

Exploring Spatial-Temporal Dynamics in Event-based Facial Micro-Expression Analysis

Nicolas Mastropasqua^{1,3*} Ignacio Bugueno-Cordova^{2,4*} Rodrigo Verschae² Daniel Acevedo^{1,3}
Pablo Negri^{1,3} Maria E. Buemi^{1,3}

¹Universidad de Buenos Aires, Facultad de Ciencias Exactas y Naturales, Argentina

²Institute of Engineering Sciences, University of O'Higgins, Chile

³CONICET-UBA, Instituto de Ciencias de la Computacion (ICC), Argentina

⁴L3S Research Center, Leibniz Universität Hannover, Germany

{nmastropasqua, dacevedo, pnegri, mebuemi}@dc.uba.ar

i.bugueno@ieee.org, rodrigo@verschae.org

Abstract

Micro-expression analysis has applications in domains such as Human-Robot Interaction and Driver Monitoring Systems. Accurately capturing subtle and fast facial movements remains difficult when relying solely on RGB cameras, due to limitations in temporal resolution and sensitivity to motion blur. Event cameras offer an alternative, with microsecond-level precision, high dynamic range, and low latency. However, public datasets featuring event-based recordings of Action Units are still scarce. In this work, we introduce a novel, preliminary multi-resolution and multi-modal micro-expression dataset recorded with synchronized RGB and event cameras under variable lighting conditions. Two baseline tasks are evaluated to explore the spatial-temporal dynamics of micro-expressions: Action Unit classification using Spiking Neural Networks (51.23% accuracy with events vs. 23.12% with RGB), and frame reconstruction using Conditional Variational Autoencoders, achieving $SSIM = 0.8513$ and $PSNR = 26.89$ dB with high-resolution event input. These promising results show that event-based data can be used for micro-expression recognition and frame reconstruction.

1. Introduction

Sensing and analyzing micro-expressions provides information about subtle human emotions. Facial Expression Recognition (FER) within Affective Computing has applications in social robotics [30] and intelligent vehicles, where Driver Monitoring Systems (DMS) use fa-

*These authors contributed equally to this work.

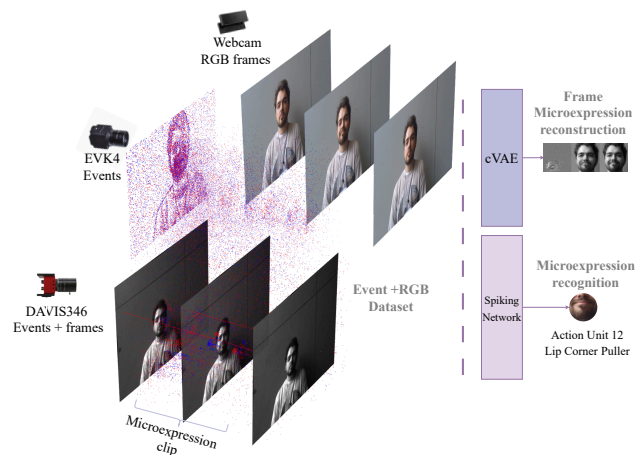


Figure 1. Our proposed preliminary dataset to study Action Unit recognition and Facial micro-expression reconstruction consists of multi-modal (RGB and Events) and multi-resolution recordings of micro-expressions using a DAVIS346 event camera, a Prophesee EVK4 event camera, and an ASUS ROG Eye S RGB camera.

cial expressions—both macro [8, 40] and micro [38]—to detect drowsiness and emotional states, supporting road safety [44]. Micro-expression analysis also contributes to Autism Spectrum Disorder detection [19].

Deep learning models trained on large FER datasets achieve high performance for posed expressions representing Ekman’s six basic emotions [12] under controlled conditions. However, performance decreases in ‘in-the-wild’ settings with varied lighting, occlusions, and poses [18]. Categorizing emotions into a few archetypal expressions

oversimplifies their diversity [29, 33], introducing label ambiguity and inconsistency across datasets [45].

An alternative and more descriptive codification is given by the Facial Action Coding System (FACS) [13], which aims to identify precise facial micro-movements called Action Units (AUs). AU-labeled datasets remain limited due to the time-intensive, expertise-dependent annotation process. Moreover, RGB cameras impose temporal constraints, as micro-expressions last between 170–500 ms, with a lower bound near the frame rate of standard cameras [43]. As a result, current sensors and methods remain insufficiently robust for AU recognition [20]. Alternative sensing modalities, such as thermal [1] and event cameras [5], address these limitations.

Event cameras provide microsecond-level temporal resolution, low latency, high dynamic range (≥ 120 dB), and reduced power consumption. They also improve privacy, as re-identification from event data is more challenging than from RGB images [2]. Although there has been an increased interest in using event cameras for facial analysis [4], publicly available event-based micro-expression datasets remain scarce. Recent efforts include VETEX [1], covering only a subset of AUs, and FACEMORPHIC [5], which is not yet fully accessible. Simulated event streams from RGB videos, using tools like v2e [16] and ESIM [31], offer an alternative but do not replicate the temporal dynamics and noise characteristics of real neuromorphic data.

Processing data from event cameras requires models that operate natively on sparse, temporally precise signals. Spiking Neural Networks (SNNs)[23], inspired by biological neurons, are naturally suited for this representation. While most deep learning models use backpropagation with labeled datasets, brain-inspired unsupervised local rules, such as Spike Timing Dependent Plasticity (STDP) allow SNNs to adapt by modulating synaptic strengths based on the Hebbian principle. However, STDP alone is often insufficient for training deep SNNs, motivating the use of surrogate gradient descent to address spike non-differentiability with smoothing functions [26].

In this work, we propose to explore the spatial-temporal dynamics of micro-expressions using event-based facial data. To this end, we introduce a preliminary multi-modal, multi-resolution dataset of performed Action Units recorded with synchronized RGB and event cameras under varying lighting conditions. Using this dataset, we evaluate two baseline tasks: (1) Action Unit recognition with Spiking Neural Networks trained via surrogate gradient descent, and (2) frame reconstruction from event data using Conditional Variational Autoencoders (cVAE), comparing low and high-resolution event inputs.

The main contributions of this work are:

- A preliminary multi-modal, multi-resolution dataset of RGB and event data under mixed lighting conditions and

posed Action Units. A larger version of this dataset is under construction and is expected to be publicly available upon request soon.

- An exploration of Spiking Neural Networks trained with surrogate gradient descent to learn spatial-temporal dynamics for Action Unit recognition on this dataset.
- An exploration of intensity frame micro-expression reconstruction models conditioned on event-based inputs, highlighting differences between low and high resolution data from DAVIS346 and EVK4, respectively, in terms of reconstruction quality.

The rest of this paper is organized as follows: Section 2 reviews related work; Section 3 introduces the novel and preliminary multi-modal dataset for micro-expression analysis; Section 4 presents the proposed baselines for exploring the spatio-temporal dynamics in facial micro-expression recognition and intensity reconstruction; Section 5 describes the training details; Section 6 reports the preliminary results; and Section 7 concludes and projects this work.

2. Related work

2.1. Event-based micro-expression datasets

While micro-expression datasets for the RGB modality are relatively small and scarce, their event-based counterpart is no exception. Recently, the community has made valuable efforts to address this problem in multiple ways: Berlincioni *et al.* [6] introduced NEFER, a dataset for emotion recognition which consists of paired RGB and Event videos annotated with the six basic emotions, as well as landmarks and bounding boxes. However, it is not very well suited for micro-expression analysis since it does not provide AU annotations. The first event dataset containing a large subset of Action Units to date was proposed by Becattini *et al.* [5]. FACEMORPHIC contains 64 subjects performing 24 AUs recorded with a 640×480 RGB camera and a 1280×720 event camera, which are temporally synchronized. Similarly, Adra *et al.* [1] proposed the VETEX dataset containing event, RGB and thermal recordings of 30 participants performing a subset of seven AUs under natural and controlled artificial illumination. However, to date, both of these datasets have yet to be made publicly available.

An alternative approach to collecting real event data consists of simulating an event stream from RGB videos. For example, Verschae *et al.* [36] used the v2e simulator [16] to obtain an event-based version of the CK+ [22] and MMI [28] datasets, both well-known benchmarks for posed FER and AU recognition. They showed that the Asynchronous Sparse Neural Network (Asynet) [25] can achieve promising results on the proposed datasets. Although working with simulated events facilitates training and testing new algorithms, a significant domain shift still emerges between

simulated and real event data. For example, compression artifacts in videos might lead to the generation of noisy events during the simulation [7]. Furthermore, simulated events share a series of discrete timestamps, unlike real data. This can accentuate the domain shift for models sensitive to timestamp distribution, such as Spiking Neural Networks [46].

2.2. Spiking Neural Networks

The neuromorphic-driven nature of Spiking Neural Networks (SNN) makes them suitable for processing event camera streams. Spiking neurons communicate with each other by sending discrete spikes when their accumulated membrane potential surpasses a threshold value. Moreover, existing neuromorphic hardware, such as TrueNorth [24] and Loihi [10], offer exciting opportunities to leverage these features.

However, SNNs still fall short of ANNs in terms of their training capabilities, generally struggling to achieve comparable performance. Training methods for SNNs can be grouped into three categories: supervised learning with gradient descent, unsupervised learning with local learning such as Spike-Timing-Dependent Plasticity (STDP), and reinforcement learning using reward-modulated plasticity [42]. Existing works have studied the application of SNNs on different tasks and datasets, from small complexity to more challenging ones. For example, Vasudevan *et al.* [35] employed SNNs trained with different learning algorithms based on backpropagation for gesture recognition, obtaining a reasonable baseline in their proposed SL-ANIMALS-DVS dataset and good results in the DVS-GESTURES dataset. Dong *et al.* [11] proposed the use of temporal batch STDP and designed an adaptive synaptic filter, together with an adaptive lateral inhibitory connection, achieving current state-of-the-art performance for unsupervised STDP-trained SNNs on MNIST and FashionMNIST. They also showed that their SNNs outperformed ANNs for smaller training sizes of MNIST. However, these techniques still struggle to train SNNs capable of performing as well as ANNs. Interestingly, state-of-the-art results are commonly achieved by converting trained ANNs to SNNs [42].

Currently, to the best of our knowledge, there are no works using SNNs for AU recognition on real event data. To evaluate its viability, we propose a standard SNN architecture with learnable membrane time constant based on the formulations given by Fang *et al.* [14].

2.3. Event-based facial reconstruction for micro-expression analysis

Reconstruction of facial appearance from event data supports tasks such as AU recognition and micro-expression analysis, where subtle motion cues need to be preserved. Several methods have addressed scene reconstruction from

event streams. Scheerlinck *et al.* [32] proposed an efficient network with reduced parameters and inference time. Pan *et al.* [27] introduced the Event-based Double Integral model to recover sharp high-frame-rate videos from blurred images and events.

For facial analysis, Barua *et al.* [3] developed a patch-based model capable of reconstructing faces and demonstrated direct face detection from events. Zou *et al.* [47] proposed a convolutional recurrent network to reconstruct temporally consistent HDR videos, validated on a dedicated dataset. Recently, Verschae *et al.* [37] proposed evTRANSFER, a framework that employs an adversarially trained encoder for facial reconstruction and reuses it for expression recognition. While the method focuses on macro-expressions, it demonstrates that reconstruction from event data can provide useful representations for facial analysis. To date, however, micro-expression reconstruction from real event streams remains unexplored.

In this work, we explore Variational Autoencoders conditioned on real event data as a baseline approach to generate reconstructed frames of micro-expressions, enabling a qualitative assessment of the temporal dynamics captured across different sensor resolutions.

3. Dataset collection

To improve the public availability of event-based micro-expression data and help the community deepen the study in this area, we collected a reduced set of videos of subjects performing a subset of Action Units according to FACS. This, in turn, allowed us to assess the viability of SNNs for AU recognition in this scenario. All videos were recorded with an ASUS RGB webcam and DAVIS346 and Prophesee EVK4 event cameras. People aged 18 to 65 years were invited to participate in the experiment.

3.1. Annotation schema

We specifically selected 21 AUs focusing on potential Driving Monitoring Systems applications. According to [39], some AUs are indicative of certain levels of drowsiness before falling asleep: AU 45 (blink), AU 2 (outer brow raiser), AU 15 (frown), AU 17 (chin raise), and 9 (nose wrinkle). Some AUs represent macro-expressions like turning the head to one side or moving the eyes. We also considered those actions, as they could be related to potential distractions while driving. Based on EMFACS [15], it is possible to decode a group of AUs activations into one of the six basic emotions. The rest of the AUs were selected to complete the codification of Happiness, Anger, and Surprise.

Unlike standard FACS-coded RGB datasets based on certified human coders, our labeling process was carried out through a guided elicitation process. For each AU, subjects were shown a video demonstration of an individual

AU Number	Description
AU 2	Outer Brow Raiser
AU 4	Brow Lowerer
AU 5	Upper Lid Raiser
AU 6	Cheek Raiser
AU 9	Nose Wrinkler
AU 12	Lip Corner Puller
AU 15	Lip Corner Depressor
AU 17	Chin Raiser
AU 23	Lip Tightener
AU 26	Jaw Drop
AU 41	Lid Droop
AU 43	Eyes Closed
AU 45	Blink
AU 51	Head Turn Left
AU 52	Head Turn Right
AU 53	Head Up
AU 54	Head Down
AU 61	Eyes Turn Left
AU 62	Eyes Turn Right
AU 63	Eyes Up
AU 64	Eyes Down

Table 1. List of the 21 selected Action Units (AUs) [13] with their corresponding descriptions, chosen based on their possible relevance to drowsiness and potential distractions in driving scenarios, as well as some emotions that could affect road safety.

performing the target AU, and then they were instructed to perform the AU at a natural speed. This method relies on the subject’s ability to mimic the demonstrated expression and our visual judgment. This is to highlight that the labels represent elicited performances intended to activate specific AUs, rather than expert-coded FACS activations.

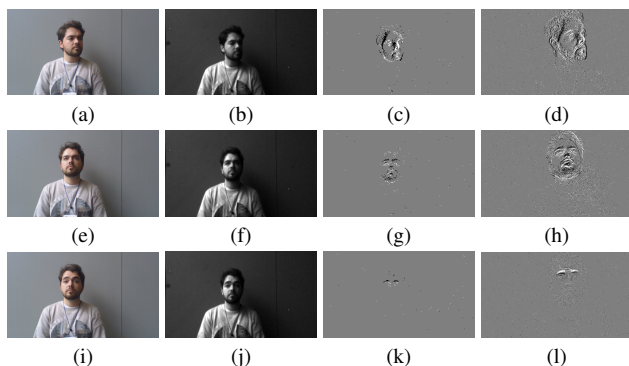


Figure 2. Multimodal Facial Micro-Expression Dataset. From left to right: ASUS webcam RGB frame; DAVIS346 frame; DAVIS346 event frame; EVK4 event frame. The frames show the apex of a) AU 51 (Head Turn Left), b) AU 53 (Head Turn Up), and c) AU 2 (Outer brow raiser).

3.2. Set-up

The experiments were set up in an auditorium with a fully windowed wall and plenty of direct sunlight during the day. Subjects participated in the experiment at different times of day, ranging from morning to late afternoon, resulting in non-uniform lighting conditions across recordings. Consequently, some participants were recorded with additional artificial ceiling illumination to compensate for suboptimal lighting during late afternoon. Facial illumination was measured for each subject, with lux levels ranging from 150 to 1400. Distractors might have been present during the recording since subjects were not completely isolated from other people passing by in the room.

To record each subject, we employed three cameras: a DAVIS346 event camera recording events and synchronized grayscale frames at 346×260 resolution at a maximum of 40 FPS, a Prophesee EVK4 event camera recording events at 1280×720 resolution, and an ASUS ROG Eye S RGB webcam with a resolution of 1920×1080 at a maximum of 60 FPS. All three devices were mounted on a fixed camera tripod with a special adaptation that allowed for horizontal alignment of the cameras. The cameras were connected via USB cable to a laptop, and a ROS script was used to ensure simultaneous recording and temporal synchronization. During the experiment, subjects sat comfortably in a chair in front of the cameras at a fixed distance of 80 cm, ensuring that their face was centered within the frame. Sequences of AUs were displayed on an external LCD monitor mounted on a desk adjacent to the camera tripod. Before the experiment, both event cameras went through a calibration process to estimate their intrinsic parameters.

3.3. Protocol

We introduced each subject to the experimental procedure, briefly discussing the goal of our work. After the subject agreed to sign the informed consent, the experimenter described the subset of AUs to be performed by the subject, then showed them short clips with clear textual descriptions of how to perform the actions. After confirming their understanding of each AU, subjects were then instructed to sit comfortably on a chair, which was adjusted to meet the necessary distance to the cameras. Before the experiment’s main run, a demonstration run was carried out to test the subject’s understanding of the entire task and the timing for performing the actions.

Once the subject was ready, the main run was executed. During each run, subjects were asked to look at the screen where a video instructed them which AU to perform and when. For each AU, the video showed a clip demonstrating the action three times. A countdown with sound indicated the subject to look at the camera and be ready to perform the action. After a sufficient period, a double beep and a message on the video instructed them to rest and relax for a few

seconds. Unless they warned the experimenter about a possible mistake, the next segment with the next AU played, and the subject performed the action in the same manner. The experimenter started recording each segment independently as soon as the countdown appeared on the video. If the subject or the experimenter made a mistake, the video was paused, and the action was re-done. Each main run lasted approximately 7 minutes, and each participant went through only one main run.

3.4. Dataset description

Our preliminary dataset consists of 7 individuals, six men and one woman, ages ranging from 18 to 27. For each of them, we selected 21 videos of elicited AU performances recorded in both RGB and Event modalities with different image resolutions, yielding a total of 510 clips with evenly distributed classes of AUs. Lighting conditions were not uniform across recordings, with lux ranging from 150 to 1400 using mainly natural illumination. A few samples obtained during our recordings can be seen in Figure 2, and the description of our dataset is summarized in Table 2.

We trimmed each clip, ensuring that only the elicited AU was present (from onset to apex and offset). In this way, we eliminated blinking and any other facial or body movements before and after the AU performance, leaving, when possible, about a second of frames of margin with no apparent activity. However, it is worth noting that some participants had difficulties isolating single Action Units. For example, some of them would move their mouth while performing AU 6 (cheek raiser) or even blink, AU 45, while doing another action. We decided to be less selective in this stage and keep these non-ideal clips, favoring quantity. As a consequence, some clips could contain co-activations of AU, but, in that case, only the elicited AU was labeled. The average duration of clips was 2.5s (including the safe margins of neutral activity before and after AU performance).

As a final step, we took advantage of the synchronized frames of the DAVIS346 and extracted the bounding box positions to crop each event frame using the MTCNN detector [41]. We note that it could also be possible to extract facial landmarks and bounding boxes from the EVK4 event frames using the synchronized stream from the ASUS RGB camera.

#Vids.	#Subj.	Modality	Resolution
147	7	RGB (ASUS)	1920 × 1080
147	7	Events (Davis)	346 × 260
147	7	RGB (Davis)	346 × 260
63	3	Events (Evk4)	1280 × 720

Table 2. Summary of our multi-modal, multi-resolution preliminary dataset. A total of 7 subjects were recorded, with only 3 using the complete camera setup.

4. Facial Micro-Expression Analysis Framework

4.1. Main problem formulation

Given an event stream recording from an event camera $E_i(t_N) = \{e_k\}_{k=1}^{N_i} = \{(x_k, y_k, t_k, p_k)\}_{k=1}^{N_i}$ and its label l_i representing the AU performed during the i -th recording of our dataset, we optimize the parameters of an SNN model F_θ for the single-label, multi-class classification problem consisting of 21 classes of AUs present in our dataset.

4.2. Event data representation

Each event $e_k = (x_k, y_k, t_k, p_k)$ is a 4-tuple representing a brightness change in pixel (x_k, y_k) at time t_k above certain threshold since the previous event at that pixel. Here, $p_k \in \{-1, 1\}$ indicates the sign of the brightness change. A widely used event representation, which allows the use of traditional computer vision techniques, consists of fixing a slice of time and integrating the events in that period into a frame. Assuming a slice of size T with starting and ending times t_k and t_l , an event frame F_p with two channels corresponding to polarities $p = 1$ and $p = -1$ is defined as:

$$F_p(x, y) = \sum_{i=t_k}^{t_l} \mathcal{I}_{p,x,y}(p_i, x_i, y_i), \quad (1)$$

where \mathcal{I} is the indicator function. One drawback of this representation is that much of the asynchrony of events is lost if T is too large, defying the nature of events. However, it is a straightforward way to transform the event stream into a synchronous frame representation, facilitating the use of common deep ANNs architectures.

4.3. Baseline Architecture for AU recognition

4.3.1. Spiking Neurons

LIF spiking neurons are the essential building block of the proposed SNNs due to their low computational cost. They can be seen as a simplified bio-inspired model resembling the dynamics of an RC circuit. The subthreshold dynamics of LIF Neurons are given by the differential equation:

$$\tau \frac{dV(t)}{dt} = -(V(t) - V_{rest}) + X(t) \quad (2)$$

where $V(t)$ is the membrane potential at time t , $X(t)$ is the input to neuron at time t , τ is the membrane time constant and V_{rest} is the resting potential. When $V(t)$ exceeds the threshold V_{th} , the neuron produces a spike and its membrane potential is set to V_{reset} . The discrete implementation of Eq. 2 can be derived as:

$$H[t] = V[t - 1] - \frac{1}{\tau}(V[t - 1] - V_{rest}) + X[t] \quad (3)$$

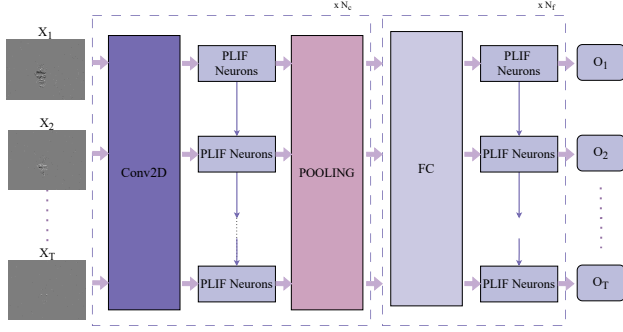


Figure 3. SNN architecture employed for AU recognition. Here, CONV2D represents a sequential block of 2D Convolution followed by Batch Normalization, and FC a sequential block of Dropout followed by a FC layer. The spiking encoder is composed of N_c blocks of Conv2D-PLIFNeurons-2DMaxPooling. The classification layers FC-PLIFNeurons are repeated N_f times. In this work, we set $N_c = N_f = 2$.

The membrane time constant τ is usually a hyperparameter of the network. However, based on the Parametric Leaky Integrate-and-Fire (PLIF) unit proposed in [14], we let $\frac{1}{\tau} = \text{sigmoid}(w)$ to be a learnable parameter, shared only within the neurons in the same layer. From Eq. 3 and setting $V_{reset} = 0$, it can be seen that LIF and PLIF neurons resemble the functionality of RNNs. They are both stateful; in the case of LIF, they "remember" the membrane potential $V(t)$; they "forget" according to the membrane time constant τ ; and they also take current input information $X(t)$ to update their next state. This highlights the potential use of PLIF SNNs to learn the spatial-temporal dynamics in micro-expressions.

4.3.2. Action-Unit recognition

We employed a convolutional architecture with PLIFs as its building block based on the formulations given by Fang et al. [14] and represented in Figure 3.

The proposed architecture consists of two sequential blocks of [Conv2D-BatchNorm-PLIFNeurons-MaxPool2D]. All Conv2D layers were set with a kernel size = 3, stride = 1, padding = 1, and output feature map of size 32. The 2D Pooling layers were set with kernel size = 2 and stride = 2. These blocks act as a learnable encoder to convert images to spikes. After the spiking encoder, we add two sequential blocks of [Dropout-FClayer-PLIFNeuron] to learn the classification task. The output feature sizes of the FC layers are 512 and 210, respectively. A final voting layer, implemented as an Average pooling layer of stride 10 and kernel size 10, takes the last output spikes and maps them to a final response corresponding to one of our 21 classes. When a sequence of T frames is consumed, we take the average along the time dimension, yielding a vector of average activations per class.




Figure 4. Conditional Variational Autoencoder architecture for face reconstruction. The model reconstructs the target frame (x) conditioned on the event-based representation (c) obtained from accumulated events.

To train the SNN using backpropagation, a surrogate gradient function is needed due to the non-derivability of the heavy-side step function that simulates the spikes. We chose arctangent as the smoothing function, and we employed Mean Squared Error (MSE) as the loss function.

4.4. Baseline Architecture for Frame Micro-expression Reconstruction

For intensity frame micro-expression reconstruction, we implement a cVAE as a baseline architecture. The model conditions on an event-based representation c and learns to reconstruct the corresponding ground-truth frame x , as shown in Figure 4.

Formally, the encoder $q_\phi(z | x, c)$ receives the target frame x together with its conditional event-based representation c and maps them into a latent Gaussian distribution. This distribution is parameterized by a mean $\mu_\phi(x, c)$ and a diagonal covariance given by $\sigma_\phi^2(x, c)$:

$$q_\phi(z | x, c) = \mathcal{N}(z | \mu_\phi(x, c), \text{diag}(\sigma_\phi^2(x, c))). \quad (4)$$

To enable backpropagation through the stochastic sampling process, the latent code z is obtained via the reparameterization trick. In this formulation, z is expressed as a deterministic function of $\mu_\phi(x, c)$, $\sigma_\phi(x, c)$, and a noise variable $\epsilon \sim \mathcal{N}(0, I)$:

$$z = \mu_\phi(x, c) + \sigma_\phi(x, c) \odot \epsilon. \quad (5)$$

The decoder $p_\theta(x | z, c)$ then takes the latent vector z concatenated with the conditional input c and generates the reconstructed frame. The output \hat{x} represents the approximation of the original target frame x :

$$\hat{x} = p_\theta(x | z, c). \quad (6)$$

In practice, the encoder consists of two convolutional layers with ReLU activations followed by fully connected layers that produce μ and σ with latent dimensionality 32. The decoder receives the sampled latent vector concatenated with the flattened conditional input and generates \hat{x} through fully connected and transposed convolutional layers with sigmoid activation. The network is optimized by minimizing

the conditional Evidence Lower Bound (ELBO):

$$\mathcal{L}(\theta, \phi; x, c) = \mathbb{E}_{q_\phi(z|x,c)} [\log p_\theta(x | z, c)] - D_{\text{KL}}(q_\phi(z | x, c) \| p(z)), \quad (7)$$

with $p(z) = \mathcal{N}(0, I)$. To evaluate the reconstruction performance, multiple image similarity metrics were calculated, including Mean Squared Error (MSE), Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index Measure (SSIM), and Normalized Cross Correlation (NCC). These metrics evaluate the pixel-wise accuracy, perceptual quality, and structural consistency of the reconstructed frames compared to the webcam ground truth.

5. Training Details

The present section describes the training details of the proposed facial micro-expression analysis framework.

For AU recognition, given the reduced size of our preliminary dataset, we performed Leave-One-Out Cross-Validation per subject, ensuring no data leakage. For both image modalities, cropped frames of facial bounding boxes were resized to 64×64 . The SNNs were trained using the ADAM [17] optimizer, Mean Squared Error (MSE) as the loss function, and cosine annealing scheduler [21] during 100 epochs. After brief hyperparameter exploration, the best found models were trained with a learning rate of 0.001, setting dropout to 0.2, and a batch size of 8 due to GPU memory constraints. To have a comparative baseline against a similar model, in terms of size and speed, we also trained a 3DCNN architecture based on [34] adapted for variable clip length. The training configuration was the same as the SNNs but with a smaller learning rate of 10^{-4} .

For these preliminary experiments, we decided to set the slice size of the event representation to 33 ms, yielding a comparable frame rate to that of typical RGB cameras. Since event data from the Prophesee EVK4 was only available for a few subjects at the moment of writing, we decided to focus the analysis of Action Unit recognition solely on the other two cameras. All experiments were implemented in PyTorch and ran on an Intel Xeon server equipped with two NVIDIA RTX 3080 GPUs. The hardware specification is detailed in Table 3.

Table 3. Hardware specifications of the server used to train the SNN for AU recognition.

CPU	Max Freq.	#Threads	Cache
Xeon® E5-2680 v3	3.30 GHz	24	30MB
GPU	Boost clock	Bandwidth	VRAM
NVIDIA® RTX 3080	1.71 GHz	760GB/sec	12GB

For frame micro-expression reconstruction training, we used an Adam optimizer with a learning rate of 10^{-3} , batch size of 16, and trained each model for 100 epochs. Two datasets were used: one using DVS-events as conditional input, and the other using EVK4-events. The loss function combines MSE reconstruction loss and the Kullback–Leibler divergence term from the VAE formulation.

The large volume of asynchronous events requires substantial computational resources and leads to long training times, particularly for the frame reconstruction task. To handle this, the cVAE models were trained on an NVIDIA® DGX-1™ system optimized for deep learning workloads. The hardware specifications are detailed in Table 4.

Table 4. Hardware specifications of the NVIDIA® DGX-1™ system used to train the Conditional Variational Autoencoder for event-based micro-expression frame reconstruction.

Hardware	Accelerator	Architecture	Boost clk
NVIDIA DGX-1	V100	Volta	1530 MHz
Memory Clock	Bandwidth	VRAM	GPU
1.75Gbit/s	900GB/sec	32GB	GV100

6. Results

This section reports the preliminary results for micro-expression recognition and intensity frame reconstruction.

6.1. Event-based Action Unit recognition

We evaluated the performance of the SNN model in AU recognition on our preliminary dataset, and we compared both image modalities. Our results show that the selected architectures perform significantly better when trained with event-based frames rather than RGB, as shown in Table 5. This is in line with previous findings [1, 9] where it was shown that the event modality outperforms RGB in AU recognition using spatio-temporal models like 3D CNNs and CNNs + LSTM/Transformer.

Model	#Par.	Modal.	Acc. (%)	Top-3 Acc. (%)
SNN	4.11M	events	51.23	63.57
		rgb	23.12	25.57
3DCNN	13.06M	events	44.89	56.46
		rgb	17.06	21.76

Table 5. Accuracy results for AU recognition after LOOCV for both models and modalities (webcam RGB and Davis Events).

Additionally, the SNN outperformed the 3DCNN model by a margin of roughly 6% in Accuracy, showing a promising direction to deepen the study of bio-inspired networks

in micro-expression analysis. Nevertheless, we emphasize that these results, although encouraging, are still preliminary since the data collected lacks diversity and size; there are just seven samples for each Action Unit.

An important limitation of the experiment is the fixed time window selected to integrate events. This parameter could play a key role in the way events are interpreted, and thus it should be addressed carefully in future experiments.

6.2. Frame micro-expression reconstruction conditioned by event frames

The qualitative results in Figure 5 show that both models attempt to reconstruct the target webcam frames from event-based inputs, preserving the general facial structure and expressions. However, differences in sharpness and noise patterns are evident between the DAVIS346 and EVK4 inputs, with EVK4-based reconstructions appearing visually smoother and closer to the ground truth.




Figure 5. Examples of event-based micro-expression reconstruction of intensity frames. From left to right, each image shows the input event-based frame, the reconstructed frame, and the ground truth frame. The images (a,b,c) use the DAVIS346 sensor as input and (e,d,f) use the EVK4 sensor.

Table 6. Reconstruction metrics for DAVIS346-events and EVK4-events to grayscale webcam frame.

Metric	DAVIS346	EVK4
Model Params.	289M	289M
Mean MSE	0.0034	0.0024
Mean RMSE	0.0566	0.0471
Mean MAE	0.0399	0.0340
Mean PSNR	25.31 dB	26.89 dB
Mean SSIM	0.7717	0.8513
Mean NCC	0.9613	0.9744

Quantitatively (Table 6), the EVK4-conditioned model achieves lower reconstruction errors across all metrics, with a Mean MSE of 0.002445 and a higher SSIM of 0.8513,

indicating better pixel-wise accuracy and perceptual quality compared to the DAVIS346-conditioned model (SSIM = 0.7717). Despite these differences, it is important to note that the reconstructed frames still lack the level of detail, texture consistency, and identity-preserving features required for reliable downstream tasks. Further advances in model architecture and training strategies will be necessary to improve reconstruction quality from event-based inputs.

7. Conclusions and Future Work

This work explored the spatial-temporal dynamics of facial micro-expressions through multimodal data acquisition with RGB and event cameras. We introduced a novel, preliminary dataset recorded under varying lighting conditions with multiple event-based sensors, providing a first step towards systematic evaluation of Action Units in this modality. The dataset was used to study two baseline tasks: (a) Action Unit recognition using Spiking Neural Networks trained with surrogate gradient descent, which achieved 51.23% accuracy with event data compared to 23.12% with RGB inputs; and (b) frame reconstruction with Conditional Variational Autoencoders, where EVK4-conditioned models outperformed DAVIS346 in terms of perceptual similarity (SSIM = 0.8513, PSNR = 26.89 dB) but still lacked detail and identity preservation.

Despite these promising results, several limitations remain. The dataset is constrained by the number of participants, resulting in gender imbalance and underrepresentation of specific age groups. Event recordings from the EVK4 were only available for a subset of subjects, reducing the consistency of cross-sensor analysis. The dataset also lacks expert-coded AU activations and activation levels, as recordings were based on elicited performances rather than certified annotations.

Future work will address these limitations by expanding the dataset with a broader participant pool, more balanced demographics, and additional recording conditions such as illumination and pose variability. Beyond dataset growth, promising directions include the design of new event-based facial representations, more advanced generative models for reconstruction, and improved methods for Action Unit recognition. Finally, multimodal fusion of RGB and event streams offers a compelling path to leverage complementary information and further advance micro-expression analysis.

Acknowledgments

This work was partially funded by FONDEQUIP Project EQM170041 and the 2025 IEEE CIS Graduate Student Research Grants. We thank all participants who contributed to the database and the University of Chile as the venue.

References

- [1] Mira Adra, Nélide Mirabet-Herranz, and Jean-Luc Dugelay. Beyond rgb: Tri-modal microexpression recognition with rgb, thermal, and event data. In *International Conference on Pattern Recognition*, 2024. 2, 7
- [2] Shafiq Ahmad, Pietro Morerio, and Alessio Del Bue. Person re-identification without identification via event anonymization. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 11098–11107, 2023. 2
- [3] Souptik Barua, Yoshitaka Miyatani, and Ashok Veeraraghavan. Direct face detection and video reconstruction from event cameras. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–9, 2016. 3
- [4] Federico Becattini, Lorenzo Berlincioni, Luca Cultrera, and A. Bimbo. Neuromorphic face analysis: a survey. *Pattern Recognit. Lett.*, 187:42–48, 2024. 2
- [5] Federico Becattini, Luca Cultrera, Lorenzo Berlincioni, Claudio Ferrari, Andrea Leonardo, and A. Bimbo. Neuromorphic facial analysis with cross-modal supervision. *ArXiv*, abs/2409.10213, 2024. 2
- [6] Lorenzo Berlincioni, Luca Cultrera, Chiara Albisani, Lisa Cresti, Andrea Leonardo, Sara Picchioni, Federico Becattini, and A. Bimbo. Neuromorphic event-based facial expression recognition. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 4109–4119, 2023. 2
- [7] Lorenzo Berlincioni, Luca Cultrera, Federico Becattini, and A. Bimbo. Neuromorphic valence and arousal estimation. *ArXiv*, abs/2401.16058, 2024. 3
- [8] Guang Chen, Lin Hong, Jinhua Dong, Peigen Liu, Jörg Conradt, and Alois Knoll. Eddd: Event-based drowsiness driving detection through facial motion analysis with neuromorphic vision sensor. *IEEE Sensors Journal*, 20:6170–6181, 2020. 1
- [9] Luca Cultrera, Federico Becattini, Lorenzo Berlincioni, Claudio Ferrari, and Alberto Del Bimbo. Spatio-temporal transformers for action unit classification with event cameras, 2024. 7
- [10] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, Gautham N. Chinya, Yongqiang Cao, Sri Harsha Choday, Georgios D. Dimou, Prasad Joshi, Nabil Imam, Shweta Jain, Yuyun Liao, Chit-Kwan Lin, Andrew Lines, Ruokun Liu, Deepak A. Mathaikutty, Steve McCoy, Arnab Paul, Jonathan Tse, Guruguanathan Venkataramanan, Yi-Hsin Weng, Andreas Wild, Yoonseok Yang, and Hong Wang. Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, 38:82–99, 2018. 3
- [11] Yiting Dong, Dongcheng Zhao, Yang Li, and Yi Zeng. An unsupervised stdp-based spiking neural network inspired by biologically plausible learning rules and connections. *Neural networks : the official journal of the International Neural Network Society*, 165:799–808, 2022. 3
- [12] Paul Ekman. Universal and cultural differences in facial expression of emotion. In *Nebraska Symposium on Motivation*, pages 207–283, 1971. 1
- [13] Paul Ekman and Wallace V. Friesen. Facial action coding system: a technique for the measurement of facial movement. 1978. 2, 4
- [14] Wei Fang, Zhaofei Yu, Yanqing Chen, Timothée Masquelier, Tiejun Huang, and Yonghong Tian. Incorporating learnable membrane time constant to enhance learning of spiking neural networks. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2641–2651, 2020. 3, 6
- [15] Wallace V. Friesen and Paul Ekman. Emfacs-7: Emotional facial action coding system. 1983. 3
- [16] Yuhuang Hu, Shih-Chii Liu, and Tobi Delbruck. v2e: From video frames to realistic dvs events. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, pages 1312–1321, 2021. 2
- [17] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *CoRR*, abs/1412.6980, 2014. 7
- [18] Thomas Kopalidis, Vassilios Solachidis, Nicholas Vretos, and Petros Daras. Advances in facial expression recognition: A survey of methods, benchmarks, models, and datasets. *Inf.*, 15:135, 2024. 1
- [19] Beibin Li, Sachin Mehta, Deepali Aneja, Claire E. Foster, Pamela Ventola, Frédéric Shic, and Linda G. Shapiro. A facial affect analysis system for autism spectrum disorder. *2019 IEEE International Conference on Image Processing (ICIP)*, pages 4549–4553, 2019. 1
- [20] Yante Li, Jinsheng Wei, Yang Liu, Janne Kauttonen, and Guoying Zhao. Deep learning for micro-expression recognition: A survey. *IEEE Transactions on Affective Computing*, 13:2028–2046, 2021. 2
- [21] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv: Learning*, 2016. 7
- [22] Patrick Lucey, Jeffrey F. Cohn, Takeo Kanade, Jason M. Saragih, Zara Ambadar, and I. Matthews. The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition - Workshops*, pages 94–101, 2010. 2
- [23] Wolfgang Maass. Networks of spiking neurons: The third generation of neural network models. *Electron. Colloquium Comput. Complex.*, TR96, 1996. 2
- [24] Paul Merolla, John V. Arthur, Rodrigo Alvarez-Icaza, Andrew S. Cassidy, Jun Sawada, Filipp Akopyan, Bryan L. Jackson, Nabil Imam, Chen Guo, Yutaka Nakamura, Bernard Brezzo, Ivan Vo, Steven K. Esser, Rathinakumar Appuswamy, Brian Taba, Arnon Amir, Myron Flickner, William P. Risk, Rajit Manohar, and Dharmendra S. Modha. A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, 345:668 – 673, 2014. 3
- [25] Nico Messikommer, Daniel Gehrig, Antonio Loquercio, and Davide Scaramuzza. Event-based asynchronous sparse convolutional networks. *ArXiv*, abs/2003.09148, 2020. 2
- [26] Emre O. Neftci, Hesham Mostafa, and Friedemann Zenke. Surrogate gradient learning in spiking neural networks: Bringing the power of gradient-based optimization to spiking neural networks. *IEEE Signal Processing Magazine*, 36: 51–63, 2019. 2

- [27] Liyuan Pan, Richard Hartley, Cedric Scheerlinck, Miaomiao Liu, Xin Yu, and Yuchao Dai. High Frame Rate Video Reconstruction Based on an Event Camera. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 44(05):2519–2533, 2022. 3
- [28] Maja Pantic, Michel F. Valstar, Ron Rademaker, and Ludo Maat. Web-based database for facial expression analysis. *2005 IEEE International Conference on Multimedia and Expo*, pages 5 pp.–, 2005. 2
- [29] Robert Plutchik. A general psychoevolutionary theory of emotion. 1980. 2
- [30] Silvia Ramis, Jose Maria Buades Rubio, and Francisco J. Perales López. Using a social robot to evaluate facial expressions in the wild. *Sensors (Basel, Switzerland)*, 20, 2020. 1
- [31] Henri Rebecq, Daniel Gehrig, and Davide Scaramuzza. Esim: an open event camera simulator. In *Conference on Robot Learning*, 2018. 2
- [32] Cedric Scheerlinck, Henri Rebecq, Daniel Gehrig, Nick Barnes, Robert E. Mahony, and Davide Scaramuzza. Fast image reconstruction with an event camera. In *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 156–163, 2020. 3
- [33] Warren D. TenHouten. Basic emotion theory, social constructionism, and the universal ethogram. *Social Science Information*, 60:610 – 630, 2021. 2
- [34] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 4489–4497, 2015. 7
- [35] Ajay Vasudevan, Pablo Negri, Camila Di Ielsi, Bernabé Linares-Barranco, and Teresa Serrano-Gotarredona. S1-animals-dvs: event-driven sign language animals dataset. *Pattern Analysis and Applications*, 25:505 – 520, 2021. 3
- [36] Rodrigo Verschae and Ignacio Bugueno-Cordova. Event-based gesture and facial expression recognition: A comparative analysis. *IEEE Access*, 11:121269–121283, 2023. 2
- [37] Rodrigo Verschae and Ignacio Bugueno-Cordova. evtransfer: A transfer learning framework for event-based facial expression recognition, 2025. 3
- [38] Malaika Vijay, Nandagopal Netrakanti Vinayak, Maanvi Nunna, and Natarajan Subramanyam. Real-time driver drowsiness detection using facial action units. *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 10113–10119, 2021. 1
- [39] Esra Vural, Müjdat Çetin, Aytül Erçil, Gwen Littlewort, Marian Stewart Bartlett, and Javier R. Movellan. Drowsy driver detection through facial movement analysis. In *ICCV-HCI*, 2007. 3
- [40] Guoliang Xiang, Song Yao, Hanwen Deng, Xianhui Wu, Xinghua Wang, Qian Xu, Tianjian Yu, Kui Wang, and Yong Peng. A multi-modal driver emotion dataset and study: Including facial expressions and synchronized physiological signals. *Eng. Appl. Artif. Intell.*, 130:107772, 2024. 1
- [41] Jia Xiang and Gengming Zhu. Joint face detection and facial expression recognition with mtcnn. *2017 4th International Conference on Information Science and Control Engineering (ICISCE)*, pages 424–427, 2017. 5
- [42] Kashu Yamazaki, Viet-Khoa Vo-Ho, D Bulsara, and Ngan T. H. Le. Spiking neural networks and their applications: A review. *Brain Sciences*, 12, 2022. 3
- [43] Wen-Jing Yan, Qi Wu, Jing Liang, Yu-Hsin Chen, and Xiaolan Fu. How fast are the leaked facial expressions: The duration of micro-expressions. *Journal of Nonverbal Behavior*, 37:217–230, 2013. 2
- [44] Dingkan Yang, Shuai Huang, Zhi Xu, Zhenpeng Li, Shunli Wang, Mingcheng Li, Yuzheng Wang, Yang Liu, Kun Yang, Zhaoyu Chen, Yan Wang, Jing Liu, Pei Zhang, Peng Zhai, and Lihua Zhang. Aide: A vision-driven multi-view, multi-modal, multi-tasking dataset for assistive driving perception. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 20402–20413, 2023. 1
- [45] Jiabei Zeng, S. Shan, and Xilin Chen. Facial expression recognition with inconsistently annotated datasets. In *European Conference on Computer Vision*, 2018. 2
- [46] Zhongyang Zhang, Shuyang Cui, Kaidong Chai, Haowen Yu, Subhasish Dasgupta, Upal Mahbub, and Tauhidur Rahman. V2ce: Video to continuous events simulator. *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12455–12461, 2023. 3
- [47] Yunhao Zou, Yinqiang Zheng, Tsuyoshi Takatani, and Ying Fu. Learning to reconstruct high speed and high dynamic range videos from events. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2024–2033, 2021. 3