

Pedestrian Action Anticipation using Contextual Feature Fusion in Stacked RNNs

Amir Rasouli

<http://www.eecs.yorku.ca/~aras/>

Iuliia Kotseruba

http://www.eecs.yorku.ca/~yulia_k/

John K. Tsotsos

<http://www.eecs.yorku.ca/~tsotsos/Tsotsos/Home.html>

Dept. of Electrical Engineering and
Computer Science

York University

Toronto, Canada

Abstract

One of the major challenges for autonomous vehicles in urban environments is to understand and predict other road users' actions, in particular, pedestrians at the point of crossing. The common approach to solving this problem is to use the motion history of the agents to predict their future trajectories. However, pedestrians exhibit highly variable actions most of which cannot be understood without visual observation of the pedestrians themselves and their surroundings. To this end, we propose a solution for the problem of pedestrian action anticipation at the point of crossing. Our approach uses a novel stacked RNN architecture in which information collected from various sources, both scene dynamics and visual features, is gradually fused into the network at different levels of processing. We show, via extensive empirical evaluations, that the proposed algorithm achieves a higher prediction accuracy compared to alternative recurrent network architectures. We conduct experiments to investigate the impact of the length of observation, time to event and types of features on the performance of the proposed method. Finally, we demonstrate how different data fusion strategies impact prediction accuracy.

1 Introduction

Autonomous driving systems suitable for urban environments require the ability to comprehend and anticipate the actions of other road users. In this context, pedestrians are of particular importance being the most vulnerable road users, especially when crossing the road. Anticipating pedestrian crossing action helps the driving systems to select the correct course of action to avoid any potential collisions and disruption of traffic flow.

Today, the dominant approaches to solving the problem of pedestrian action prediction are trajectory-based. These algorithms rely on the motion patterns of pedestrians and predict their trajectories at some time in the future [1, 2, 3]. Although pedestrian dynamics are important, they are not sufficient for making sense of pedestrian behavior and predicting their actions as they are often subject to error. For example, a pedestrian intending to cross the street could be standing at the intersection (with no motion history), walking alongside the road or abruptly changing their walking pattern prior to crossing [4] (see Figure 1). In



Figure 1: Examples of pedestrians prior to making crossing decision. Green and red colors indicate whether the pedestrian will or will not cross.

addition, pedestrians exhibit highly variable motion patterns which can be influenced by various environmental factors such as signals, the ego-vehicle motion, road structure, etc. All of these factors add to the complexity of predicting pedestrian actions [53]. Thus a statistical inference on pedestrian trajectories alone may not sufficient for predicting their actions. To remedy this problem, some algorithms, in addition to pedestrians’ trajectories, use information such as head orientation [20], social interactions [0, 15] or destination locations [0, 34] to predict pedestrians’ forthcoming actions. However, to achieve a more robust prediction in a complex traffic scene, there is a need for a more general approach that exploits various sources of contextual information.

To this end, we propose a novel algorithm for pedestrian action anticipation in video sequences recorded with an on-board camera inside a moving vehicle. Given that the main point of interaction between autonomous vehicles and pedestrians is at the time of crossing, here we particularly focus on the pedestrian crossing anticipation, i.e. we determine whether an observed pedestrian will cross in front of the vehicle. For this purpose, instead of focusing on a subset of contextual elements, we take a more holistic approach and combine information from multiple sources, including the appearance, pose and dynamics of pedestrians, their surrounding context, and the ego-vehicle speed. We introduce a novel stacked recurrent neural network architecture in which data from different modalities are gradually fused at different levels of processing. We conduct experiments to evaluate the performance of the proposed algorithm and discuss how different data arrangement strategies affects its performance.

2 Related Works

Action recognition and prediction are widely studied topics in the field of computer vision [0, 8, 00, 22, 27, 29, 41, 42]. Among commonly used methods for modeling sequential data are recurrent neural networks (RNNs). In particular, their variants, Gated Recurrent Units (GRUs) [00] and Long Short-Term Memory (LSTM) units [16] have a wide range of applications for activity recognition [40, 42], action prediction [0, 27] and video captioning [3, 19]. In their simplest form, RNNs are used in a single-layer format [19, 22, 27, 42] for temporal reasoning with the objective of classifying activities or predicting future behavior. RNNs are also used in multi-stream architectures [0, 24, 40] in which, for example, data of different modalities is processed in separate streams and combined at the end for inference [4]. Some approaches increase the depth of RNNs in space by stacking them on top of

each other [12, 17, 26, 28, 30], so that the input to the RNN at each level consists of the hidden states of the RNN in the previous layer. A variation of this is stacked RNNs that are organized in a hierarchy [9, 13, 23, 30, 45] where the length of information flow in each successive layer is reduced to minimize the computation cost. In both stacked and hierarchical architectures, the raw inputs, e.g. image features, enter the network at the bottom layer, and the representations of the first layer are propagated through the network.

RNN-based architectures show particularly promising results in action prediction applications such as future scene generation [6, 24, 29], activity class prediction [12, 18, 27], event anticipation [9, 40, 46], and trajectory estimation [0, 8, 22]. The latter is the most commonly used in the field of pedestrian action prediction in natural scenes. These techniques often use encoder-decoder architectures in which one RNN unit is responsible for encoding observations, often by tracking the locations of pedestrians, and generating representations which are used by a decoder, often another RNN unit, to infer future predictions.

Motion histories are often augmented with additional contextual information to achieve more robust predictions. For example, in [10, 15, 35] the social interactions between agents are taken into account for forecasting their future trajectories. These approaches, however, are applied to top-down view scenes (recorded with stationary camera) where complex social interactions can be observed easily. In the context of driving, action prediction is more challenging due to the ego-vehicle motion which may also impact the behavior of other agents. In [9] the authors use a two-stream encoder-decoder LSTM model which simultaneously predicts pedestrian trajectory and future speed of the ego-vehicle. The authors combine the estimated vehicle speed with the representations of pedestrians' trajectories to predict their future motion. Although scene information is used for speed estimation, no visual information involving the pedestrians, aside from their trajectories, is used for inference. In some works, in addition to scene dynamics, authors use cues such as pedestrian's head orientation and proximity to the curb [20] or pedestrian's potential destination locations [0, 34], however, these approaches still predominantly rely on pedestrian's trajectories which might not be sufficient for the reasons mentioned earlier.

Some algorithms address pedestrian action prediction as a classification problem and use various contextual elements such as road structure [6, 38], head orientation and signals [32] to predict pedestrian intention of crossing. These models, however, do not take into account all information relevant for reasoning, e.g. scene dynamics [32], do not propose a mechanism for visual processing of the scenes [38] or are evaluated in a very limited context [9].

Contributions. This paper contributes the followings: 1) We propose a novel stacked recurrent network architecture with multilevel feature fusion for predicting pedestrian crossing action. The proposed algorithm benefits from a combination of various visual and motion features for prediction¹. 2) We perform extensive experiments to evaluate the performance of the proposed algorithm against alternative RNN architectures. 3) We examine the effects of time to event and observation duration on the accuracy of crossing prediction. 4) We investigate the influence of various sources of contextual information on the performance of the proposed algorithm. 5) In the end, we show how changing the order of feature inputs during training and inference according to their level of complexity affects prediction.

¹The code is available at <https://github.com/aras62/SF-GRU>

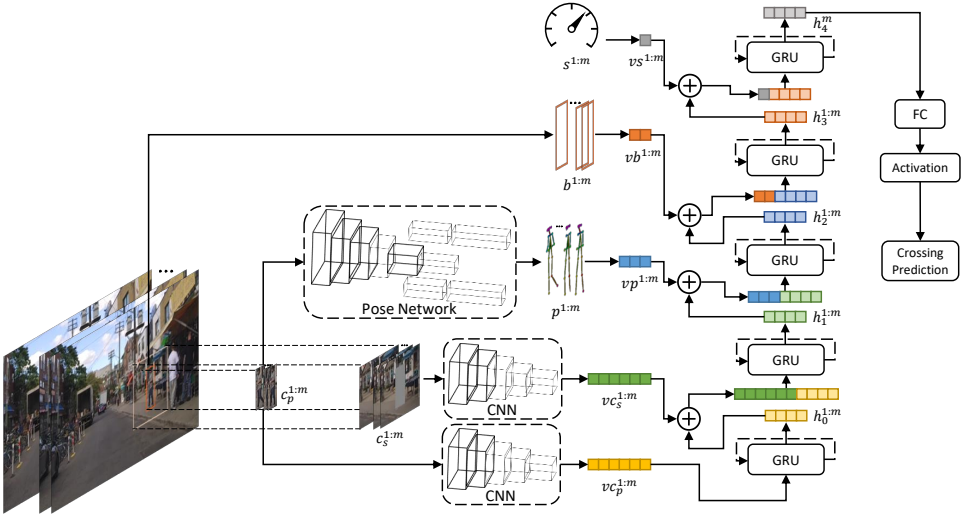


Figure 2: The architecture of the proposed algorithm *SF-GRU* comprised of five GRUs each of which processes a concatenation of features of different modalities and the hidden states of the GRU in the previous level. The information is fused into the network gradually according to the complexity of the features. Each feature input consists of m sequential observations. From bottom to top layers features are fused as follows: pedestrian appearance $c_p^{1:m}$, surrounding context $c_s^{1:m}$, poses $p^{1:m}$, bounding boxes $b^{1:m}$ and ego-vehicle speed $s^{1:m}$. \oplus refers to concatenation operation.

3 Approach

We define pedestrian crossing prediction as a binary classification problem in which the objective is to determine whether a pedestrian i will cross the street given the observed context up to some time m . The prediction relies on five sources of information including the local context $\{C_{p_i}, C_{s_i}\}$, where $C_{p_i} = \{c_{p_i}^1, \dots, c_{p_i}^m\}$ and $C_{s_i} = \{c_{s_i}^1, \dots, c_{s_i}^m\}$ refer to visual features of the pedestrian and their surroundings respectively, pedestrian pose $P_i = \{p_i^1, \dots, p_i^m\}$, 2D bounding box locations $B_i = \{b_i^1, \dots, b_i^m\}$, where b_i is a two-point coordinate $[(x1_i, y1_i)(x2_i, y2_i)]$ corresponding to the top-left and bottom-corner of the bounding box around the pedestrian, and the speed of the ego-vehicle $S = \{s^1, \dots, s^m\}$.

3.1 Architecture

Recurrent neural networks (RNNs) are extensions of feedforward networks. RNNs have recurrent hidden states allowing them to learn temporal dependencies in sequence data. This inherent temporal depth has been shown to greatly benefit tasks, such as pedestrian trajectory prediction, that apply single-layer RNNs to point coordinates in a space. In addition to temporal depth, spatial depth of RNNs can also be increased by stacking multiple layers of RNN units on top of one another. This approach is an effective way of improving sequential data modeling in complex tasks [13], in particular, video sequence analysis [17, 18] in which the network models dependencies between visual features of consecutive video frames.

Given the multimodal nature of pedestrian action anticipation which relies on both dynamics and visual scene information, we employ a hybrid approach. We use a stacked RNN architecture similar to [28] in which we gradually fuse the features at each level according to their complexity. In other words, we input the visual features of the scene that can benefit more from spatial depth of the network at the bottom layers and the dynamics features, e.g. trajectories and speed, at the higher levels of the network (see Figure 2). Below, we describe the procedures to generate each data type we use in the proposed model.

Local context. At each time step of the observation, for each pedestrian, we use their appearance and surroundings. The former is captured using images cropped to the size of 2D bounding boxes around the pedestrian in the frame. For the surroundings, we extract a region around the pedestrian by scaling up the 2D bounding box coordinates, and squarifying the dimensions so the width of the scaled bounding box matches its height. This gives us a wider viewing angle of the scene around the pedestrian which may include street, other pedestrians, signals or traffic. In the surround crop, we suppress the pedestrian appearance by setting the pixel values in the original bounding box coordinate to neutral gray. Both appearance and surround crops are processed using a convolutional neural network (CNN) which produces two feature vectors $vc_p^{1:m}$ and $vc_s^{1:m}$.

Pose. A pose network is used to generate 18 body joints coordinates, each corresponding to a point in 2D space, for each pedestrian. The joint coordinates are normalized and concatenated into a 36D feature vector $vp^{1:m}$.

2D bounding box. We transform the bounding box coordinates into relative displacement from the initial position forming a feature vector $vb^{1:m}$. This can be seen as the velocity of the pedestrian at every time step.

Speed. Speed is a vector of the ego-vehicle speed recordings for each time step $vs^{1:m}$ in km/h .

Multimodal feature fusion. For the joint modeling of our sequence data, we use gated recurrent units (GRUs) [14] which are simpler compared to LSTMs and, in our case, achieve similar performance. Recalling the equation of GRU, the j^{th} level of the stack is given by,

$$\begin{aligned} r_j^t &= \sigma(W_j^{xr}x_j^t + W_j^{hr}h_j^{t-1}), \\ z_j^t &= \sigma(W_j^{xz}x_j^t + W_j^{hz}h_j^{t-1}), \\ \tilde{h}_j^t &= \tanh(W_j^{xh}x_j^t + W_j^{hh}(r_j^t \odot h_j^{t-1})), \\ h_j^t &= (1 - z_j^t) \odot h_j^{t-1} + z_j^t \odot \tilde{h}_j^t, \end{aligned} \quad (1)$$

where $\sigma(\cdot)$ is the sigmoid function, r^t and z^t are reset and update gates, and matrices $W^{\cdot\cdot}$ are weights between two units. For $j = 0$ (the bottom level of the stack), $x_0^t = vc_p^t$ and for $j > 0$, $x_j^t = h_{j-1}^t + vy^t[j-1]$ where $y^t = \{vc_s^t, vp^t, vb^t, vs^t\}$. The final prediction is achieved by a linear transformation of h_n^t where n is the number of levels (in our case 5) in the proposed stacked architecture. In the training phase we use the *binary cross-entropy* loss function.

3.2 Implementation

In our architecture, we use GRUs [14] with 256 hidden units. For local context, we crop the pedestrian samples C_p using the 2D bounding box annotations, resize them so the larger dimension is equal to 224 and pad them with zeros to preserve the aspect ratio. For surround context, C_s , we use 1.5x (set empirically) scaled version of the 2D bounding boxes. The parts of the cropped images that include pedestrians of interest are suppressed by neutral gray with

RGB of (128, 128, 128). We resize these images to 224×224 . The local context images are processed using VGG16 [69] (without fully connected (*fc*) layers) pretrained on ImageNet [66] followed by an average global pooling generating a feature vector of size 512 per crop. For pedestrian poses, we use [7] which is pretrained on the COCO dataset [25]. The network generates 18-joint pose per pedestrian sample.

Training The model is trained using ADAM [40] optimizer with a learning rate of 5×10^{-6} for 60 epochs with batch size of 32 and *L2* regularization of 0.0001. The context and pose features are precomputed. In addition, we augment the data at training time by horizontally flipping the images and sub-sampling the over-represented class to equalize the number of crossing and non-crossing samples.

4 Experiments

Dataset. There are not many datasets suitable for the purpose of pedestrian crossing prediction. One such dataset is JAAD [52] which contains videos of pedestrians prior or during crossing. Unfortunately, the number of samples in this dataset is small, no vehicle information is available and sequences are short snippets which are not suitable for long-term predictions.

For the purpose of this work, we used our newly collected pedestrian intention estimation (PIE) dataset². The dataset comprises 1842 pedestrian tracks captured using an on-board monocular camera while driving in urban environments with various street structures and crowd densities. The samples represent people who are close to the curbs or are at intersections and may or may not have the intention of crossing, e.g. waiting for a bus. Overall, the ratio of non-crossing to crossing events is 2.5 to 1. All video sequences are collected during daylight under clear weather conditions. The videos are continuous allowing us to observe the pedestrians from the moment they appear in the scene until they go out of the field of view of the camera.

The dataset contains bounding box annotations for pedestrians as well as vehicle sensor data such as speed and heading angle per frame. For each pedestrian sample we identified an event point. For those who cross in front of the ego-vehicle, the event is the moment they start crossing. For other samples, the events are set at the time when the pedestrians go out of the field of view of the camera. We randomly split the data into train-test sets with ratio of 60-40 respectively.

Metrics. As in [40], we report all the evaluation results using the following metrics: *Accuracy*, *F1 score*, *precision* and *recall*. We also use *Area Under Curve (AUC)* metric which, in the case of binary event anticipation, reflects the balanced accuracy of the algorithms.

4.1 Evaluations

Predicting crossing event. We evaluate the performance of our proposed algorithm, stacked with multilevel fusion GRU (*SF-GRU*), against the following models:

Static. This model is inspired by [52] and has two VGG16 branches (without *fc* layers and with a global pooling layer) pretrained on ImageNet. One network processes the local context corresponding to the pedestrian crop c_p^m and the other processes the surroundings c_s^m at the last frame of the observation. The outputs of both networks are combined and fed into a

²http://data.nvision2.eecs.yorku.ca/PIE_dataset/



Figure 3: Examples of the predictions produced by the proposed algorithm *SF-GRU* and top competing methods, namely *GRU*, *M-GRU*, *H-GRU*, and *S-GRU*. In the examples, *GT* stands for ground truth and *green* and *red* colors indicate whether the pedestrian will cross in front of the ego-vehicle or not respectively. The instances where the color of the algorithm labels matches the GT means that their predictions are correct.

fc layer for the final prediction.

GRU. A single-layer GRU [10] trained and tested only on pedestrians’ appearances C_p and their surroundings C_s . We also use this model with all sources of information which are concatenated and fed into the network at the same time.

Multi-stream GRU (M-GRU). Following the approach in [9], this architecture processes different types of features separately using different GRUs, and feeds the concatenation of the last hidden states of all GRUs into a dense layer for prediction.

Hierarchical GRU (H-GRU). This model has a hierarchical structure similar to [10]. H-GRU processes each feature type using a separate GRU, concatenates the hidden states of all units and then feeds them into another GRU whose last hidden state is used for prediction.

Stacked GRU (S-GRU). This is a five-level stacked GRU architecture as described in [28] which receives the feature inputs at the bottom layer. The inputs to the subsequent GRUs in the higher levels are the hidden states of the GRUs in the previous layers.

All evaluations are done on observation sequences of 0.5s (15 frames) duration. The samples are selected with 2s time to event (TTE), the minimum time within which pedestrians make crossing decision according to [67].

The results are summarized in Table 1. We can see that using the visual information of the local context, even as a single image in the static method can lead to approximately 60% accuracy which can be improved by 9% by performing temporal reasoning using a GRU.

Using all sources of information, the proposed algorithm *SF-GRU* performs best on all metrics except recall. For this metric single-layer *GRU* performs slightly better (by 1.2%) at the expense of more than 6% drop in precision. In addition, the results show that no performance improvement is achievable by simply adding layers to the network or separating the processing of features with different modalities.

When to predict a crossing event. The prediction of crossing events may vary depending on TTE as the scene dynamics changes, in particular, when the ego-vehicle motion impacts the way people make a crossing decision. Here we examine the prediction ability of

Models	Features	Acc	AUC	FI	Prec	Recall
Static	c_p^m, c_s^m	0.592	0.589	0.419	0.328	0.582
GRU	C_p, C_s	0.681	0.644	0.475	0.407	0.570
GRU	C_p, C_s, P, B, S	0.811	0.812	0.685	0.593	0.812
M-GRU	C_p, C_s, P, B, S	0.804	0.792	0.665	0.585	0.770
H-GRU	C_p, C_s, P, B, S	0.819	0.805	0.685	0.612	0.776
S-GRU	C_p, C_s, P, B, S	0.801	0.770	0.643	0.588	0.709
SF-GRU (ours)	C_p, C_s, P, B, S	0.844	0.829	0.721	0.657	0.800

Table 1: Evaluation results of the algorithms using observation length of 0.5s and time to event (TTE) of 2s. Abbreviations in *features* column are: pedestrian appearance C_p , surround context C_s , pose P , bounding box B , and ego-vehicle speed S . c_p^m and c_s^m stand for appearance and surround context in the last observation frame respectively.




Figure 4: The performance of the algorithms with respect to varying time to event (TTE) points with 0.5s observation length.

the temporal algorithms with respect to TTE. We alter TTE points from 0s to 3s with steps of approximately 0.16s, a total of 19 different points. To maintain the consistency of data across different time frames, we only sample from pedestrian tracks equal to or longer than 3.5s (the maximum TTE time in the experiment + observation length). All other parameters including the observation sequence length remain the same as before.

The proposed algorithm *SF-GRU* performs best for the most part at different TTE points (see Figure 4). At early TTE times where the intention of pedestrians becomes obvious, all algorithms perform similarly well. However, as expected, the performance of the algorithms degrades gradually (some at a faster rate than others) as the observations are moved further away from the time of the event. We can also see that the single-layer *GRU* only performs better than *M-GRU* and *S-GRU* up to 2s TTE after which its performance drops rapidly.

The effect of observation length on prediction. Longer observation time can potentially provide more information but at the same time may add noise. We examine the effect of observation length on the proposed algorithm *SF-GRU* with respect to different TTE points. For the same reason as mentioned in the previous experiment, we only sample from tracks with length equal to or longer than 4.5s (the longest observation length + the largest TTE value). In total, we examine 16 different combinations.

As shown in Figure 5, on most metrics the improvement gain is only on samples very close (0s) or far away (3s) from the event. In these cases, precision can be improved by longer observations at the expense of reducing the recall. In critical decision regions of 1-2s, however, a small gain is achieved by increasing observation from 0.3s to 0.5s after which




Figure 5: The changes in the performance of *SF-GRU* according to varying observation length and time to event (TTE).

<i>Features</i>	<i>Acc</i>	<i>AUC</i>	<i>F1</i>	<i>Prec</i>	<i>Recall</i>
C_p	0.660	0.622	0.448	0.380	0.546
C_{p+s}	0.666	0.650	0.483	0.397	0.618
C_p, C_s	0.692	0.645	0.475	0.417	0.552
C_p, C_s, P	0.745	0.705	0.554	0.498	0.624
C_p, C_s, P, D	0.796	0.765	0.636	0.580	0.703
C_p, C_s, P, B	0.816	0.781	0.661	0.619	0.709
C_p, C_s, P, B, S	0.844	0.829	0.721	0.657	0.800

Table 2: The impact of different sources of information on the performance of *SF-GRU*. The feature types are as follows: C_p pedestrian context (appearance), C_s surround context, C_{p+s} full context, P pose, D displacement (center coordinates), B bounding box, and S speed.

point the performance drops rapidly. This could be due to noise in longer observations caused by accumulation of the changes in the scene dynamics. For instance, within 1.5s observation window, the speed of the vehicle can change significantly which can have a considerable effect on predicting pedestrian crossing behavior.

Feature types and prediction accuracy. We examine the contribution of each feature type on the performance of the proposed algorithm. In addition to the features discussed earlier, we also evaluate two other types of features: displacement D (the center coordinates of the bounding boxes) and full context C_{p+s} which is the pedestrian appearance and surround context in a single frame, not as decoupled features as proposed earlier.

As shown in Table 2, we can see that adding contextual information in addition to pedestrian appearance to the network improves the overall performance by more than 18%. We also see that decoupling appearance and surround context boosts the accuracy by almost 3% owing to precision gain. Another observation is that using bounding box coordinates instead of center coordinates improves the results by 2%. This can be due to the fact that the changes in the scale of the bounding boxes in a sequence can add another layer of information, e.g. the movement of pedestrian or the changes in their distance to the ego-vehicle.

The order of fusion and performance. In this experiment, we investigate how different fusion strategies alter the performance. Since reporting on all possible permutations of different sources of information is prohibitive, we only include a subset of these permutations to show the fluctuations in the overall performance.

A summary of the results is provided in Table 3. Here, it is shown that when more complex features such as local context are infused into higher levels of the network, the performance gets worse. By inputting different feature types in the right order, that is by moving simpler features, such as speed, to the higher levels of the stack, the performance improves by up to 9% on accuracy, 10% on recall and more than 15% on precision. This

Features	Acc	AUC	F1	Prec	Recall
P, S, B, C_p, C_s	0.753	0.737	0.590	0.509	0.703
S, B, C_p, C_s, P	0.784	0.759	0.624	0.557	0.709
B, C_p, C_s, P, S	0.798	0.776	0.647	0.579	0.733
S, C_p, C_s, P, B	0.810	0.785	0.661	0.602	0.733
C_p, B, C_s, S, P	0.813	0.803	0.679	0.619	0.788
C_p, C_s, P, B, S	0.844	0.829	0.721	0.657	0.800

Table 3: Feature fusion strategies and their impact on the performance of the proposed algorithm *SF-GRU*. The feature types are as follows: C_p pedestrian context (appearance), C_s surround context, P pose, B bounding box, and S speed.

can be due to the fact that more complex visual features, which benefit more from deeper spatial analysis, are inputted at the bottom layers of the network while simpler features such as trajectory coordinates are entered at the higher levels.

5 Conclusion

We presented a novel stacked RNN architecture in which different sources of information including pedestrian and the vehicle dynamics, pedestrian appearance and their surroundings are fused gradually at different levels of processing. Using empirical evaluations we showed that our approach performs best compared to alternative RNN architectures. In addition, we demonstrated how different sources of information and data fusion strategies into the network can impact crossing action prediction. We highlighted that the performance is optimal when more complex features are fed to the bottom layers of the network and the simpler ones at the higher levels. Although the proposed architecture was presented in the context of pedestrian crossing prediction, other applications of similar nature e.g. activity recognition may also benefit from using this approach.

Acknowledgements

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), the NSERC Canadian Robotics Network (NCRN), the Air Force Office for Scientific Research (USA), and the Canada Research Chairs Program through grants to JKT.

References

- [1] A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese. Social LSTM: Human trajectory prediction in crowded spaces. In *CVPR*, pages 961–971, 2016.
- [2] H. Bai, S. Cai, N. Ye, D. Hsu, and W. S. Lee. Intention-aware online POMDP planning for autonomous driving in a crowd. In *ICRA*, pages 454–460, 2015.
- [3] Lorenzo Baraldi, Costantino Grana, and Rita Cucchiara. Hierarchical boundary-aware neural encoder for video captioning. In *CVPR*, pages 1657–1666, 2017.

-
- [4] Apratim Bhattacharyya, Mario Fritz, and Bernt Schiele. Long-term on-board prediction of people in traffic scenes under uncertainty. In *CVPR*, pages 4194–4202, 2018.
 - [5] S. Bonnin, T. H. Weisswange, F. Kummert, and J. Schmuedderich. Pedestrian crossing prediction using multiple context-based models. In *Intelligent Transportation Systems (ITSC)*, pages 378–385, 2014.
 - [6] Wonmin Byeon, Qin Wang, Rupesh Kumar Srivastava, and Petros Koumoutsakos. ContextVP: Fully context-aware video prediction. In *ECCV*, pages 781–797, 2018.
 - [7] Zhe Cao, Tomas Simon, Shih-En Wei, and Yaser Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, pages 7291–7299, 2017.
 - [8] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *CVPR*, pages 4724–4733. IEEE, 2017.
 - [9] Fu-Hsiang Chan, Yu-Ting Chen, Yu Xiang, and Min Sun. Anticipating accidents in dashcam videos. In *ACCV*, pages 136–153, 2016.
 - [10] Lei Chen, Jiwen Lu, Zhanjie Song, and Jie Zhou. Part-activated deep reinforcement learning for action prediction. In *ECCV*, pages 421–436, 2018.
 - [11] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. *arXiv:1412.3555*, 2014.
 - [12] Jeffrey Donahue, Lisa Anne Hendricks, Sergio Guadarrama, Marcus Rohrbach, Subhashini Venugopalan, Kate Saenko, and Trevor Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *CVPR*, pages 2625–2634, 2015.
 - [13] Yong Du, Wei Wang, and Liang Wang. Hierarchical recurrent neural network for skeleton based action recognition. In *CVPR*, pages 1110–1118, 2015.
 - [14] Harshala Gammulle, Simon Denman, Sridha Sridharan, and Clinton Fookes. Two stream lstm: A deep fusion framework for human action recognition. In *WACV*, pages 177–186, 2017.
 - [15] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *CVPR*, pages 2255–2264, 2018.
 - [16] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
 - [17] Mohamed Ilyes Lakhall, Albert Clapés, Sergio Escalera, Oswald Lanz, and Andrea Cavallaro. Residual stacked rnns for action recognition. In *ECCV*, pages 534–548, 2018.
 - [18] Ashesh Jain, Avi Singh, Hema S Koppula, Shane Soh, and Ashutosh Saxena. Recurrent neural networks for driver activity anticipation via sensory-fusion architecture. In *ICRA*, pages 3118–3125, 2016.

- [19] Wenhao Jiang, Lin Ma, Yu-Gang Jiang, Wei Liu, and Tong Zhang. Recurrent fusion network for image captioning. In *ECCV*, pages 499–515, 2018.
- [20] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *ICLR*, 2015.
- [21] Julian Francisco Pieter Kooij, Nicolas Schneider, Fabian Flohr, and Dariu M Gavrilă. Context-based pedestrian path prediction. In *ECCV*, pages 618–633, 2014.
- [22] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B. Choy, Philip H. S. Torr, and Manmohan Chandraker. DESIRE: Distant future prediction in dynamic scenes with interacting agents. In *CVPR*, pages 336–345, 2017.
- [23] Wenbo Li, Longyin Wen, Ming-Ching Chang, Ser Nam Lim, and Siwei Lyu. Adaptive rnn tree for large-scale human action recognition. In *ICCV*, pages 1444–1452, 2017.
- [24] Xiaodan Liang, Lisa Lee, Wei Dai, and Eric P. Xing. Dual motion gan for future-flow embedded video prediction. In *ICCV*, pages 1744–1752, 2017.
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755, 2014.
- [26] Jun Liu, Amir Shahroudy, Dong Xu, Alex C Kot, and Gang Wang. Skeleton-based action recognition using spatio-temporal lstm network with trust gates. *PAMI*, 40(12): 3007–3021, 2018.
- [27] Tahmida Mahmud, Mahmudul Hasan, and Amit K. Roy-Chowdhury. Joint prediction of activity labels and starting times in untrimmed videos. In *ICCV*, 2017.
- [28] Joe Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *CVPR*, pages 4694–4702, 2015.
- [29] Marc Oliu, Javier Selva, and Sergio Escalera. Folded recurrent neural networks for future video prediction. In *ECCV*, pages 5773–5782, 2018.
- [30] Pingbo Pan, Zhongwen Xu, Yi Yang, Fei Wu, and Yueting Zhuang. Hierarchical recurrent neural encoder for video representation with application to captioning. In *CVPR*, pages 1029–1038, 2016.
- [31] Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, and Yoshua Bengio. How to construct deep recurrent neural networks. *arXiv:1312.6026*, 2013.
- [32] Amir Rasouli, Iuliia Kotseruba, and John K. Tsotsos. Are they going to cross? A benchmark dataset and baseline for pedestrian crosswalk behavior. In *ICCVW*, pages 206–213, 2017.
- [33] Amir Rasouli, Iuliia Kotseruba, and John K Tsotsos. Agreeing to cross: How drivers and pedestrians communicate. In *Intelligent Vehicles Symposium (IV)*, pages 264–269, 2017.

- [34] Eike Rehder, Florian Wirth, Martin Lauer, and Christoph Stiller. Pedestrian prediction by planning using deep neural networks. In *ICRA*, pages 1–5, 2018.
- [35] Alexandre Robicquet, Amir Sadeghian, Alexandre Alahi, and Silvio Savarese. Learning social etiquette: Human trajectory understanding in crowded scenes. In *ECCV*, pages 549–565, 2016.
- [36] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 115(3):211–252, 2015.
- [37] S. Schmidt and B. Farber. Pedestrians at the kerb—recognising the action intentions of humans. *Transportation Research Part F: Traffic Psychology and Behaviour*, 12(4):300–310, 2009.
- [38] F. Schneemann and P. Heinemann. Context-based detection of pedestrian crossing intention for autonomous driving in urban environments. In *IROS*, pages 2243–2248, 2016.
- [39] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556*, 2014.
- [40] B. Singh, T. K. Marks, M. Jones, O. Tuzel, and M. Shao. A multi-stream bi-directional recurrent neural network for fine-grained action detection. In *CVPR*, pages 1961–1970, 2016.
- [41] Tomoyuki Suzuki, Hirokatsu Kataoka, Yoshimitsu Aoki, and Yutaka Satoh. Anticipating traffic accidents with adaptive loss and large-scale incident db. In *CVPR*, pages 3521–3529, 2018.
- [42] Vivek Veeriah, Naifan Zhuang, and Guo-Jun Qi. Differential recurrent neural networks for action recognition. In *ICCV*, pages 4041–4049, 2015.
- [43] Amnon Shashua Yoav Levine, Or Sharir. Benefits of depth for long-term memory of recurrent networks. In *ICLR*, 2018.
- [44] YoungJoon Yoo, Kimin Yun, Sangdoon Yun, JongHee Hong, Hawook Jeong, and Jin Young Choi. Visual path prediction in complex scenes with crowded moving objects. In *CVPR*, pages 2668–2677, 2016.
- [45] Haonan Yu, Jiang Wang, Zhiheng Huang, Yi Yang, and Wei Xu. Video paragraph captioning using hierarchical recurrent neural networks. In *CVPR*, pages 4584–4593, 2016.
- [46] Kuo-Hao Zeng, Shih-Han Chou, Fu-Hsiang Chan, Juan Carlos Niebles, and Min Sun. Agent-centric risk assessment: Accident anticipation and risky region localization. In *CVPR*, pages 2222–2230, 2017.