

# Coupling Intent and Action for Pedestrian Crossing Behavior Prediction

Yu Yao, Ella Atkins, Matthew Johnson-Roberson, Ram Vasudevan and Xiaoxiao Du

University of Michigan

{brianyao, ematkins, mattjr, ramv, xiaodu}@umich.edu

## Abstract

Accurate prediction of pedestrian crossing behaviors by autonomous vehicles can significantly improve traffic safety. Existing approaches often model pedestrian behaviors using trajectories or poses but do not offer a deeper semantic interpretation of a person’s actions or how actions influence a pedestrian’s intention to cross in the future. In this work, we follow the neuroscience and psychological literature to define pedestrian crossing behavior as a combination of an unobserved inner will (a probabilistic representation of binary intent of crossing vs. not crossing) and a set of multi-class actions (*e.g.*, walking, standing, *etc.*). Intent generates actions, and the future actions in turn reflect the intent. We present a novel multi-task network that predicts future pedestrian actions and uses predicted future action as a prior to detect the present intent and action of the pedestrian. We also designed an attention relation network to incorporate external environmental contexts thus further improve intent and action detection performance. We evaluated our approach on two naturalistic driving datasets, PIE and JAAD, and extensive experiments show significantly improved and more explainable results for both intent detection and action prediction over state-of-the-art approaches. *Our code is available at: [https://github.com/umautobots/pedestrian\\_intent\\_action\\_detection](https://github.com/umautobots/pedestrian_intent_action_detection)*

## 1 Introduction

Pedestrians are among the most exposed and vulnerable road users in traffic accidents and fatalities in urban areas [Cantillo *et al.*, 2015; Yao *et al.*, 2019b; Yao *et al.*, 2020b]. In particular, pedestrians have been found more likely to be involved in crashes by vehicles and at greater risks during street crossing [Kim and Ulfarsson, 2019]. Therefore, for the safe navigation and deployment of intelligent driving systems in populated environments, it is essential to accurately understand and predict pedestrian crossing behaviors.

Many existing work that study pedestrian intention [Rehder and Kloeden, 2015; Bai *et al.*, 2015; Karasev *et al.*, 2016] referred to the term “intention” as goals in pedestrian

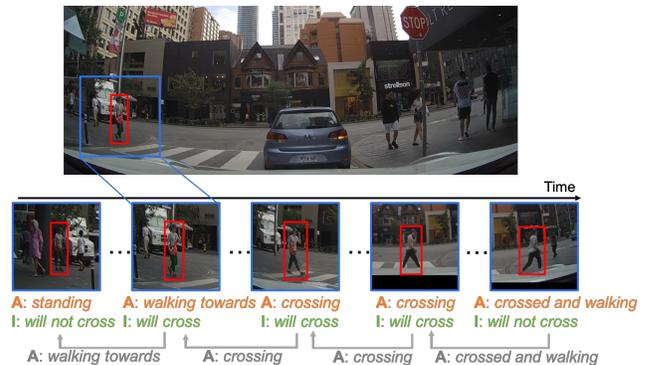


Figure 1: A sample video sequence illustrating our future action (A) guided intent (I) detection. A person goes through the sequence of first standing far from the intersection with no intent of crossing, and then starts going towards the crosswalk with crossing intent. The person then keeps crossing until reaching the state of crossed and walking. Successfully predicting future action of going towards and crossing can help detect the “will cross” intent.

trajectories and evaluated using trajectory or path location errors. However, this definition only models intent as motion but did not offer a deeper semantic interpretation of both the pedestrian’s physical and mental activities. In this work, we follow neuroscience and psychology literature [Haggard *et al.*, 2002; Haggard, 2008; Schröder *et al.*, 2014] and define “intent” as the inner consciousness or “will” of a person’s decision to act. We hypothesize that a human’s implicit intent affects the explicit present and future actions, and future actions in turn reflect the intent. Thus, we propose to leverage future action prediction as priors to perform online intent and action detection and prediction. We developed a multi-task network to encode pedestrian visual and location features, in addition to features from the future action predictor, to output probabilities of binary crossing intent (will cross or not) and multi-class semantic actions (*e.g.*, “waiting” or “going towards the crosswalk”). We also developed an attentive relation module to extract object-level features (different from previously commonly used semantic maps) to incorporate environmental context in predicting pedestrian crossing behaviors. The main contributions of our work are two-fold:

- We redefine the pedestrian behavior as a combination of a binary-class crossing intent and a multi-class semantic

action and designed a novel multi-task model to detection and predict both jointly. We show that our method achieves superior (improved by  $\sim 5 - 25\%$ , based on multiple metrics) and real-time ( $< 6ms$ ) intent and action prediction performance over the state-of-the-art on two in-the-wild benchmark driving datasets.

- Inspired by human visual learning, we designed an attentive relation network (ARN) to extract environmental information based on pre-detected traffic objects and their state. We show that this object-level relation feature is more effective for pedestrian intent detection task, compared to existing methods using semantic maps. Our ARN also offers interpretable attention results on the relations between the traffic objects in the urban scene and a pedestrian’s crossing intent and actions.

## 2 Related Work

**Intent detection.** Intention detection algorithms offer an intuitive understanding to human behaviors and assist many application areas, including user purchase and consumption [Wang *et al.*, 2020; Ding *et al.*, 2018], Electroencephalography (EEG)-based motion monitoring [Fang *et al.*, 2020], and human-robot interaction [Li and Zhang, 2017; Chen *et al.*, 2019], among others. To model pedestrian behaviors in driving applications, prior work typically model intent as motion and use poses and/or trajectories to predict future goals or locations [Yao *et al.*, 2019a; Fang and López, 2018; Yao *et al.*, 2020a; Bouhsain *et al.*, 2020]. The convolutional neural network (CNN) have been used to detect pedestrian intents in videos [Saleh *et al.*, 2019]. The state-of-the-art for pedestrian crossing intent detection and most relevant to our work include  $PIE_{int}$  [Rasouli *et al.*, 2019], stacked fusion GRU (SF-GRU) [Rasouli *et al.*, 2020a] and multi-modal hybrid pedestrian action prediction (MMH-PAP) [Rasouli *et al.*, 2020b].  $PIE_{int}$  uses a Convolutional LSTM network to encode past visual features and then concatenates encoded features and bounding boxes to predict pedestrian intent at the final frame. SF-GRU used pre-computed pose feature and two CNNs to extract pedestrian appearance and context feature separately. MMH-PAP used pre-computed semantic maps to extract environmental feature for intent prediction. However, none of these work explored the relation between action and intent, nor can they predict actions in addition to intent.

**Action recognition.** Recurrent neural networks (RNNs) and 3-D CNNs have been used for action recognition in TV shows, indoor videos and sports videos *etc.* [Simonyan and Zisserman, 2014; Tran *et al.*, 2015; Wang *et al.*, 2016; Carreira and Zisserman, 2017; Feichtenhofer *et al.*, 2019]. These methods typically make one prediction per video and are mainly used for off-line applications, such as video tagging and query. For real-time applications, online action detectors have been developed to classify video frames with only past observations [Gao *et al.*, 2017; Shou *et al.*, 2018; Gao *et al.*, 2019]. However, none of these work explored pedestrian crossing actions, nor do they explore the causes of the action (the intent). The temporal recurrent network (TRN) [Xu *et al.*, 2019] is most relevant to our work as it provided evidence that predicting future action can better recognize the

action that is currently occurring. Different from TRN which focused on video action, our work focuses on pedestrian behaviors and predicts future actions to assist intent detection. While TRN leveraged future action to detect current action, our method views future action as the consequence of current intent and uses predicted action as a prior for intent detection.

## 3 Approach

Given a naturalistic driving dataset containing 2-D RGB traffic video clips, let  $\{O_1, O_2, \dots, O_t\}$  denote  $t$  time steps of past observations (*i.e.*, image frames) in a sequence. The goal of our work is to detect the probability of a pedestrian’s crossing intent  $i_t \in [0, 1]$  and the pedestrian’s action  $a_t \in \mathcal{A}$  at time  $t$ , and also predict future semantic actions  $[a_{t+1}, \dots, a_{t+\delta}] \in \mathcal{A}$ , where  $\mathcal{A}$  denotes the collection of possible actions related to crossing, as described in Section 4. The proposed architecture is illustrated in Figure 2. The main part of our network contains a multi-task, encoder-decoder-based intent detection and action predictor, which outputs the detection results of the pedestrian crossing intent at the current time step and probabilistic forecast of future pedestrian behaviors (Section 3.1). Additionally, an Attentive Relation Network (ARN) was developed to explore the surrounding environment for traffic objects that are in relation to the pedestrians (Section 3.2).

### 3.1 Intent Detection with Future Action Prediction

The multi-task intent and action prediction model consists of two parts: an intent-action encoder and a future action predictor. A multi-class cross-entropy-based loss function is used to optimize the network.

**Intent-action encoder.** Given 2-D RGB traffic video clips, visual features were first extracted from image patches around each target pedestrian using a pre-trained deep convolutional neural network (CNN) following [Rasouli *et al.*, 2019], as illustrated in the upper-left of Fig. 2. An image patch is a square region defined by an enlarged pedestrian bounding box (“bbox”) that includes a pedestrian as well as some context information such as ground, curb, crosswalk, etc. Meanwhile, the bounding box coordinates of the image patch were passed through a fully-connected (FC) network to obtain location-scale embedding features. Then, the intent-action encoder (ENC) cells took the concatenated feature vectors as input and recurrently encoded the features over time. To ensure the neural network learns features that represent both crossing intent and more complicated semantic actions, a multi-layer perceptron (MLP)-based intent classifier and an action classifier were designed to use the encoder hidden state to predict the present intent and action simultaneously.

**Action predictor.** Prior work such as the temporal recurrent neural network (TRN) [Xu *et al.*, 2019] have shown that using future action predictions can improve the performance of classifying an action in the present. Inspired by this, our action predictor module was designed to utilize features from future action predictions and pass this information to the next ENCCell. In the action predictor module, the hidden state at each iteration was classified by an MLP network to predict the action at each future time step. Then, the average of all decoder hidden vectors was computed and used as the “future”



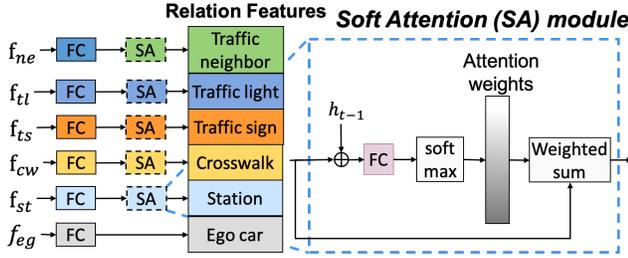


Figure 3: Our attentive relation network (ARN) with soft attention (SA). FC is fully-connected layer. Different colors correspond to different traffic objects.  $f_{(\cdot)}$  represents feature vectors for all objects, e.g.,  $f_{ne} = [f_{ne}^1, \dots, f_{ne}^{N_{ne}}]$ , where  $N_{ne}$  is the total number of observed traffic participants.  $h_{t-1}$  is the previous hidden state.

puted as described above and an attentive relation network was designed to extract the relational features of all objects for a target pedestrian, as illustrated in Fig. 3. Each feature vector was first embedded to the same size by a corresponding FC network (e.g., neighbor feature vector is embedded by  $FC_{ne}$ ). For object types that may have multiple instances (e.g., there can be more than one cars in the scene), a soft attention network (SAN) [Chen *et al.*, 2017] was adopted to fuse the features. The SAN predicts a weighing factor for each instance, which represents the importance of that instance to the desired task, and then use a weighted sum to merge all instances. The ego car feature is not processed by a SAN, since current work assumes there is only one ego car at any given time. Then, the extract features of each traffic object type were concatenated to form the final relation feature  $f_{rel} = [f_{ne}, f_{tl}, f_{ts}, f_{cw}, f_{st}, f_{eg}]$ . Note that all elements of a feature vector were set to zeros if the corresponding object type was not observed in a scene. Finally, the relation features extracted through the ARN were concatenated with the visual, bounding box, and future action features as described in Section 3.1 to be passed together into the ENC network for intent and action encoding.

Compared to using visual features (e.g., semantic maps) to represent environmental information as used in prior work [Rasouli *et al.*, 2020a; Rasouli *et al.*, 2020b], our ARN directly uses pre-detected traffic object features, which is far more straightforward and informative. This is also inspired by human visual learning, as human drivers make decision based on observing objects rather than image pixels. The ARN is a very simple and fast network, since no deep CNNs were needed for visual feature extraction.

## 4 Experiments

We evaluate the effectiveness of our method for intent detection and action prediction on two publicly available naturalistic traffic video benchmarks, Pedestrian Intent Estimation (PIE) [Rasouli *et al.*, 2019] and Joint Attention in Autonomous Driving (JAAD) [Kotseruba *et al.*, 2016]. The PIE dataset was collected with an on-board camera covering six hours of driving footage. There are 1,842 pedestrians (880/243/719 for training/validation/test) with 2-D bounding boxes annotated at 30Hz with behavioral tags. Ego-vehicle velocity readings were provided using gyroscope readings

from the camera. The JAAD dataset contains 346 videos with 686 pedestrians (188/32/126) captured from dashboard cameras, annotated at 30Hz with crossing intent annotations. Note that the original pedestrian action annotation in JAAD and PIE contains only binary categories, “standing” and “walking”. To make the annotations more informative for intent detection, we leveraged the intent and crossing annotation to automatically augment the existing two into seven categories: standing, waiting, going towards (the crossing point), crossing, crossed and standing, crossed and walking, and other walking.

### 4.1 Experimental Settings

**Evaluation metrics.** Pedestrian crossing intent detection is traditionally defined as a binary classification problem (cross/not cross) and we include previous benchmark evaluation metrics including accuracy, F1 score, precision, and the receiver operating characteristic (ROC) area under curve (AUC) [Rasouli *et al.*, 2020b]. We introduce another metric in this work – the difference between the average scores  $s$  of all positive samples  $P$  and all negative samples  $N$ ,  $\Delta_s = \frac{1}{|P|} \sum_{i \in P} s_i - \frac{1}{|N|} \sum_{i \in N} s_i$ . This  $\Delta_s$  metric reflects the margin between the detected positive and negative labels. The greater the  $\Delta_s$ , the larger the margin, thus, the more easily (better) the model can separate the crossing (positive) and not crossing (negative) classes. Note that accuracy, F1 score and precision can be biased towards the majority class if the test datasets were imbalanced (which is the case with both datasets used). Besides, these three metrics were computed at a hard threshold at 0.5 (score > 0.5 is classified as positive, otherwise negative). On the other hand, both AUC and  $\Delta_s$  are less sensitive to the data imbalance thus can more accurately reflect the performance comparisons between methods.

**Implementation details.** The proposed neural network model was implemented using PyTorch. We used gated recurrent units (GRU) with hidden size 128 for both ENCcell and DECcell, and a single-layer MLP as the classifier for intent detection, action detection and action prediction. The input image patch was prepared following [Rasouli *et al.*, 2019] to realize fair comparison, resulting in  $224 \times 224$  image size. An ImageNet-pre-trained VGG16 network acted as the backbone CNN to extract features from image patches [Rasouli *et al.*, 2019]. During training, the pedestrian trajectories were segmented to training samples with constant segment length  $T$  in order to make batch training feasible. We applied a weighted sampler during training to realize balanced mini-batch. All models were trained with sample length  $T = 30$ , prediction horizon  $\delta = 5$ , learning rate  $1e - 5$ , and batch size 128 on a NVIDIA Tesla V100 GPU. During inference, the trained model was run on one test sequence at a time and the per-frame detection results were collected.

**Ablations.** We compared multiple variations of our method: The “*I-only*” model is a basic model without action detection and prediction modules, which is similar to  $PIE_{int}$  [Rasouli *et al.*, 2019] except that we used an average pooling feature extraction and a simpler GRU encoder rather than the ConvLSTM network in  $PIE_{int}$ ; The “*I+A*” model is a multi-task model with both intent and present ac-

tion detection modules and the “ $I+A+F$ ” adds future action prediction; The “ $I+A+F+R$ ” or “*full*” is our full model with the attentive relation network. We also present an “*A-only*” and an “ $A+F$ ” ablations to show that future action prediction improves present action detection.

## 4.2 Performance Evaluation and Analysis

Two data sampling settings exist in previous work and we report results on both for direct and fair comparison: 1) [Rasouli *et al.*, 2019] used all original data in PIE dataset, and 2) [Rasouli *et al.*, 2020a; Rasouli *et al.*, 2020b] sampled sequences from each track between 1-2 seconds up to the time of the crossing events. The second setting offers more future action changes than the first, since the first setting contains sequences long before the crossing event and actions tend not to change much until closer to the event of crossing.

### Quantitative Results on the Original Data Setting

The top three rows of Table 1 shows the intent detection results of the baseline methods. *Naive-1* achieves 0.82 accuracy, 0.90 F1 score and 0.82 precision, even higher than the previously state-of-the-art  $PIE_{int}$  method due to the imbalanced test set of PIE dataset. This observation verifies our previous discussion on the inadequacy of using only accuracy, F1 score and precision as evaluation metrics. The  $\Delta_s$  of naive methods was 0.00 because it cannot distinguish crossing and not crossing classes, while  $\Delta_s$  score of  $PIE_{int}$  is slightly higher at 0.06.

The bottom rows of Table 1 below the middle line show ablation results on the variations of our proposed method. The *I-only* baseline, with action detection and prediction removed, achieves similar results with  $PIE_{int}$  in terms of AUC and  $\Delta_s$ , indicating our choice of an average pooling and simple GRU network with less parameters can perform similarly to the ConvLSTM encoder in  $PIE_{int}$ . The  $I+A$  model significantly increased AUC and  $\Delta_s$ , indicating an improved capability in distinguishing between crossing and not crossing intent. However, its lower accuracy and F1 scores show that it missed some crossing intents. One explanation is that more negative (not crossing) samples previously misclassified by the *I-only* model are now correctly classified when the action type is known in the  $I+A$  model. For example, a person standing close to the curb but not facing the road can be regarded as just “standing” there instead of “waiting”. The  $I+A$  model recognizes such differences and will correctly classify this example as “not crossing”, while the *I-only* model may lose that information. The  $I+A+F$  model further improves the AUC and  $\Delta_s$  as well the accuracy and F1 scores from the  $I+A$  ablation. This verifies our assumption that future action is a more straightforward indicator to intent, *e.g.*, a high confidence on “waiting” or “crossing” action indicates higher probability that a pedestrian has crossing intent. The higher  $\Delta_s$  results show that adding future action stream helps the model to better distinguish different crossing intents. Our *full* model,  $I+A+F+R$ , significantly improves the accuracy, AUC and  $\Delta_s$ , indicating the effectiveness of adding the environmental information from ARN in detecting intent.

| Method                                     | intent         |               |             |                |                       | action         |
|--|----------------|---------------|-------------|----------------|-----------------------|----------------|
|  | Acc $\uparrow$ | F1 $\uparrow$ | Prec        | AUC $\uparrow$ | $\Delta_s$ $\uparrow$ | mAP $\uparrow$ |
| Naive - 1                                  | <b>0.82</b>    | <b>0.90</b>   | 0.82        | 0.50           | 0.00                  | N/A            |
| $PIE_{int}$ [Rasouli <i>et al.</i> , 2019] | <u>0.79</u>    | 0.87          | 0.86        | 0.73           | 0.06                  | N/A            |
| Ours ( <i>I-only</i> )                     | <b>0.82</b>    | <b>0.90</b>   | 0.86        | 0.72           | 0.07                  | N/A            |
| Ours ( <i>A-only</i> )                     | N/A            | N/A           | N/A         | N/A            | N/A                   | 0.18           |
| Ours ( $I+A$ )                             | 0.74           | 0.83          | 0.90        | 0.76           | 0.11                  | 0.21           |
| Ours ( $A+F$ )                             | N/A            | N/A           | N/A         | N/A            | N/A                   | 0.20           |
| Ours ( $I+A+F$ )                           | 0.79           | 0.87          | 0.90        | 0.78           | 0.16                  | 0.23           |
| Ours ( $I+A+F+R$ ( <i>full</i> ))          | <b>0.82</b>    | <u>0.88</u>   | <b>0.94</b> | <b>0.83</b>    | <b>0.41</b>           | <b>0.24</b>    |

Table 1: Intent and action detection results on the PIE dataset using the original data setting. **Bold** is best and underline is second best.

| Method                                  | PIE         |             |             |             | JAAD        |             |             |             |
|---|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|
|   | Acc         | F1          | AUC         | Prec        | Acc         | F1          | AUC         | Prec        |
| Naive-1 or 0                            | 0.81        | <b>0.90</b> | 0.50        | 0.81        | 0.79        | <b>0.88</b> | 0.50        | <b>0.79</b> |
| I3D [Carreira and Zisserman, 2017]      | 0.63        | 0.42        | 0.58        | 0.37        | 0.79        | 0.49        | 0.71        | 0.42        |
| SF-GRU [Rasouli <i>et al.</i> , 2020a]  | 0.87        | 0.78        | 0.85        | 0.74        | 0.83        | 0.59        | 0.79        | 0.50        |
| MMH-PAP [Rasouli <i>et al.</i> , 2020b] | <b>0.89</b> | 0.81        | <b>0.88</b> | 0.77        | 0.84        | 0.62        | 0.80        | 0.54        |
| Ours ( <i>full</i> )                    | 0.84        | <b>0.90</b> | <b>0.88</b> | <b>0.96</b> | <b>0.87</b> | 0.70        | <b>0.92</b> | 0.66        |

Table 2: Results on PIE and JAAD datasets using “event-to-crossing” setting. Detection precision (Prec) is presented and  $\Delta_s$  was removed to compare with baselines using their metrics.

### Quantitative Results on the “Event of Crossing” Setting

Table 2 shows the results of our method on PIE and JAAD following training and evaluation method in [Rasouli *et al.*, 2020a; Rasouli *et al.*, 2020b] considering only sequences between 1-2s to the event of crossing. We again observed that the *Naive-1* (for PIE) or *-0* (for JAAD) methods achieved higher accuracy, F1 score and precision on both datasets due to the imbalanced testing samples, since PIE and JAAD are dominated by positive and negative samples respectively. AUC is a more reliable metric in this case. Our method outperforms the state-of-the-art SF-GRU and MMH-PAP methods and achieved the highest 0.88 AUC on PIE and 0.92 AUC on JAAD datasets, which indicated that predicting the change of action through our future-action-conditioned module indeed helps improve the detection of the crossing intent.

### Ablation Study on ARN

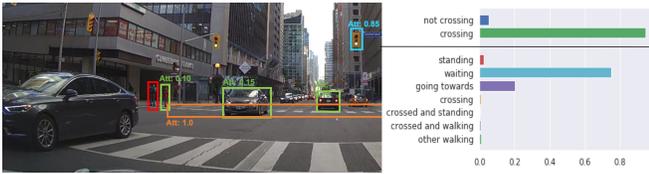
Table 3 studies the impact of each different traffic object type in the ARN. As shown, adding traffic lights and/or crosswalks significantly improves the results, while the rest of the objects by themselves do not seem to help much. This result indicated that the ARN network learned that the two major factors of determining pedestrian crossing intent include traffic lights and crosswalks, which fits human intuition.

### Qualitative Examples

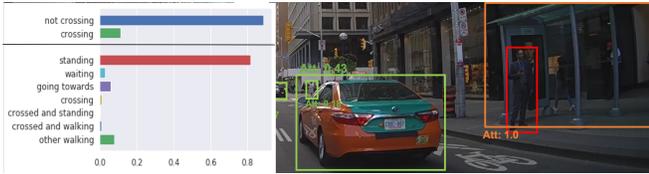
Fig. 4 shows three examples of running our *full* model on the PIE dataset. As shown in Fig 4(a), our  $I+A+F$  model has higher confidence in “crossing” prediction than  $I+A$  and *I-only*, demonstrating again the effectiveness of predicting future actions as the prior of current intent detection. Our *full* model predicts highest score for “crossing” intent and the “waiting” action by paying higher attention to the front crosswalk, traffic light, and the traffic cars and the nearby pedestrian that has same intent and action. The pedestrian traffic light was not detected, thus, no attention was assigned. Fig. 4(b) shows an example where a pedestrian is waiting for

| Traffic neighbor | Cross walk | Traffic light | Traffic sign | Station | Ego car | Acc         | F1          | AUC         | $\Delta_s$  |
|------------------|------------|---------------|--------------|---------|---------|-------------|-------------|-------------|-------------|
| -                | -          | -             | -            | -       | -       | 0.79        | <b>0.87</b> | 0.78        | 0.16        |
| ✓                | -          | -             | -            | -       | -       | 0.73        | 0.83        | 0.70        | 0.16        |
| -                | ✓          | -             | -            | -       | -       | 0.77        | 0.84        | <b>0.83</b> | 0.31        |
| -                | -          | ✓             | -            | -       | -       | <b>0.81</b> | <b>0.88</b> | <u>0.81</u> | <u>0.39</u> |
| -                | -          | -             | ✓            | -       | -       | 0.78        | 0.86        | 0.73        | 0.12        |
| -                | -          | -             | -            | ✓       | -       | 0.71        | 0.80        | 0.77        | 0.12        |
| -                | -          | -             | -            | -       | ✓       | 0.79        | <b>0.87</b> | 0.73        | 0.10        |
| ✓                | ✓          | ✓             | ✓            | ✓       | ✓       | <b>0.82</b> | <b>0.88</b> | <b>0.83</b> | <b>0.41</b> |

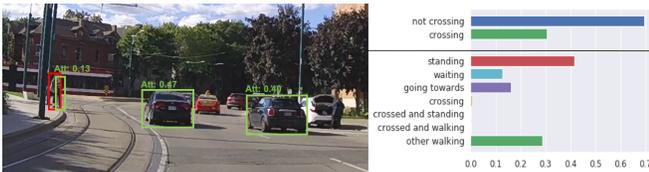
Table 3: Ablation study on ARN on the PIE dataset. The first row shows our I+A+F model, while the following are models with each single traffic object type, including ego motion. The bottom row shows our full model with all traffic object types in PIE dataset.



(a) A pedestrian waiting to cross at an intersection. The predicted crossing scores are: 0.65 (*I-only*), 0.70 (*I+A*), 0.75 (*I+A+F*) and **0.95** (*full*).



(b) A pedestrian waiting for bus at a bus stop without crossing intent. The predicted non-crossing scores are: 0.31 (*I-only*), 0.36 (*I+A*), 0.35 (*I+A+F*) and **0.89** (*full*).



(c) A pedestrian waiting to cross a street without traffic light/sign or crosswalk. The predicted crossing scores are: 0.65 (*I-only*), 0.68 (*I+A*), **0.80** (*I+A+F*) and 0.31 (*full*).

Figure 4: Qualitative examples of our full model on PIE dataset. The target pedestrians are shown by red boxes. Up to top-5 attentive traffic participants, traffic lights and infrastructures from ARN are shown in green, cyan and orange boxes together with the intermediate attention probabilities in our model. The predicted crossing-intent scores and action scores are presented in bar charts.

bus at a bus stop. All the models without ARN predicted low score for non-crossing intent, while our *full* model improves the result significantly since it captures the bus stop (orange box) and also the taxi (green box) in front of it.

Fig 4(c) shows a “failure” case of our *full* method. In this scene, two pedestrians are waiting to cross the street without crosswalk, traffic light or traffic sign. Our *I+A+F* model predicted the highest score for crossing intent, since it focused

on the context near the target pedestrian and the temporal actions in the future. The *full* model, on the contrary, predicted low scores, since it did not find the infrastructures related to a crossing event (e.g., crosswalk, traffic light, traffic sign etc.). Moreover, the *full* model paid higher attention to the two cars than the nearby pedestrian who was also waiting to cross, making it more difficult to recognize the crossing intent of the target pedestrian. This challenging case indicated that while the current model with ARN is successful at incorporating relation features, it sometimes lacks the ability to recognize when to rely on relation feature and when to switch rely on target pedestrian’s feature, which is an interesting future work area to further explore. Our current ARN also shares the same drawback with prior supervised methods in that it cannot extract features from environmental objects that were unseen during training. Semi-supervised or unsupervised methods may be developed to solve this issue.

### 4.3 Discussion on Computation Time

We report the computation time to show the efficiency of our models. All experiments were conducted on a single NVIDIA Tesla V100 GPU. The mean and standard deviation ( $\mu \pm \sigma$ ) of the per-frame inference time of our models are  $0.5 \pm 0.1ms$  (*I-only*),  $0.5 \pm 0.1ms$  (*I+A*),  $2.2 \pm 0.1ms$  (*I+A+F*),  $6.0 \pm 1.2ms$  (*full*), respectively. Our *full* model, being slightly slower than its ablations, is still fast enough for real-time applications. The larger standard deviation was due to the fact that the number of traffic objects varied largely in different scenes. As a comparison, we ran the original implementation of previous *PIE<sub>int</sub>* on the same machine and observed a higher  $9 \pm 0.5ms$  inference time.

## 5 Conclusion

In this work, we proposed a multi-task action and intent prediction network for pedestrian crossing behavior given naturalistic driving videos. We showed through extensive experiments that our method outperforms the state-of-the-art by a large margin and can jointly and recurrently predict present and future actions and intent for every pedestrian in every frame, in real-time. Through our novel future action predictor module, we verified our hypothesis that there is a relation between intent and action, that intent affects actions and that future action is a useful prior to accurately predict intent. We also studied the relations among various traffic objects and showed that our ARN module assigns attention to nearby objects that affect crossing intent and improve the performance.

## Acknowledgments

This work was supported by a grant from Ford Motor Company via the Ford-UM Alliance under award N028603. This material is based upon work supported by the Federal Highway Administration under contract number 693JJ319000009. Any options, findings, and conclusions or recommendations expressed in the this publication are those of the author(s) and do not necessarily reflect the views of the Federal Highway Administration.

## References

- [Bai *et al.*, 2015] Haoyu Bai, Shaojun Cai, Nan Ye, David Hsu, and Wee Sun Lee. Intention-aware online pomdp planning for autonomous driving in a crowd. In *ICRA*, 2015.
- [Bouhsain *et al.*, 2020] Smail Ait Bouhsain, Saeed Saadatnejad, and Alexandre Alahi. Pedestrian intention prediction: A multi-task perspective. *arXiv preprint arXiv:2010.10270*, 2020.
- [Cantillo *et al.*, 2015] Victor Cantillo, Julian Arellana, and Manuel Rolong. Modelling pedestrian crossing behaviour in urban roads: A latent variable approach. *Transportation research part F: traffic psychology and behaviour*, 32:56–67, 2015.
- [Carreira and Zisserman, 2017] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, 2017.
- [Chen *et al.*, 2017] Long Chen, Hanwang Zhang, Jun Xiao, Liqiang Nie, Jian Shao, Wei Liu, and Tat-Seng Chua. Sca-cnn: Spatial and channel-wise attention in convolutional networks for image captioning. In *CVPR*, 2017.
- [Chen *et al.*, 2019] Cen Chen, Xiaolu Zhang, Sheng Ju, Chilin Fu, Caizhi Tang, Jun Zhou, and Xiaolong Li. Antprophet: an intention mining system behind alipay’s intelligent customer service bot. In *IJCAI*, 2019.
- [Ding *et al.*, 2018] Xiao Ding, Bibo Cai, Ting Liu, and Qiankun Shi. Domain adaptation via tree kernel based maximum mean discrepancy for user consumption intention identification. In *IJCAI*, 2018.
- [Fang and López, 2018] Zhijie Fang and Antonio M López. Is the pedestrian going to cross? answering by 2d pose estimation. In *IV*, 2018.
- [Fang *et al.*, 2020] Zhijie Fang, Weiqun Wang, Shixin Ren, Jiaying Wang, Weiguo Shi, Xu Liang, Chen-Chen Fan, and Zengguang Hou. Learning regional attention convolutional neural network for motion intention recognition based on eeg data. In *IJCAI*, 2020.
- [Feichtenhofer *et al.*, 2019] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *ICCV*, 2019.
- [Gao *et al.*, 2017] Jiyang Gao, Zhenheng Yang, and Ram Nevatia. Red: Reinforced encoder-decoder networks for action anticipation. *BMVC*, 2017.
- [Gao *et al.*, 2019] Mingfei Gao, Mingze Xu, Larry S Davis, Richard Socher, and Caiming Xiong. StartNet: Online detection of action start in untrimmed videos. In *ICCV*, 2019.
- [Haggard *et al.*, 2002] Patrick Haggard, Sam Clark, and Jeri Kalogeras. Voluntary action and conscious awareness. *Nature neuroscience*, 5(4):382–385, 2002.
- [Haggard, 2008] Patrick Haggard. Human volition: towards a neuroscience of will. *Nature Reviews Neuroscience*, 9(12):934–946, 2008.
- [Karasev *et al.*, 2016] Vasilij Karasev, Alper Ayvaci, Bernd Heisele, and Stefano Soatto. Intent-aware long-term prediction of pedestrian motion. In *ICRA*, 2016.
- [Kim and Ulfarsson, 2019] Sungyop Kim and Gudmundur F Ulfarsson. Traffic safety in an aging society: Analysis of older pedestrian crashes. *J. Transp. Safety & Security*, 11(3), 2019.
- [Kotseruba *et al.*, 2016] Iuliia Kotseruba, Amir Rasouli, and John K Tsotsos. Joint attention in autonomous driving (JAAD). *arXiv preprint arXiv:1609.04741*, 2016.
- [Li and Zhang, 2017] Songpo Li and Xiaoli Zhang. Implicit intention communication in human–robot interaction through visual behavior studies. *IEEE Trans. Human-Mach. Syst.*, 47(4):437–448, 2017.
- [Rasouli *et al.*, 2019] Amir Rasouli, Iuliia Kotseruba, Toni Kunic, and John K Tsotsos. PIE: A Large-Scale Dataset and Models for Pedestrian Intention Estimation and Trajectory Prediction. In *ICCV*, 2019.
- [Rasouli *et al.*, 2020a] Amir Rasouli, Iuliia Kotseruba, and John K Tsotsos. Pedestrian action anticipation using contextual feature fusion in stacked rnns. In *BMVC*, 2020.
- [Rasouli *et al.*, 2020b] Amir Rasouli, Tiffany Yau, Mohsen Rohani, and Jun Luo. Multi-modal hybrid architecture for pedestrian action prediction. *arXiv preprint arXiv:2012.00514*, 2020.
- [Rehder and Kloeden, 2015] Eike Rehder and Horst Kloeden. Goal-directed pedestrian prediction. In *ICCVW*, 2015.
- [Saleh *et al.*, 2019] Khaled Saleh, Mohammed Hossny, and Saeid Nahavandi. Real-time intent prediction of pedestrians for autonomous ground vehicles via spatio-temporal densenet. In *ICRA*, 2019.
- [Schröder *et al.*, 2014] Tobias Schröder, Terrence C Stewart, and Paul Thagard. Intention, emotion, and action: A neural theory based on semantic pointers. *Cognitive science*, 38(5), 2014.
- [Shou *et al.*, 2018] Zheng Shou, Junting Pan, Jonathan Chan, Kazuyuki Miyazawa, Hassan Mansour, Anthony Vetro, Xavier Giro-i Nieto, and Shih-Fu Chang. Online detection of action start in untrimmed, streaming videos. In *ECCV*, 2018.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *NeurIPS*, 2014.
- [Tran *et al.*, 2015] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *ICCV*, 2015.
- [Wang *et al.*, 2016] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *ECCV*, 2016.
- [Wang *et al.*, 2020] Shoujin Wang, Liang Hu, Yan Wang, Quan Z Sheng, Mehmet Orgun, and Longbing Cao. Intention2basket: A neural intention-driven approach for dynamic next-basket planning. In *IJCAI*, 2020.
- [Xu *et al.*, 2019] Mingze Xu, Mingfei Gao, Yi-Ting Chen, Larry S Davis, and David J Crandall. Temporal recurrent networks for online action detection. In *ICCV*, 2019.
- [Yao *et al.*, 2019a] Yu Yao, Mingze Xu, Chiho Choi, David J Crandall, Ella M Atkins, and Behzad Dariush. Ego-centric vision-based future vehicle localization for intelligent driving assistance systems. In *ICRA*, 2019.
- [Yao *et al.*, 2019b] Yu Yao, Mingze Xu, Yuchen Wang, David J Crandall, and Ella M Atkins. Unsupervised traffic accident detection in first-person videos. In *IROS*, 2019.
- [Yao *et al.*, 2020a] Yu Yao, Ella Atkins, Matthew Johnson-Roberson, Ram Vasudevan, and Xiaoxiao Du. Bitrap: Bi-directional pedestrian trajectory prediction with multi-modal goal estimation. *arXiv preprint arXiv:2007.14558*, 2020.
- [Yao *et al.*, 2020b] Yu Yao, Xizi Wang, Mingze Xu, Zelin Pu, Ella Atkins, and David Crandall. When, where, and what? a new dataset for anomaly detection in driving videos. *arXiv preprint arXiv:2004.03044*, 2020.