

# Transformer-based Fusion of 2D-pose and Spatio-temporal Embeddings for Distracted Driver Action Recognition

Erkut Akdag<sup>†</sup>, Zeqi Zhu, Egor Bondarev, Peter H.N. de With  
VCA Group, Department of Electrical Engineering, Eindhoven University of Technology  
P.O. Box 513, Eindhoven 5612AZ, The Netherlands

<sup>†</sup> corresponding author

{e.akdag, z.zhu, e.bondarev, p.h.n.de.with}@tue.nl

## Abstract

Classification and localization of driving actions over time is important for advanced driver-assistance systems and naturalistic driving studies. Temporal localization is challenging because it requires robustness, reliability, and accuracy. In this study, we aim to improve the temporal localization and classification accuracy performance by adapting video action recognition and 2D human-pose estimation networks to one model. Therefore, we design a transformer-based fusion architecture to effectively combine 2D-pose features and spatio-temporal features. The model uses 2D-pose features as the positional embedding of the transformer architecture and spatio-temporal features as the main input to the encoder of the transformer. The proposed solution is generic and independent of the camera numbers and positions, giving frame-based class probabilities as output. Finally, the post-processing step combines information from different camera views to obtain final predictions and eliminate false positives. The model performs well on the A2 test set of the 2023 NVIDIA AI City Challenge for naturalistic driving action recognition, achieving the overlap score of the organizer-defined distracted driver behaviour metric of 0.5079.

## 1. Introduction

In recent years, the advanced driving-assistance systems (ADAS) development has garnered significant attention for its promise to improve road safety and reduce accidents. However, accurately identifying and locating hazardous events, particularly in complex driving scenarios, remains a challenge for ADAS development. The recognition of distracted driver activity is an important objective with practical relevance in this area, and it is a crucial cause for fatalities in traffic accidents. For instance, distracted driver activities cause nine fatalities in the United States daily,

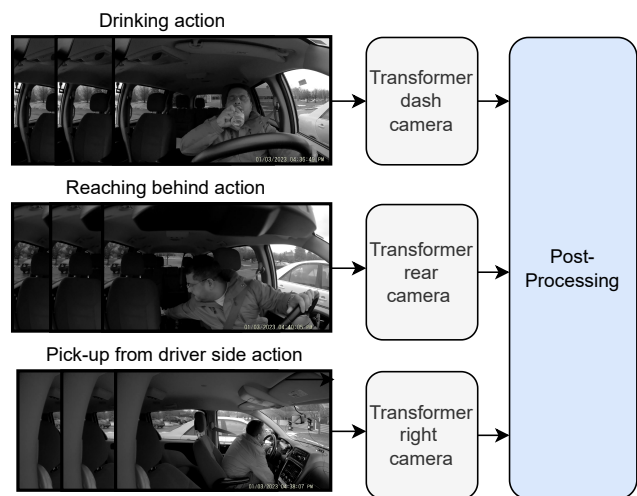


Figure 1. Overview of the proposed transformer-based method. From left top to bottom: example images of the "drinking", "reaching behind", and "pick-up from driver side" action classes.

based on National Highway Traffic Safety Administration statistics [1]. Therefore, identifying and addressing such driver behaviour is vital for ADAS and can potentially reduce deaths in traffic incidents.

With recent deep learning advances, the research on modelling systems is increasingly based on examining a substantial volume of video data. One important analysis domain is detection of actions in untrimmed or extended videos, which enables recognizing distracted driver behavior and avoiding accidents. In the last decade, immense research has been conducted, especially on advanced video action recognition techniques to facilitate robust action classification. Nevertheless, this problem still needs further work to be solved due to its challenging nature. For instance, complex motion patterns, actions with similar visual appearances, and varying duration make this problem quite challenging.

Distracted driver actions take place in different direc-

tions in the space surrounding the driver. To understand these actions for reliable recognition, several cameras are used in different positions. The actions captured by the individual camera views are learned, and then combined Fig. 1. In this paper, we propose a transformer-based fusion of 2D-pose and spatio-temporal features to efficiently predict the temporal localization of distracted driver activity. Specifically, we first extract 2D-pose and spatio-temporal features from each video. Secondly, to take advantage of multi-modality of the feature data, we apply a transformer-based model for the feature fusion. 2D-Pose features are given as a position embedding at the encoder stage, named "POSEition" embedding throughout the paper, while spatio-temporal features are the main input to the encoder. During the training step, we use the density-guided label-smoothing technique proposed earlier in [2]. In the final step, the data from transformer models which are trained on different camera views, is combined and a post-processing step is applied for obtaining temporal localization of distracted driver actions. The method is independent of camera numbers inside the vehicle, making it a generic solution. Summarizing, this paper provides the following contributions

- Novel solution for the distracted driver action recognition, based on a transformer model that is independent of the amount of in-vehicle cameras.
- Efficient feature extraction from 2D-pose estimation including the key points of the face and hand.
- Fusion of 2D-pose features and video action features by the encoder module with multi-head attention.

The remainder of the paper is organized as follows. Section 2 reviews related work on 2D-pose estimation and human activity recognition. In Section 3, we describe the proposed methodology in detail. Section 4 presents the experimental settings and obtained results, including the ablation studies. Finally, concluding remarks are given in Section 5.

## 2. Related Work

We review the closely related literature on human activity recognition from three perspectives: video action recognition, distracted driver action recognition, and modalities for feature representation.

**Video action recognition** Action recognition has been the subject of significant research in recent years, with various approaches being used [10, 21, 37, 39, 41]. This field has progressed with modern benchmarks on the Kinetics dataset, offering a larger number of categories and videos. Over the past two decades, the trend in action recognition has been to gather larger datasets and develop larger models (such as C3D [35], I3D [4], and SlowFast [11]). In addition to this, several backbone architectures, such as

VGG [32], ResNet [12], and DenseNet [14], have been developed to enhance precision and efficiency. While CNNs remain the primary models for this task, Vision Transformers are gradually gaining prominence due to their high potential. The ViT model (Vision Transformer) [6] applies the transformer architecture directly to image classification and has shown promising results. In recent years, ViT and its variants, including MViT [9] and Swin Transformer [24], have achieved outstanding performance in action recognition.

**Distracted driver action recognition** Distracted driving is a severe problem that has received increasing attention in recent years. Researchers have investigated various approaches to mitigate the effects of distracted driving. The work in [8] introduces the distracted driver dataset (DDD) involving 44 drivers. Authors train multiple convolutional neural network (CNN) architectures and use a learnable weighted ensemble technique for real-time driver distraction identification. Another work [27] publishes the Drive&Act dataset, including 12 distraction-related actions captured by six cameras. The authors show that their video and body pose-based action recognition algorithm gives competitive results. In the study by [18], the authors propose a contrastive learning approach to learn a metric to differentiate normal driving from abnormal driving. To achieve this, they introduce the Driver Anomaly Detection (DAD) dataset, which consists of 31 subjects that perform various activities in a real car.

**Modalities for feature representation** In previous studies, various feature modalities have been explored and integrated to enhance the performance of video action recognition. An RGB frame, as an atomic video component, is the most common input modality [12, 16, 34] for image classification. Ng et al. [40] extracted features independently from each RGB frame using 2D-ConvNet, followed by LSTM cells [13] for a sequence of ConvNet outputs, to encode state and capture temporal ordering and long-range dependencies. Du et al. [36] have used segmented RGB volumes from the video as inputs for 3D-ConvNet to classify actions, as temporal information is propagated across all the layers in the network, allowing 3D-ConvNet to model temporal information better than 2D-ConvNet. In addition to RGB frames, Simonyan and Zisserman [31] have introduced optical flow to explicitly describe the human motion between video frames. Complementary modalities of motion significantly boost model performance in action recognition. Recently, the human skeleton data has received increasing attention for its action-oriented nature and compactness. Most works [3, 5, 7, 15, 19, 23, 25] extract body skeletons from raw images via 2D-CNN-based approaches and convert the coordinates in a skeleton sequence to a pseudo-image with transformations. The skeleton features can be fused with other modalities and further processed

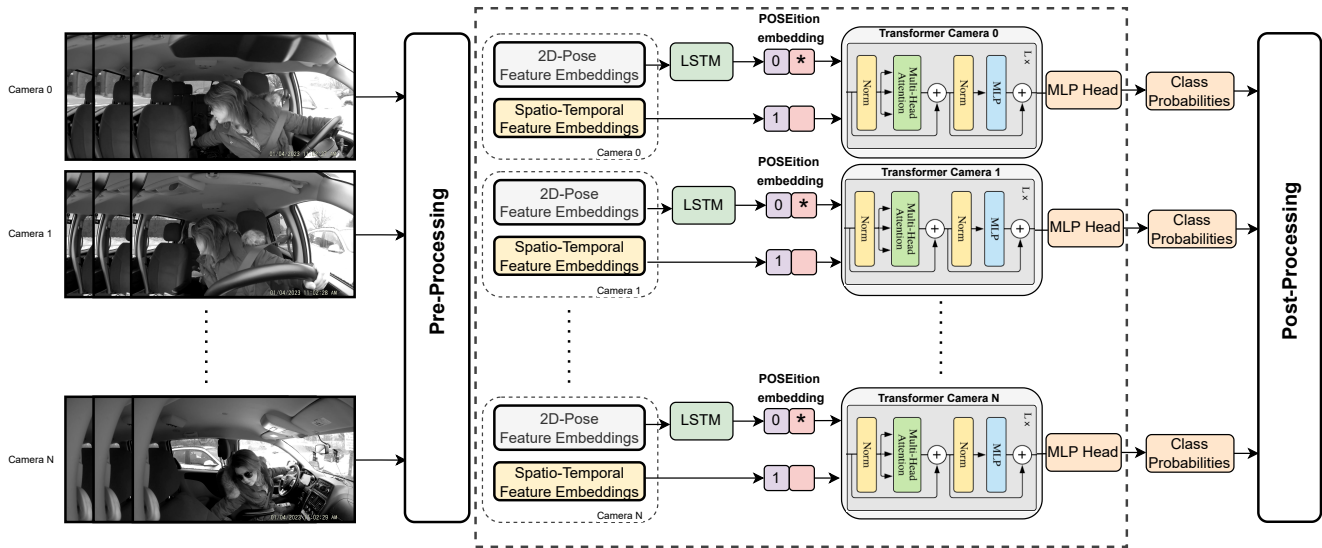


Figure 2. Overview of the proposed architecture. Left: different camera-view inputs to the pre-processing step. Middle: extracted 2D-pose and spatio-temporal embeddings are supplied to the transformer architecture. 2D-Pose embedding is considered as the "POSEition" embedding of the encoder, while the spatio-temporal embedding is the main input of the encoder. Right: 1D-class probabilities obtained after the MLP head per camera view are analyzed for finding the significant peaks for each class in a video. Note that the modules shown with a bold dashed border are used for training only.

by graph convolutional network (GCN)-based [38] or 3D-CNN-based approaches [7, 23], achieving superior performance on action recognition benchmarks.

Previous works have demonstrated that representations from different modalities, such as RGB and skeletons, are complementary. In this work, we propose to integrate the pose features of the driver face and body parts with spatio-temporal video features to enhance performance in action recognition.

### 3. Methodology

We propose a transformer-based end-to-end fusion solution to exploit spatio-temporal features aligned with 2D-pose features, as shown in Figure 2. In a pre-processing step, we re-size each video frame to  $512 \times 512$  pixels and feed our model to extract 2D-pose features and the spatio-temporal features. The proposed architecture takes 2D-pose features as position embedding, named "POSEition" and spatio-temporal features as the main input to the encoder of the transformer-based fusion model. In the final step, post-processing of the output probabilities of frames in each video identifies significant peaks, related to each class in the corresponding video.

#### 3.1. 2D-Pose Features

Previous studies [7, 42] have highlighted the critical importance of 2D-pose extraction in the pre-processing step of skeleton-based action recognition. Even small pertur-

bations in the joint coordinates can lead to vastly different predictions. To ensure accurate pose extraction, we adopt a 2D Top-Down pose estimator [33], which has been shown to outperform its 2D Bottom-Up counterparts on standard benchmarks [22]. As described in Figure 3, we first utilize a Faster R-CNN [30] human detector which is pre-trained on the COCO dataset, to identify the region of the driver in the raw RGB frame. Next, we input the detected region to the Top-Down pose estimator HRNet [33], which generates the whole-body joint coordinates of the driver. Since the focus is solely on a single driver, the Top-Down pose estimator offers higher accuracy than the Bottom-Up estimator, while requiring the same processing time.

For the whole-body pose estimation, a total of 133 key



Figure 3. Top-Down pose estimation.

points are used (17 for the body, 6 for the feet, 68 for the face, and 42 for the hands) [20]. However, many of these key points are irrelevant and do not aid in action recognition. For instance, when the driver’s body is occluded by the steering wheel (as shown in Fig. 6), the estimation of key points for both legs can be either inaccurately estimated or completely missed, which is likely to interfere with the final prediction results. Therefore, we select a subset of critical key points from body parts that correspond highly to distracted driving activities, such as the face, shoulders and hands. As illustrated in Tab. 1, actions such as drinking, phone calling, eating, texting, reaching behind, and picking-up objects are strongly associated with hand movements, while head pose and facial expression are highly relevant for actions like singing, yawning, talking, drinking, and phone calling.

The driver’s skeleton can be represented as a set of joint coordinates, where  $J_i = (x_i, y_i, c_i)$  denotes the location and confidence score of a detected joint  $i$ ,  $i = 1..N_j$ . In time, a skeleton sequence of  $T_c$  frames can be represented as a  $T_c \times N_j \times 3$  array, where  $N_j$  is the number of selected joints per skeleton. The joint coordinates alone cannot ex-




Figure 4. Relative distances between hand and face points (left), and set of facial feature points (right).



Figure 5. Head pose of a driver during a drinking action.

PLICITLY convey motion information between body parts to the network. To address this, we focus on highly active body parts (face and hands) and incorporate additional features into the pose feature embeddings. These features include the pose of the head  $P^h$ , the relative positions of the left and right hands to the face  $D^h$ , and the relative position of the upper lip to the lower lip  $D^l$ , as depicted in Fig. 4 and Fig. 5.

To estimate the driver’s head pose, we implement the Perspective-n-Point (PnP) pose computation method [26]. The resulting rotation and translation vectors are used to draw the orthogonal line  $P_h$  of the driver’s face, which indicates the viewing direction of the driver. In the visual representation provided in Fig. 5, the variation in the length and position of the blue line in accordance with the driver’s head movement, provides a clear indication of the driver’s head pose.

Table 1. Classification of driver actions based on the presence of the driver’s head, mouth, and hands movements. The ✓ symbol indicates the presence of a movement, while an empty cell indicates its absence.

Class	Head	Mouth	Hands
Normal driving			✓
Drinking	✓	✓	✓
Phone call (right hand)		✓	✓
Phone call (left hand)		✓	✓
Eating			✓
Texting (right hand)	✓		✓
Texting (left hand)	✓		✓
Reaching behind	✓		✓
Adjusting control panel			✓
Picking up (driver)	✓		✓
Picking up (passenger)	✓		✓
Talking to passenger (right)	✓	✓	
Talking to passenger (back)	✓	✓	
Yawning		✓	✓
Hand on head	✓		✓
Singing	✓	✓	✓

To capture the relative distance between two body joints  $J_i$  and  $J_j$  with Cartesian coordinates  $(x_i, y_i)$  and  $(x_j, y_j)$  in the 2D plane, we calculate the distance between joints  $J_i$  and  $J_j$  using Eq. (1). By considering the relative distance between joints, we can describe the temporal dynamics of the driver’s mouth and hands, allowing for a more comprehensive understanding representation of the driver’s behaviour in terms of mouth and hands. The distance between the joints  $J_i$  and  $J_j$  is specified by

$$D(J_i, J_j) = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}. \quad (1)$$

To incorporate movement information of the above ac-



Figure 6. Missed (left) and flawed (right) leg estimations for the driver.

tive body parts, we collect the aforementioned pose and distance features at each frame and arrange them as a chain of motion vectors over time  $T_c$ . The motion vector is described by  $fm_t = (P^h, D_0^h, \dots, D_L^h, D_0^l, \dots, D_K^l)_t$ , where  $t$  denotes the time step,  $L$  is the number of relative distances between hand joints and face joints,  $K$  is the number of relative distances between upper-lip joints and bottom-lip joints. We then concatenate the selected skeleton coordinates with the motion vector and use the resulting feature embeddings of the  $i^{th}$  video,  $f\text{-pose}_t^i = (J_0, \dots, J_{N_j}, P^h, D_0^h, \dots, D_L^h, D_0^l, \dots, D_K^l)_t^i$ , as the additional input stream for the proposed network. This allows the model to leverage both the spatial and temporal dynamics of the driver’s behavior for improving the prediction accuracy.

### 3.2. Spatio-temporal Features

To extract spatio-temporal features, we employ the SlowFast video action recognition network pre-trained on the Kinetics-400 dataset [17]. This network has a multi-branch architecture that captures both short-term and long-term relations, by extracting one low-rate input sequence (fast pathway) and one high-rate input sequence (slow pathway) through temporal sub-sampling of the video frames. The fast pathway helps to describe long-term actions, while the slow pathway is better suited for short-term actions. Considering multiple temporal resolutions enables to extract features robustly for both types of actions. The SlowFast architecture takes every video segment of the provided training videos as input. Formally, assume that  $V_s^i \in \mathbb{R}^{h \times w \times T_c}$  denotes a video segment that consists of  $T_c$  frames at the time interval  $[t, t + T_c - 1]$  of the  $i^{th}$  video with horizontal and vertical resolution of  $h$  and  $w$ , respectively. The output spatio-temporal feature vector,  $f\text{-spatem}_t^i = S(V_s^i) \in \mathbb{R}^{N_f}$ , is used for further capturing spatio-temporal embeddings for the related video segment, where  $S$  is the functional form of the SlowFast backbone and  $N_f$  is the resulting

feature size. In the proposed solution, SlowFast extracts spatio-temporal embeddings for each 64-frame segment of the video clip with a temporal stride of unity. The size of the output feature vector,  $N_f$ , is 2,304 for each video segment.

### 3.3. Fusion of 2D-pose and Spatio-temporal Embeddings

Feature concatenation and ensembling have been widely used in various machine learning applications, and their effectiveness have been demonstrated in numerous studies. Feature concatenation involves combining multiple features or variables into a single input, which can help capturing complex relationships between them. Although feature concatenation and ensembling can be efficient techniques for improving the performance of machine learning models, there are also some potential drawbacks that should be considered. For instance, if the concatenated features are not carefully selected or pre-processed, this can increase data noise which can harm the model’s performance. These drawbacks can make it more challenging to interpret the contributions of individual features to the model’s predictions. Furthermore, if the data modalities used for ensembling are too similar or have similar weaknesses, this may not improve performance. For instance, ensembling only spatio-temporal models or 2D-pose-based models may not improve the results as expected. Therefore, we propose an efficient transformer-based architecture to fuse the 2D-pose features and spatio-temporal features to address the aforementioned drawbacks. By using this fusion approach, we improve the robustness of the model, and obtain a more accurate overall prediction and reduce the risk of over-fitting.

The dashed box in the middle of the architecture in Figure 2 illustrates the transformer-based model for fusing the 2D-pose and spatio-temporal embeddings. In previous sections, we have obtained the 2D-pose output feature vector  $f\text{-pose}_t^i = (J_0, \dots, J_{N_j}, P^h, D_0^h, \dots, D_L^h, D_0^l, \dots, D_K^l)_t^i$  and the spatio-temporal output feature vector  $f\text{-st}_t^i$ . The proposed architecture includes one LSTM layer between the 2D-pose output feature vector and the POSEition embedding input of the encoder for two main reasons. Firstly, the POSEition embedding dimensionality should be consistent with the main encoder input size of spatio-temporal embedding. Secondly, the LSTM layer brings additional temporal information from the extracted 2D-pose features. Therefore, the LSTM module takes  $f\text{-pose}_t^i$  as input and provides the output feature vector size  $N_f=2,304$  for the POSEition embedding. While the primary input of the encoder module part of the transformer-based fusion model is spatio-temporal features, 2D-pose feature embeddings become the POSEition embedding input. The transformer encoder consists of the multi-head self-attention (MSA) with normalization and MLP blocks. The MLP contains two layers with a GELU activation  $GELU(x) = x\varphi(x)$ , where  $\varphi(x)$  is the




Figure 7. Illustration of one selected frame based on the sliding window approach.

cumulative distribution function of the standard Gaussian distribution. Layer normalization (LN) is applied prior to each block and residuals are concatenated after each block. The output of the transformer encoder is fed into the MLP head for computing the class probabilities per camera view.

The proposed method uses a temporal sliding-window approach during training phase to obtain features for each video frame. In this problem, sliding window can be defined as a video segment and the size of the sliding window is equal to the number of frames in that video segment. After obtaining the features for each frame, method combines and utilizes them during the training. As stated in Figure 7, assuming the  $C_n^k$  denotes a video segment with a segment size  $k$  that consists of frame-based class probabilities of  $n^{th}$  frame of a video after applying the sliding window technique. Each frame has  $k$  different class probability values from the sliding window technique, where the  $f_n$  refers to the selected  $n^{th}$  frame within a video segment size  $k$ . In the proposed architecture, video segment size is set  $k=64$ . Considering the selected frame  $f_n$  the label with the maximum number of frames in a given video segment is regarded as the label for the entire video segment for the batch sampling. However, there is a problem with using multi-label video segments during training, making it difficult for the network to classify them as single labels. Typically, researchers deal with this by either using the most commonly observed frame-level label as the segment-level label during training or by discarding these segments altogether. However, authors in [2] show the usefulness of the density-guided label smoothing technique for improving the network’s overall performance by re-using information located within boundary segments. Based on this approach, we train our network with the cross-entropy loss for classification. Including the density-guided label smoothing, the complete loss function is given as

$$\mathcal{L}(x) = - \sum_k p_t^i(k) q''(k|x), \quad (2)$$

where  $(i, t) \in B$  defines a unique video segment in a training-batch  $B$ . We apply the generalized form of the softmax function with the temperature parameter  $\beta$ , enabling the control over smooth vs. sharp transitions of the class label change. After training, we supply test video segments to our network, extracting segment class probabilities,  $p_t^i(y)$ , which are further subject to subsequent post-processing, explained in the next section.

### 3.4. Post-Processing

Multi-camera setup requires the combination of individual stream probabilities into global frame-level probabilities. This means that the class probabilities for a frame from all camera views should be fused to extract the class probabilities of the multi-camera scene. To achieve this, we average all the segment-level class probabilities associated with a video segment that contains the frame under consideration, which is represented by

$$P_{\text{frame}}(y) = \frac{\sum_{c=1}^c P_{\text{frame}}^{\text{seg}}(y)}{c}. \quad (3)$$

where  $P_{\text{frame}}(y)$  is the output frame-level class probabilities within a associated video segment,  $P_{\text{frame}}^{\text{seg}}(y)$  is the segment-level class probabilities for that selected frame under consideration, and  $c$  denotes the total number of cameras within the multi-camera setup, which is set to  $c=3$  in this problem. As a result, the likelihood value assigned to a combination of a class label and a selected scene frame includes information about the likelihood of other class probabilities for that frame, providing 1-D signal output for specified frame including all class probabilities.

Raw predictions of temporally localized and classified actions can be obtained by detecting consistent peaks in the frame-level scene probabilities for each class and scene separately. To identify consistent peaks in obtained 1-D signal, we first use median temporal filtering to remove noise while maintaining sharp edges that are helpful for accurate action localization. To achieve this, we determine the fastest positive and negative changes that precede and succeed at the highest peak to determine the start and end times of the detected driver behaviour. However, some actions may appear similar, resulting in high probabilities for more than one action being assigned to the same frame in individual scenes. Since the target dataset specifies only one action for a given time instance in all scenes, overlapping predictions can increase the false-positive rate and decrease the overall precision of the methodology. To overcome this, we first calculate the intersection-over-union (IoU) for all possible pairs of predictions within the same scene. Then, among the complete set of predictions with an IoU overlap greater than a pre-defined threshold ( $o_{max}$ ), they only keep the prediction with the highest peak in the final output.

## 4. Experiments

### 4.1. Track3 Dataset

The 2023 NVIDIA AI City Challenge [28] has released naturalistic driving dataset [29] consisting of 210 video clips, of around 10-min. length each with a resolution of  $1920 \times 1080$  pixels. Videos recorded of the driver are from various angles using three synchronized cameras mounted inside a car. The dataset features 35 drivers, some wearing appearance blocks such as sunglasses and hats. Each driver performs 16 different tasks such as phone calls, eating, reaching back, in random order, within each video clip. The dataset is split into three parts, A1, A2, and B. The A1 dataset includes labels indicating the start time, end time, and the class of distracting behavior performed, while the A2 dataset has no labels. The main goal of the challenge is to localize and classify the distracted behavior activities, performed by the driver during a given time frame on the A2 test set. The B dataset is reserved for later testing, to determine the top-performers of the challenge.

### 4.2. Implementation Details

For training, we use four GTX-1080Ti GPUs and set the batch size to 80. This batch size allows incorporating 5 training samples of each action class. We adopt the Adam optimizer with initial learning rate  $5 \times 10^{-6}$  and weight decay  $5 \times 10^{-4}$ . The proposed architecture uses the cross-entropy loss with the density-guided label smoothing as the loss function. The temperature parameter,  $\beta$ , is set to  $\beta=5$ . The size of the resulting output feature vector of the 2D-pose feature extraction step is 460. One layer LSTM takes this and provides a 2,304 dimensional vector as output into the POSEition embedding. The resulting output spatio-temporal feature vector has a size of 2,304, which is compatible with the LSTM output layer. We apply a temporal median filter with a filter size of 351 frames for temporal localization and post-processing step. We set the minimum height and width to 0.1 and 200 frames, respectively, to detect peaks, based on empirical observations.

### 4.3. Evaluation metrics

The evaluation metric of the 2023 NVIDIA AICITY challenge Track3 is the overlap score. Given a ground-truth class activity  $g$  with start time  $gs$  and end time  $ge$ , and a predicted class activity  $p$  with start time  $ps$  and end time  $pe$ , with the added condition that the start time  $ps$  and the end time  $pe$  are in the range  $[gs-10s, gs+10s]$  and  $[ge-10s, ge+10s]$ , respectively. To find the overlap between  $g$  and  $p$ , the overlap score is defined as the ratio between the time intersection and the time union of the two activities, i.e.,

$$os(p, g) = \frac{\max(\min(pe, ge) - \max(gs, ps), 0)}{\max(ge, pe) - \min(gs, ps)}, \quad (4)$$

where the final score is the average overlap score among all matched and unmatched activities. In addition to the overlap score, we include the  $F_1$  score, precision, and recall metrics to evaluate more metrics in the experimental results, based on the changes performed in different experiments. The  $F_1$  score is specified by

$$F_1 = 2 \times \frac{P_r \cdot R_c}{P_r + R_c}. \quad (5)$$

The precision  $P_r$  is calculated as the number of true positives  $N_{TP}$  divided by the total number of true positives  $N_{TP}$  and false positives  $N_{FP}$ , which is given by:

$$P_r = \frac{N_{TP}}{N_{TP} + N_{FP}}. \quad (6)$$

The recall  $R_c$  is calculated as the number of true positives  $N_{TP}$  divided by the total number of true positives  $N_{TP}$  and false negatives  $N_{FN}$ , which is given by:

$$R_c = \frac{N_{TP}}{N_{TP} + N_{FN}}. \quad (7)$$

### 4.4. Experimental Results

We have obtained an  $os$  score of 0.5079 in the 2023 NVIDIA Challenge Track3. Note that we have created our local evaluation system during the competition to verify the model’s effectiveness and extend our experiments accordingly. The local evaluation system scores have been almost identical to the public leaderboard scores. In the public leaderboard, we have achieved an  $os$  score of 0.4849 but with the extended experiments and ablation studies, the proposed method achieved an  $os$  score of 0.5079. We present some of our results in Tab. 2, which demonstrates the effectiveness of the proposed approach in evaluating different

Table 2. Experimental results for the proposed solution with two different settings.

Method	$os$ score	$F_1$ score	Precision	Recall
Solution with 2D-Pose skeleton	0.4929	0.6359	0.6591	0.5708
Solution 2D-Pose skeleton&motion	0.5079	0.6452	0.6789	0.5783

Table 3. Experimental results of different video action recognition models for the “dash camera view” dataset.

Method	$os$ score	$F_1$ score	Precision	Recall
SlowFast [11]	0.3459	0.4705	0.4965	0.4472
X3D [10]	0.3136	0.4367	0.4689	0.4379
Fine-tuned X3D	0.3203	0.4399	0.4796	0.4361

Table 4. Experimental results of ablation studies.

Method	<i>os</i> score	$F_1$ score	Precision	Recall
Baseline spatio-temp.	0.3703	0.5126	0.6120	0.4410
+ concat feat. vect. 2D-pose skel.	0.4274	0.5987	0.6364	0.5652
+ concat feat. vect. 2D-pose skel.&mot.	0.4420	0.6330	0.6912	0.5838
Solution w. 2D-pose skel.	0.4322	0.5859	0.6397	0.5403
Solution w. 2D-pose skel.&mot.	0.4493	0.6381	0.6896	0.5846

metrics. This table provides two different 2D-pose features, while one of them includes only skeleton-based points, resulting in a 212-dimensional 2D-pose feature vector. The other feature vector consists of motion vectors additionally based on the selected skeleton points, resulting in a 460-dimensional 2D-pose feature vector. The proposed solution improves the performance by using the 2D-pose features, including both skeleton-based and motion vector-based key points from the selected skeleton points.

#### 4.5. Ablation Studies

We conduct most of the ablation experiments for the “dash camera view” training dataset to understand the importance of each architecture component. This dataset can provide sufficient information for most of the actions. Throughout the ablation study, the following questions are addressed:

1. *Does another video action recognition model work better compared to the SlowFast network?* We have employed the X3D [10] video action recognition backbone network to compare to the SlowFast results. Both X3D and SlowFast feature extractors are models pre-trained on the Kinetics-400 dataset. Tab. 3 illustrates that the X3D model does not improve the results. The 2D-pose features have not been integrated in this comparison.

2. *Does fine-tuning on the video action recognition model improve performance?* We have examined fine-tuning of the X3D feature extractor model on the “dash camera view” dataset by using transfer learning. However, it does not improve the score. A possible reason for this can be data shortage as shown in Tab. 3. Note that we have prepared the naturalistic driver dataset in the format of the Kinetics-400 dataset for this experiment.

3. *Does standard vector concatenation of 2D-pose and spatio-temporal features improve performance?* We have started conducting the experiments with standard feature vector concatenation of the extracted 2D-pose features, based on the coordinates and spatio-temporal features without having the transformer-based architecture. In these

experimental settings, we have taken the mean of each 64 chunks from the 2D-pose feature vector and obtained one feature vector per frame. This improves the contribution compared to using only spatio-temporal features, as shown in Tab. 4.

4. *Which 2D-pose features are better?* We have applied two different types of 2D-pose features. The first type contains skeleton-joint coordinates  $f\text{-pose}_t^i = (J_0, \dots, J_{N_j})_t^i$ , resulting in a 212-dimensional vector. The second type includes motion vectors based on the selected skeleton coordinates  $f\text{-pose}_t^i = (J_0, \dots, J_{N_j}, P^h, D_0^h, \dots, D_L^h, D_0^l, \dots, D_K^l)_t^i$ , resulting in a 460-dimensional vector. Tab. 4 indicates the performance improvement of the motion-vector model.

## 5. Conclusion

In this paper, we have proposed a transformer-based solution for the fusion of 2D-pose and spatio-temporal features to classify and localize the distracted driver behaviors in the time domain. First, we have created both 2D-pose features and spatio-temporal features similar in dimension with the help of a LSTM module. Next, 2D-pose embeddings and spatio-temporal embeddings are fused within the encoder, including the MHA to obtain the frame-level class probabilities for the different camera views. Finally, in the post-processing step, all frame-level probabilities per camera are combined and the final prediction for each distracted driver class is obtained. The proposed solution has been evaluated on the A2 test set of the 2023 NVIDIA AI City Challenge’s Track3, yielding a 0.5079 *os* score.

## Acknowledgements

This work is supported by the European ITEA project SMART on intelligent traffic flow systems.

## References

- [1] National Highway Traffic Safety Administration. *Traffic safety facts research note: Distracted driving 2019 (dot hs 813 111)*. National Highway Traffic Safety Administration., 2021. 1
- [2] Tunc Alkanat, Erkut Akdag, Egor Bondarev, et al. Density-guided label smoothing for temporal localization of driving actions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3174–3182, 2022. 2, 6
- [3] Carlos Caetano, Jessica Sena de Souza, Jefersson dos Santos, and William Schwartz. Skelemotion: A new representation of skeleton joint sequences based on motion information for 3d action recognition, 07 2019. 2
- [4] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6299–6308, 2017. 2



- [5] Vasileios Choutas, Philippe Weinzaepfel, Jérôme Revaud, and Cordelia Schmid. Potion: Pose motion representation for action recognition. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7024–7033, 2018. [2](#)
- [6] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [2](#)
- [7] Haodong Duan, Yue Zhao, Kai Chen, Dahua Lin, and Bo Dai. Revisiting skeleton-based action recognition. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2959–2968, 2022. [2](#), [3](#)
- [8] Hesham M Eraqi, Yehya Abouelnaga, Mohamed H Saad, and Mohamed N Moustafa. Driver distraction identification with an ensemble of convolutional neural networks. *Journal of Advanced Transportation*, 2019, 2019. [2](#)
- [9] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6824–6835, 2021. [2](#)
- [10] Christoph Feichtenhofer. X3d: Expanding architectures for efficient video recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 203–213, 2020. [2](#), [7](#), [8](#)
- [11] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6202–6211, 2019. [2](#), [7](#)
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. [2](#)
- [13] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9:1735–80, 12 1997. [2](#)
- [14] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. [2](#)
- [15] Hamid Reza Vaezi Joze, Amirreza Shaban, Michael L Iuzzolino, and Kazuhito Koishida. Mmtm: Multimodal transfer module for cnn fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13289–13299, 2020. [2](#)
- [16] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014. [2](#)
- [17] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. *arXiv preprint arXiv:1705.06950*, 2017. [5](#)
- [18] Okan Kopuklu, Jiapeng Zheng, Hang Xu, and Gerhard Rigoll. Driver anomaly detection: A dataset and contrastive learning approach. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 91–100, 2021. [2](#)
- [19] Chao Li, Qiaoyong Zhong, Di Xie, and Shiliang Pu. Co-occurrence feature learning from skeleton data for action recognition and detection with hierarchical aggregation. *arXiv preprint arXiv:1804.06055*, 2018. [2](#)
- [20] Weiwei Li, Rong Du, and Shudong Chen. Semantic-structural graph convolutional networks for whole-body human pose estimation. *Information*, 13:109, 02 2022. [4](#)
- [21] Ji Lin, Chuang Gan, and Song Han. Tsm: Temporal shift module for efficient video understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 7083–7093, 2019. [2](#)
- [22] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, Lubomir D. Bourdev, Ross B. Girshick, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. *CoRR*, abs/1405.0312, 2014. [3](#)
- [23] Hong Liu, Juanhui Tu, and Mengyuan Liu. Two-stream 3d convolutional neural network for skeleton-based action recognition. *arXiv preprint arXiv:1705.08106*, 2017. [2](#), [3](#)
- [24] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al. Swin transformer v2: Scaling up capacity and resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12009–12019, 2022. [2](#)
- [25] Diogo C Luvizon, David Picard, and Hedi Tabia. 2d/3d pose estimation and action recognition using multitask deep learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5137–5146, 2018. [2](#)
- [26] Eric Marchand, Hideaki Uchiyama, and Fabien Spindler. Pose estimation for augmented reality: A hands-on survey. *IEEE Transactions on Visualization and Computer Graphics*, 22(12):2633–2651, 2016. [4](#)
- [27] Manuel Martin, Alina Roitberg, Monica Haurilet, Matthias Horne, Simon Reiß, Michael Voit, and Rainer Stiefelwagen. Drive&act: A multi-modal dataset for fine-grained driver behavior recognition in autonomous vehicles. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2801–2810, 2019. [2](#)
- [28] Milind Naphade, Shuo Wang, David C. Anastasiu, Zheng Tang, Ming-Ching Chang, Yue Yao, Liang Zheng, Mohammed Shaiqur Rahman, Meenakshi S. Arya, Anuj Sharma, Qi Feng, Vitaly Ablavsky, Stan Sclaroff, Pranamesh Chakraborty, Sanjita Prajapati, Alice Li, Shangru Li, Krishna Kunadharaju, Shenxin Jiang, and Rama Chellappa. The 7th AI City Challenge. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2023. [7](#)
- [29] Mohammed Shaiqur Rahman, Jiyang Wang, Senem Velipasalar Gursay, David Anastasiu, Shuo Wang, and Anuj Sharma. Synthetic Distracted Driving (SynDD2) dataset for analyzing distracted behaviors and various gaze zones of a driver, 2022. *arXiv:2204.08096*. [7](#)
- [30] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal net-

- works. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(06):1137–1149, jun 2017. 3
- [31] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *Advances in Neural Information Processing Systems*, 1, 06 2014. 2
- [32] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2
- [33] Ke Sun, Bin Xiao, Dong Liu, and Jingdong Wang. Deep high-resolution representation learning for human pose estimation. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5686–5696, 2019. 3
- [34] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott E. Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. *CoRR*, abs/1409.4842, 2014. 2
- [35] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. 2
- [36] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015. 2
- [37] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6450–6459, 2018. 2
- [38] Sijie Yan, Yuanjun Xiong, and Dahua Lin. Spatial temporal graph convolutional networks for skeleton-based action recognition. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence and Thirtieth Innovative Applications of Artificial Intelligence Conference and Eighth AAAI Symposium on Educational Advances in Artificial Intelligence*, AAAI’18/IAAI’18/EAAI’18. AAAI Press, 2018. 3
- [39] Ceyuan Yang, Yinghao Xu, Jianping Shi, Bo Dai, and Bolei Zhou. Temporal pyramid network for action recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 591–600, 2020. 2
- [40] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, and George Toderici. Beyond short snippets: Deep networks for video classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. 2
- [41] Yue Zhao, Yuanjun Xiong, and Dahua Lin. Trajectory convolution for action recognition. *Advances in neural information processing systems*, 31, 2018. 2
- [42] Dingyuan Zhu, Ziwei Zhang, Peng Cui, and Wenwu Zhu. Robust graph convolutional networks against adversarial attacks. In *Proceedings of the 25th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1399–1407, 2019. 3