(\mathcal{A}^*)$ .

## 4 Experiments

**Datasets.** We experiment with four smile datasets: **(a)** UvA-NEMO dataset[8]: This dataset contains 597 spontaneous and 643 posed smile videos of 400 participants. It is recorded in resolution  $1920 \times 1080$  at 50 frames per second. Among the four datasets, this one has a higher resolution and frames per second (FPS) and the highest number of videos. **(b)** BBC dataset: This dataset has 20 videos, divided into 10 posed and 10 spontaneous. The videos are recorded in resolution  $314 \times 286$  with 25 FPS. **(c)** MMI facial expression dataset[29]: It contains spontaneous and posed facial expressions. Similar to [8, 30], we use 138 spontaneous and 49 posed smile videos. The spontaneous smile part is recorded in resolution  $640 \times 480$  and at 29 FPS, and the posed smile videos are recorded in  $720 \times 576$  pixels at 25 FPS. **(d)** SPOS dataset [21]: There is two types of image sequences, gray and near-infrared sequences and all the images are in  $640 \times 480$  pixels with 25 FPS. There are in total 80 smiles, 66 spontaneous, and 14 deliberate smiles. We use the gray image sequences in our research. Detailed statistics of smile datasets are presented in the supplementary material.

**Train/test split & evaluation.** We follow the train/test split and 10-cross-validation settings from [8] for the UvA-NEMO dataset. For the BBC, MMI, and SPOS datasets, we follow [8, 29] and use 10-fold, 9-fold, and 7-fold cross-validation, respectively. All the datasets maintain no subject overlap in training and testing splits. To evaluate the model, we use prediction accuracy. We run ten trials to report the average result.

**Implementation details<sup>1</sup>.** We extract 478 landmarks per frame with Attention Mesh [15]. Our model is trained with these extracted landmarks. We use a batch size of 16 and train the model for 300 epochs. We use AdamW optimizer with a learning rate of  $5e - 4$ . We sample 10 FPS for the UvA-NEMO, MMI, and BBC dataset and 1 FPS for the SPOS dataset. In training, we sample 16 continuous frames per video. During testing, we average the prediction score for 5 samples per video. In the Relativity Network, we use a CurveNet block having four CIC layers [30]. We have a convolution 1d layer to linearly mix the 478 landmark features per frame to 32 features and then pass it to the Trajectory network  $\mathcal{V}(\cdot)$ . It has nine regular transformer blocks. Six of them perform spatial self-attention, and three are for self-attention in the time domain. We use the *PyTorch* library to perform our experiments.

<sup>1</sup>Codes are available at: <https://github.com/Tushar-Faroque/MeshSmileNet>

Method	UVA-NEMO	MMI	SPOS	BBC
Cohn'04 [8]	77.3	81.0	73.0	75.0
Dibeklioglu'10 [9]	71.1	74.0	68.0	85.0
Pfister'11 [10]	73.1	81.0	67.5	70.0
Wu'14 [28]	91.4	86.1	79.5	90.0
Dibeklioglu'15 [9]	89.8	88.1	77.5	90.0
Wu'17 [28]	<b>93.9</b>	92.2	81.2	90.0
Mandal'17 [25]	80.4	-	-	-
Mandal'16 [20]	78.1	-	-	-
RealSmileNet'20 [26]	82.1	92.0	86.2	90.0
PSTNet [23]	72.9	94.3	87.1	95.0
P4Transformer [20]	74.9	91.3	82.9	85.0
Vanilla ViT [8]	78.4	99.0	93.5	95.0
Ours	85.0	<b>99.0</b>	<b>94.4</b>	<b>95.0</b>

Table 1: Comparison of state-of-the-art methods.

Method	UVA-NEMO	MMI	SPOS	BBC
Dibeklioglu'12 [9]	72.1	78.2	53.2	55.0
Dibeklioglu'15 [9]	77.0	73.9	46.9	55.0
Ours	<b>85.0</b>	<b>99.0</b>	<b>94.4</b>	<b>95.0</b>

Table 2: Using Attention Mesh landmark extractors on other related methods [9, 9].

Individual Group	Real	
UVA-NEMO	SmileNet [26]	Ours
Young	79.6	<b>80.4</b>
Adult	79.4	<b>82.4</b>
Male	77.8	<b>81.4</b>
Female	80.0	<b>80.2</b>

Table 3: Results on diverse subject groups.

## 4.1 Main Results

**Comparison methods.** We compare our work with three categories of methods. (a) Semi-automatic: These approaches [9, 9, 9, [18], [21], [26], [29]] require manual annotation of facial landmarks for the first frame in a smile video. Also, the extract smile features are hand-engineered by experts. (b) Fully-automatic: Unlike the semi-automatic case, these approaches [20, [5]] require no manual annotation for the first frame. They use CNN layers on the raw video frames which allows many redundant smile features. (c) Baseline: We propose three baselines that could be a reasonable but incomplete version of our method. After extracting landmarks from each frame using Attention Mesh [25], we represent a video by 4D point clouds. Then, we apply point cloud sequence classification methods like PSTNet [23] and P4Transformer [20] to classify smile videos. Also, we feed raw video frames to the vanilla vision transformer (ViT) [8] model for video classification. PSTNet, P4Transformer, and ViT represent three baselines of our method.

**Performance benchmark.** In Table 1, we present a benchmark comparison of our methods with state-of-the-art approaches based on four datasets. We separate ‘Semi-automatic’, ‘Fully-Automatic’, and ‘Baseline’ methods by horizontal lines. Our observations are as follows: (1) Compared to semi-automatic methods, we achieve state-of-the-art results in the MMI, SPOS, and BBC datasets. In the UvA-NEMO dataset, the best performance was reported by Wu *et al.* [28], but it requires hand-crafted features and heavy preprocessing steps. Our model is the most competitive automatic method in classifying spontaneous smiles, which classifies spontaneous smiles from facial landmarks with a single forward pass. (2) When comparing past fully automatic methods, our method significantly outperforms them. Mandal *et al.* [20] and Yang *et al.* [5] classify spontaneous smile with a CNN that inputs raw frames. However, in this task, most facial features are irrelevant, such as identity information, but they redundantly cover key facial landmarks that are beneficial for the spontaneous smile classification task. As a result, their approaches are difficult to understand the spontaneous smile-related features. Unlike them, our **MeshSmileNet** resolves the difficulty by learning relativity and trajectory of landmarks extracted from raw frames. (3) Our model outperforms baseline methods by a huge margin. Again, the vanilla ViT applies self-attention to raw frames in the spatial domain and time domain. Though global interactions are achieved, variations of the same spatial features need to be modeled from many patch features. It creates difficulty for the spontaneous smile classification. PSTNet [23] and P4Transformer [20] have landmark inputs. They use spatial and temporal convolution and P4Transformer [20] to aggregate features from nearby landmarks. In our task, landmarks are not equally impor-



Input	CurveNet	Attention	UVA-NEMO	MMI	SPOS	BBC
Frames	No	Both*	78.4	99.0	93.5	95.0
Landmarks	No	Both*	82.2	96.7	92.2	95.0
Landmarks	Yes	Spatial	72.3	98.5	90.1	95.0
Landmarks	Yes	Time	82.4	98.5	90.9	95.0
Landmarks	Yes	Both*	<b>85.0</b>	<b>99.0</b>	<b>94.4</b>	<b>95.0</b>

Table 4: Ablation on model architectures. \* means self-attention is used in both spatial and time direction.

Frame Rate	UVA-NEMO	MMI	SPOS	BBC
1	67.0	98.6	<b>94.4</b>	<b>95.0</b>
3	74.5	98.6	90.3	<b>95.0</b>
5	<b>85.0</b>	<b>99.0</b>	92.4	<b>95.0</b>
10	82.2	97.6	92.4	90.0

Table 5: Impact of frame rates across difference datasets.

tant to each other. Our **MeshSmileNet** selects nearby landmarks interacting with each other before aggregation and improves spontaneous smile classification performance.

**Significant test.** We follow [2] to determine if our methods significantly outperform the automatic baselines. With the 10-fold partitions from the UVA-NEMO (each testing fold contains around 100 videos, and there are 1240 testing videos in total), we use the paired t-test to compare our methods against the baselines in Table 6. Since p-values are  $< 10^{-3}$ , we conclude that our method exhibits statistically significant results than the baselines.

**Role of landmark extractor.** We work on FaceMash, which extracts more landmark locations (478 vs. 11) than the landmark tracker from past methods. In Table 2, we compare the performance of using FaceMash landmarks on Dibeklioglu’12 [5] and Dibeklioglu’15 [6]. Their feature extraction process is extremely sensitive to landmark locations because they usually track facial landmarks across the video with a pre-defined set of rules and statistics. Moreover, they require manual annotation of landmarks in the first frame. Here, when scaling their approach using the automatic landmark extractor, FaceMash, they perform lower than our proposed model. It tells that more landmarks do not help until features are learned end-to-end using deep learning methods (instead of hand-engineered features).

**Subject group result.** UVA-NEMO dataset provides smile data of different subject groups, e.g., young, adult, male and female. We apply our method in each individual groups and compare the result with *RealSmileNet* [11] in Table 3. Our method outperforms [11] which confirms the robustness of the proposed method. This becomes possible because of excluding redundancies of raw frames which were used in [11].

## 4.2 Ablation studies

**Ablation on architecture.** In Table 4, we perform ablation studies based on architecture components. Firstly, we feed raw frames directly to the ViT model for video classification where no CurveNet is used. It achieves poor performance because of redundant features. Secondly, we directly apply landmark inputs to the vanilla ViT model with self-attention in both spatial and time domains. It addresses redundant feature issues and shows considerable performance improvement. Thirdly, we include the relativity network (implemented as CurveNet), which deals with spatial geometry relations of landmarks. Then, we separately try spatial and time self-attention in the trajectory network. For both cases, we notice performance improvement over the second case because of measuring the variation of the constructed geometry feature through the trajectory network. Finally, we apply all components as our final recommended method, outperforming all previously ablated architectures. Using landmarks as input instead of raw frames and the interplay between relativity and trajectory networks help our method achieve outstanding results.

**Ablation on frame rate.** We study the impact of frame rates (Table 5) for our model on four datasets. Using 5 FPS per video, our model performs the best in UVA-NEMO, MMI, and BBC. However, 1 FPS serves the best for the SPOS. Note SPOS dataset provides only

	Ours vs. ViT	Ours vs. PSTNet	Ours vs. P4Transformer
t-statistic	8.87	22.97	10.38
p-value	$9.6 \times 10^{-6}$	$7.5 \times 10^{-8}$	$2.6 \times 10^{-6}$

Table 6: Significance tests with baselines.

Method	PSTNet	P4Transformer	Ours
DLIB (□)	69.7	62.9	80.3
Attention Mesh (■)	<b>72.9</b>	<b>74.9</b>	<b>85.0</b>

Table 7: DLIB vs. Attention Mesh landmark extractor.

the onset phase of the smile videos. In contrast, UVA-NEMO, MMI, and BBC provide all three phases (onset, apex, and offset) in the smile videos. The onset phase usually is in slow motion. Thus, low FPS finds the best performance for SPOS. Nevertheless, FPS to be used can be considered as a hyper-parameter for validation.

**Ablation on landmark extractor.** We compare Attention Mesh landmark extractor (■) with widely used DLIB extractor (□) (with 68 landmark locations). In Table 7, we show the results on automatic baselines and with our method. We exclude ViT from the comparison since ViT does not use facial landmarks. Because of providing more landmark locations in Attention Mesh (478 vs. 68), it beats the DLIB extractor. However, our method performs well against the baselines for both extractors’ cases, showing our generalization ability.

## 5 Conclusion

In this paper, we propose a framework for real/fake smile recognition, which leverages facial landmarks from video data to learn spatial and temporal. In the past, models relied on raw smile frame input, manual feature engineering, phase segmentation, etc., which hampers the efficiency of the overall smile classification. In contrast, we focus on the relatively and trajectory of facial landmarks using CurveNet and the self-attention mechanism network, respectively. This way, we learn spatial and temporal features from them and classify spontaneous and posed smiles. We executed experiments on UvA-NEMO, MMI, BBC, and SPOS smile datasets and achieve state-of-the-art results.

**Acknowledgement.** This work was supported by North South University (NSU) Conference Travel and Research Grants (CTRG) 2021–2022 (Grant ID: CTRG-21-SEPS-10). Authors also would like to thank Machine Learning & Artificial Intelligence Future Science Platforms (MLAI-FSP), CSIRO for funding this work.

## References

- [1] Zara Ambadar, Jeffrey F Cohn, and Lawrence Ian Reed. All smiles are not created equal: Morphology and timing of smiles perceived as amused, polite, and embarrassed/nervous. *Journal of nonverbal behavior*, 33(1):17–34, 2009.
- [2] Marian Stewart Bartlett, Gwen Littlewort, Mark G Frank, Claudia Lainscsek, Ian R Fasel, Javier R Movellan, et al. Automatic recognition of facial actions in spontaneous expressions. *J. Multim.*, 1(6):22–35, 2006.
- [3] Jeffrey F Cohn and Karen L Schmidt. The timing of facial motion in posed and spontaneous smiles. *International Journal of Wavelets, Multiresolution and Information Processing*, 2(02): 121–132, 2004.
- [4] Hamdi Dibeklioglu, Roberto Valenti, Albert Ali Salah, and Theo Gevers. Eyes do not lie: Spontaneous versus posed smiles. In *Proceedings of the 18th ACM international conference on Multimedia*, pages 703–706, 2010.

- [5] Hamdi Dibeklioğlu, Albert Ali Salah, and Theo Gevers. Are you really smiling at me? spontaneous versus posed enjoyment smiles. In *European Conference on Computer Vision*, pages 525–538. Springer, 2012.
- [6] Hamdi Dibeklioğlu, Albert Ali Salah, and Theo Gevers. Recognition of genuine smiles. *IEEE Transactions on Multimedia*, 17(3):279–294, 2015.
- [7] Thomas Dietterich. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural computation*, 10:1895–1923, 10 1998.
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. 10 2020.
- [9] Paul Ekman and Wallace V Friesen. Facial action coding system. *Environmental Psychology & Nonverbal Behavior*, 1978.
- [10] Paul Ekman, Richard J Davidson, and Wallace V Friesen. The duchenne smile: Emotional expression and brain physiology: Ii. *Journal of personality and social psychology*, 58(2):342, 1990.
- [11] Hehe Fan, Yi Yang, and Mohan Kankanhalli. Point 4d transformer networks for spatio-temporal modeling in point cloud videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14204–14213, 2021.
- [12] Hehe Fan, Xin Yu, Yi Yang, and Mohan Kankanhalli. Deep hierarchical representation of point cloud videos via spatio-temporal decomposition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- [13] Hehe Fan, Xin Yu, Yuhang Ding, Yi Yang, and Mohan Kankanhalli. Pstnet: Point spatio-temporal convolution on point cloud sequences. *arXiv preprint arXiv:2205.13713*, 2022.
- [14] Hongyang Gao and Shuiwang Ji. Graph u-nets. In *international conference on machine learning*, pages 2083–2092. PMLR, 2019.
- [15] Ivan Grishchenko, Artsiom Ablavatski, Yury Kartynnik, Karthik Raveendran, and Matthias Grundmann. Attention mesh: High-fidelity face mesh prediction in real-time. *arXiv preprint arXiv:2006.10962*, 2020.
- [16] Jonathan Ho, Nal Kalchbrenner, Dirk Weißenborn, and Tim Salimans. Axial attention in multi-dimensional transformers. 12 2019.
- [17] Davis E King. Dlib-ml: A machine learning toolkit. *The Journal of Machine Learning Research*, 10:1755–1758, 2009.
- [18] Bappaditya Mandal and Nizar Ouarti. Spontaneous versus posed smiles—can we tell the difference? In *Proceedings of International Conference on Computer Vision and Image Processing*, pages 261–271. Springer, 2017.
- [19] Bappaditya Mandal, David Lee, and Nizar Ouarti. Distinguishing posed and spontaneous smiles by facial dynamics. In Chu-Song Chen, Jiwen Lu, and Kai-Kuang Ma, editors, *Computer Vision - ACCV 2016 Workshops - ACCV 2016 International Workshops, Taipei, Taiwan, November 20-24, 2016, Revised Selected Papers, Part I*, volume 10116 of *Lecture Notes in Computer Science*, pages 552–566. Springer, 2016. doi: 10.1007/978-3-319-54407-6\_37. URL [https://doi.org/10.1007/978-3-319-54407-6\\_37](https://doi.org/10.1007/978-3-319-54407-6_37).

- [20] Bappaditya Mandal, David Lee, and Nizar Ouarti. Distinguishing posed and spontaneous smiles by facial dynamics. In *Asian Conference on Computer Vision*, pages 552–566. Springer, 2016.
- [21] Tomas Pfister, Xiaobai Li, Guoying Zhao, and Matti Pietikäinen. Differentiating spontaneous from posed facial expressions within a generic facial expression recognition framework. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 868–875. IEEE, 2011.
- [22] Karen L Schmidt, Zara Ambadar, Jeffrey F Cohn, and L Ian Reed. Movement differences between deliberate and spontaneous facial expressions: Zygomaticus major action in smiling. *Journal of nonverbal behavior*, 30(1):37–52, 2006.
- [23] Karen L Schmidt, Sharika Bhattacharya, and Rachel Denlinger. Comparison of deliberate and spontaneous facial movement in smiles and eyebrow raises. *Journal of nonverbal behavior*, 33(1):35–45, 2009.
- [24] Javier Selva, Anders Johansen, Sergio Escalera, Kamal Nasrollahi, Thomas Moeslund, and Albert Clapés. Video transformers: A survey. 01 2022.
- [25] Michel Valstar, Maja Pantic, et al. Induced disgust, happiness and surprise: an addition to the mmi facial expression database. In *Proc. 3rd Intern. Workshop on EMOTION (satellite of LREC): Corpora for Research on Emotion and Affect*, page 65. Paris, France., 2010.
- [26] Michel F Valstar, Maja Pantic, Zara Ambadar, and Jeffrey F Cohn. Spontaneous vs. posed facial behavior: automatic analysis of brow actions. In *Proceedings of the 8th international conference on Multimodal interfaces*, pages 162–170, 2006.
- [27] Michel F Valstar, Hatice Gunes, and Maja Pantic. How to distinguish posed from spontaneous smiles using geometric features. In *Proceedings of the 9th international conference on Multimodal interfaces*, pages 38–45, 2007.
- [28] Ping-ping Wu, Hong Liu, Xue-wu Zhang, and Yuan Gao. Spontaneous versus posed smile recognition via region-specific texture descriptor and geometric facial dynamics. *Frontiers of Information Technology & Electronic Engineering*, 18(7):955–967, 2017.
- [29] Pingping Wu, Hong Liu, and Xuewu Zhang. Spontaneous versus posed smile recognition using discriminative local spatial-temporal descriptors. In *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1240–1244. IEEE, 2014.
- [30] Tiange Xiang, Chaoyi Zhang, Yang Song, Jianhui Yu, and Weidong Cai. Walk in the cloud: Learning curves for point clouds shape analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 915–924, 2021.
- [31] Yan Yang, Md Zakir Hossain, Tom Gedeon, and Shafin Rahman. Realsmilenet: a deep end-to-end network for spontaneous and posed smile recognition. In *Proceedings of the Asian Conference on Computer Vision*, 2020.