

Agent-Centric Risk Assessment: Accident Anticipation and Risky Region Localization

Kuo-Hao Zeng^{*†} Shih-Han Chou^{*} Fu-Hsiang Chan^{*} Juan Carlos Niebles[†] Min Sun^{*} [†]Stanford University *National Tsing Hua University {khzeng, jniebles}@cs.stanford.edu {happy810705, corgi1205}@gmail.com sunmin@ee.nthu.edu.tw

Abstract

For survival, a living agent (e.g., human in Fig. 1(a)) must have the ability to assess risk (1) by temporally anticipating accidents before they occur (Fig. 1(b)), and (2) by spatially localizing risky regions (Fig. 1(c)) in the environment to move away from threats. In this paper, we take an agent-centric approach to study the accident anticipation and risky region localization tasks. We propose a novel softattention Recurrent Neural Network (RNN) which explicitly models both spatial and appearance-wise non-linear interaction between the agent triggering the event and another agent or static-region involved. In order to test our proposed method, we introduce the Epic Fail (EF) dataset consisting of 3000 viral videos capturing various accidents. In the experiments, we evaluate the risk assessment accuracy both in the temporal domain (accident anticipation) and spatial domain (risky region localization) on our EF dataset and the Street Accident (SA) dataset. Our method consistently outperforms other baselines on both datasets.

1. Introduction

A very important goal for living agents in the world is survival. In order to survive, they naturally have the ability to assess risk. For instance, humans exhibit emotional responses while taking or observing risky actions [23], in an unconscious process that appears to happen without sophisticated reasoning [20]. Furthermore, humans have the ability to turn their attention on the risky areas of the environment more often than others [22], as risk does not distribute uniformly across the environment. Such risk localization is very important for the agent to move away to safety. On the other hand, humans also have the ability to assess longerterm risk by imagining future situations. In this case, highlevel reasoning techniques (imagination, simulation) can be used to assess risk in a longer term. Such anticipation ability is also critical for the agent to react before accident occurs. We are inspired by these key capabilities of human intelligence and perception to study the problem of risk assessment from a computer vision perspective.



Figure 1. Illustration of risk assessment. (a) we show an image overlaid with a human agent from different frames and label the risky region before accident occurs in an orange box. (b) risk map for environment and accident anticipation probability through time t. (c) risky region (orange box) at the instant when accident occurs.

Towards this goal, we introduce the problem of risk assessment from an agent-centric point of view. That is, given the observed past and current behavior of each agent in a video, we tackle the problem by answering two key questions centered around each agent. First, will the agent encounter an accident in the near future? This corresponds to the task of *accident anticipation*, where we would like to predict an accident before it occurs. Second, in which region in the environment might the accident take place? This corresponds to the task of *risky region localization*, where we would like to spatially localize the regions in the scene that might be involved in a future accident.

We face two key but difficult challenges in this problem. First, note that similar visual appearances will frequently correspond to vastly different levels of risk, as risk is dependent on context and interactions between the agent and the environment. Therefore, we must explicitly consider appearances and spatial relations between agents and regions in the scene. A second challenge is that of capturing long term temporal dependencies and causalities that underlie risk events. This could be tackled by explicitly forecasting relationships between the agent and the environment.

Some early attempts focus on assessing risk related to the environment [28] or correlating the statistical occurrance of activities to static scenes [2]. Instead, we aim at assessing risk explicitly triggered by the actions of an agent and its interaction with the environment, by anticipating accidents and localizing risky regions in the scene. The task of accident anticipation is related to early activity recognition [8, 25], and event anticipation [10, 15]. However, these are primarily categorization approaches that discriminate actions into separate semantic classes. In our case, we are not as interested in the semantic categories of the actions, but in reasoning about the probability of accident in the near future. Risky region localization has less precursors in the vision literature. The closest is work on humanobject interaction from an action recognition perspective, but these methods usually model object categories explicitly and their correlation to action classes [33, 15].

We introduce a novel model for agent-centric risk assessment. Our model encodes the behavior of an agent into a distributed representation using a Recurrent Neural Network (RNN). Given the agent representation, we introduce a novel dynamic parameter predictor inspired by Noh et al. [21] to measure the riskiness of each region with respect to the agent. The parameters efficiently consider relative spatial relations and coupled appearances between the agent and the region. Next, our model takes the agent representation and appearance of the risky regions as the input of another temporal-level RNN for accident anticipation. Moreover, the hidden representation of the temporal-level RNN is used to imagine and simulate the future trajectory of the agent. The future trajectory can be used as new inputs to our model so that we can assess the risk in long-term.

Our main contributions are: (*i*) We utilize the dynamic parameter layer to efficiently model the relative spatial relation and coupled appearance between agent and region. (*ii*) We use the generative property of RNN to self-train it to encode the behavior of the agent as well as generate (i.e., imagine) its future trajectory. (*iii*) The imagined future trajectory becomes new inputs to our model to assess risk in a longer term. (*iv*) To the best of our knowledge, the new Epic Fail (EF) video dataset is the first agent-centric risk assessment dataset for computer vision research.

2. Related Work

We give an overview of related work on risk assessment from visual observations, early event recognition and anticipation, as well as parameter prediction for deep networks.

Risk assessment given visual observation has not been widely explored. Valenzuela *et al.* [28] propose to assess landslide risk from topographic images. Since landslide is caused by intense rain in localities where there was unplanned occupation of slopes of hills and mountains, detecting these slopes in topographic images helps us to predict the risk of landslide. Arietta *et al.* [2] propose to use streetlevel images to predict the crime rate (risk of crime) at each geographic location. Koshla *et al.* [11] predict crime rates in an area without real-time criminal activity information, by correlating the appearance of a scene to properties such distance to public places, businesses, etc. However, these approaches assess risk caused either by the environment or by priors on social activities, whereas we focus on assessing risk explicitly triggered by the observed actions of an agent and its interactions with the environment.

Risk assessment is related to predicting the possibility of catastrophic events occurring in the future. In early activity recognition, the focus is to predict activities before they are completed, such as recognizing a smile as early as the corners of the mouth curve up. For example, Ryoo [25] introduces a probability model for early activity prediction; Hoai et al. [8] propose a max-margin model to handle partial observation; and Lan et al. [16] propose the hierarchical movemes representation for predicting future activities. In activity anticipation, the goal is to predict events even before they occur. For instance, Jain *et al.* [10] propose to fuse multiple sensors to anticipate the actions of a driver; Chan et al. [4] introduce a dynamic soft-attention-based RNN to anticipate accidents on the road from dashcam videos; and Vondrick et al. [29] propose to learn temporal knowledge from unlabeled videos for anticipation. However, these focus on activity categories and do not study risk assessment of objects and regions in the video.

Anticipation has been applied in tasks other than event anticipation. Kitani et al. [13] propose to forecast human trajectory by surrounding physical environment (e.g., road, pavement, etc.) and show that the forecasted trajectory can be used to improve object tracking accuracy. Walker et al. [30] propose to forecast dense pixel trajectories from a static image. Yuen and Torralba [34] propose to predict motion from still images. Julian et al. [31] propose a novel visual appearance prediction method based on mid-level visual elements with temporal modeling methods. Event anticipation is also popular in the robotic community. Wang et al. [32] propose a latent variable model for inferring human intentions. Koppula and Saxena [15] address the problem by observing RGB-D data, and apply their method to assist humans in daily tasks. Finally, human activity anticipation can also improve human-robot collaboration [14, 18].

Parameter prediction in deep networks is a relatively new idea. Ba *et al.* [3] propose a zero-shot classifier for unseen classes by predicting the parameters of a classifier using text information. Noh *et al.* [21] propose to dynamically predict the parameters for image question answering depending on the given textual question. Inspired by [21], we introduce a novel dynamic parameter predictor layer for estimating spatial riskiness depending on the agent behavior.

3. Agent-centric Risk Assessment

We now define the task of agent-centric risk assessment and present our model. Given a video frame at time t, we observe information about the agent and multiple regions.



Figure 2. Illustration of our method. Panel (a) shows the pre-process to obtain appearance and location information for both agent (a, p) and regions (R, L). Panel (b) shows all the components (Sec.3,2, 3.3, and 3.4) in our model to predict riskiness of all regions S and anticipated accident probability y. Acc. denotes accident. Panel (c) illustrates how the imagined agent location \hat{p}_{t+K} triggers our model to reassess risk (S, y). In panel (b), a diamond shape node denotes a switch. It is used to control our model to imagine or take observation. In panel (c), the transformation block corresponds to Eq. 8. All dash arrows represent information across frames. Note that the anticipated accident probability increases from 0.5 at frame t to 0.7 at frame t+1.

We assume we have access to the appearance vector a_t and bounding box location $p_t = [x_t, y_t, w_t, h_t]$ of the agent. We also capture information about a set of N candidate risky regions, $R_t = \{r_t^i\}_{i=0}^N$ and $L_t = \{l_t^i\}_{i=0}^N$, where r_t^i is the appearance and l_t^i the location of region i. When we observe a video sequence from t = 0 to the current frame \hat{t} , our accumulated agent information is $\{(a_t, p_t)\}_{t=0}^{\hat{t}}$ and our accumulated region information is $\{(R_t, L_t)\}_{t=0}^{\hat{t}}$. The goal is to predict two outputs corresponding to the tasks of accident anticipation and risky region localization. The first is the accident anticipation probability $y_{\hat{t}} \in [0, 1]$ at current frame \hat{t} . The second is the riskiness score of all candidate regions at current frame \hat{t} , $S_{\hat{t}} = \{s_{\hat{t}}^i\}_{i=0}^N$, where $s_{\hat{t}}^i \in [0, 1]$ is the risk probability for the *i*-th region. Next, we give an overview of how our model infers $y_{\hat{t}}$ and $S_{\hat{t}}$.

3.1. Model Overview

Our model consists of three main components. The first is the agent-region interaction component. We propose to dynamically predict parameters to infer riskiness of a region s depending on the behavior of the agent and relative location of the region concerning the agent's location. The second is the Holistic Accident Anticipation Module incorporating information from both agent and risky regions to infer the accident anticipation probability. Finally, the recurrent component with two Recurrent Neural Networks (RNNs). One RNN aggregates behavior of the agent, while the other aggregates the holistic accident anticipation information. In the following, we describe each component in details.

3.2. Agent-Region Interaction Module

The goal of this module is to infer the risk probability S^i for each region in a frame. Consider for example the

unicycler agent in frame t of Fig. 2(a). Intuitively, the risk of region i should be dependent on: the appearance of the region r^i to verify if the objects in a region are risky, as a region covering the stairs; the appearance of the agent a, as the stairs might be riskier for a unicycler than for a pedestrian; and spatial relationship between the agent and the region u^i , as stairs close to the unicycler indicate more risk In this way, we write risk probability s^i as:

$$s^{i} = g(w_{r}^{T} \cdot r^{i}) \in [0, 1],$$
 (1)

where g is a sigmoid to ensure valid probability estimates. Note that this indicates that region riskiness only depends on region appearance r^i . To encode dependencies on a and u^i , we propose to dynamically predict the parameter w_r :

$$w_r(a, u^i) = \sigma \left(W_f \cdot \begin{bmatrix} a & \sigma \left(W_u \cdot u^i \right) \end{bmatrix}^T \right), \tag{2}$$

where σ is a rectified linear unit (ReLU), and W_f, W_u are the parameters of two fully connected layers. We encode the agent-region spatial relationship u^i with a 9-dimensional vector that we compute from the agent bounding box p and region bounding box l^i . Fig. 3 illustrates the components of u^i which concatenates: the normalized relative position of region center $(\Delta x_c, \Delta y_c)$, top-left corner $(\Delta x_{min}, \Delta y_{min})$ and bottom-right corner $(\Delta x_{max}, \Delta y_{max})$; the region relative width Δw and height Δy ; and Intersection over Union (IoU) of the agent box and region boxes.

3.3. Holistic Accident Anticipation Module

The goal of this module is to produce an accident anticipation score y for the current frame. Intuitively, the probabilty of accident y depends on: the appearance of the agent a, as some agents might be more prone to accidents than



Figure 3. Relative configuration of all regions with respect to the agent. Risk assessment of all regions are illustrated with respect to the agent (green box). In our agent-centric perspective, the orange box indicates a risky region and the blue boxes indicate non-risky regions. We normalize the horizontal and vertical axes separately such that the width and height of the agent is unit one. All 9 cues $(\Delta x_c, \Delta y_c, \Delta x_{min}, \Delta y_{min}, \Delta x_{max}, \Delta x_{max}, \Delta w, \Delta h$, IoU) in the configuration are visualized.

others; and the appearance R and risk level S of all regions in the scene, as some specific types of regions might lead to accidents more frequently than others. We encapsulate this intuition by first building a holistic representation q, which we obtain by concatenating the agent appearance a with the consolidated region information \bar{r} :

$$q = \begin{bmatrix} a & \bar{r} \end{bmatrix}^T. \tag{3}$$

We consolidate the region information by weighting each region according to its inferred risk probability:

$$\bar{r} = \phi(S, R) = \sum_{i} s^{i} \cdot r^{i}.$$
(4)

Note that \bar{r} has the same dimension even when the number of regions varies at each frame. The holistic representation q is used to infer accident anticipation probability y,

$$y = \operatorname{softmax}(W_y \cdot q) \in [0, 1]^2,$$
(5)

where W_y is the model parameter, and y[0], y[1] denote the probability of non-accident and accident, respectively.

3.4. Recurrent Temporal Memory for Anticipation

The model we described so far operates on a single frame and does not aggregate the knowledge of the sequence of past observations. Intuitively, incorporating this sequence should help the model understand how the agent and regions move and how their relation with each other evolve in time. To model these sequences, we introduce two RNNs to operate as memory components in our framework.

First, we aggregate the agent appearance and behavior information in the Agent-RNN (RNN_A) , which takes $\{(a_t, p_t)\}_{t=0}^{\hat{t}}$ as inputs and produces an encoding in its hidden vector $\alpha_{\hat{t}}$. We propagate this information by incorporating α in Eq. 2 and Eq. 3, instead of the appearance information *a*. So, Eq. 2 and Eq. 3 can be rewritten as follows.

$$w_r(\alpha, u^i) = \sigma \left(W_f \cdot \begin{bmatrix} \alpha & \sigma \left(W_u \cdot u^i \right) \end{bmatrix}^T \right), \tag{6}$$

$$q = \begin{bmatrix} \alpha & \bar{r} \end{bmatrix}^T. \tag{7}$$

Second, we aggregate the environment risk information by modeling the sequence of holistic representations q. We achieve this by an Accident-Anticipation-RNN (RNN_{AA}), which takes $\{q_t\}_{t=0}^{\hat{t}}$ as input and produces an encoding in its hidden vector $o_{\hat{t}}$. We propagate this information by incorporating o in Eq. 5, instead of the direct use of q.

As a result, our model can predict the accident probability $y_{\hat{t}}$ and region risk scores $S_{\hat{t}}$ as a function of the observations from t = 0 to \hat{t} . In practice, we use LSTM cells [9] to better handle temporal dependencies.

3.5. Imagining Future Risk

One interesting capability for humans is to assess risk by imagining future situations. In the case of Fig. 2, we can imagine the agent moving towards the stairs, which may result in an accident in the near future. We are interested in encoding such *imagination* capability to our model to better anticipate accidents and predict region risk. With the formulation so far, we have a model that can predict the probability of an accident happening in the near feature $t_f > \hat{t}$ from past observations t = 0 to \hat{t} . We include a mechanism in our model that simulates or imagines the future trajectory and location of the agent K frames into the future, which we denote as $\hat{p}_{\hat{t}+K}$. The idea is that once the model predicts the location of the agent in the future, we can ultimately produce new risk scores for all regions \hat{s} and a new accident anticipation probability \hat{y} .

In practice, we use the holistic representation $o_{\hat{t}}$ to infer a 4-dimensional transformation $\mathbf{c} = [c_x, c_y, c_w, c_h]$ that converts the agent location $p_{\hat{t}}$ to the imagined location $\hat{p}_{\hat{t}+K}$:

$$\mathbf{c} = W_c \cdot o_{\hat{t}}, \tag{8}$$
$$\hat{p}_{\hat{t}+K} = \begin{bmatrix} c_x \cdot w_{\hat{t}} + x_{\hat{t}} & c_y \cdot h_{\hat{t}} + y_{\hat{t}} & e^{c_w} \cdot w_{\hat{t}} & e^{c_h} \cdot h_{\hat{t}} \end{bmatrix}. \tag{9}$$

We train the parameters W_c with ground truth transformations \mathbf{c}^* that map ground truth locations $p_{\hat{t}}$ and $p_{\hat{t}+K}$.

Once the model imagines the location of the agent $\hat{p}_{\hat{t}+K}$, we can update the agent-region relationships to $\hat{u}_{\hat{t}+K}$ by recomputing these features using the imagined location. Similarly, we can produce new \hat{w}_r , $\hat{S}_{\hat{t}+K}$, $\hat{r}_{\hat{t}+K}$, $\hat{q}_{\hat{t}+K}$, $\hat{o}_{\hat{t}+K}$ and finally a new $\hat{y}_{\hat{t}+K}$. Note that $\hat{y}_{\hat{t}+K}$ corresponds to the accident anticipation probability that the model produces from the observations at $t = 0 \dots \hat{t}$ and one step of imagining the future position of the agent at time $\hat{t} + K$. In other words, by using this imagination mechanism, the model is able to assess risk without observing any new information. More importantly, the same process can be applied multiple times to imagine further into the future. That is, we can obtain $\hat{y}_{\hat{t}+nK}$ by recursively estimating $\hat{p}_{\hat{t}+nK}$ from $\hat{p}_{\hat{t}+(n-1)K}$ and repeating the process outlined above.

Final prediction. Finally, we estimate risk $(y_{\hat{f}}^F, S_{\hat{f}}^F)$ by fus-

ing the current risk with the imagined risk as follows,

$$y_t^F = \sum_{n=0}^I \lambda_n \hat{y}_{\hat{t}+nK}, \text{ and } S_t^F = \sum_{n=0}^I \lambda_n \hat{S}_{\hat{t}+nK},$$
 (10)

where, λ_n are hyper-parameters, and with slight abuse of notation, we use $\hat{y}_{\hat{t}}$ as $y_{\hat{t}}$ and $\hat{S}_{\hat{t}}$ as $S_{\hat{t}}$.

3.6. Multi-task Learning

The goal of the learning process is to fit all the parameters in our model: W_f , W_u , W_y , W_c , and the parameters of our recurrent models RNN_A and RNN_{AA} . During training, we have access to a set of positive videos that depict accidents and a set of negative videos that depict normal non-accident events. We assume each positive video depicts an accident at time t = T and is annotated with the ground truth agent locations p_0, \ldots, p_T . We also have access to the region bounding boxes ρ_0, \ldots, ρ_T that encapsulate the part of the environment involved in the accident with the agent.

We fit the model parameters using these training examples by minimizing a loss function \mathcal{L} over the multiple tasks that our model performs: accident anticipation, risky region localization, and agent location imagination as follows:

$$\mathcal{L}(Y, S, C) = \mathcal{L}^{A}(Y) + \mathcal{L}^{R}(S) + \mathcal{L}^{P}(C).$$
(11)

Accident Anticipation. We follow [4] to use regular crossentropy for non-accident sequences and exponential crossentropy loss for accident sequences. The exponential loss emphasizes on predictions for times t that are closer to T.

$$L^{A}(y_{t}) = \begin{cases} -\log(y_{t}[0]) & \text{for non-accident} \\ -e^{-(T-t)}\log(y_{t}[1]) & \text{for accident} \end{cases}$$
(12)

$$\mathcal{L}^{A}(Y) = \sum_{t=0}^{T} L^{A}(y_{t}), \tag{13}$$

where $Y = \{y_t\}_{t=0}^T$.

Risky Region Localization. Since *s* is the output of a sigmoid function, we use sigmoid cross entropy loss for risky region localization as follows,

$$L^{R}(y) = \begin{cases} -\log(1-s) & \text{for non-risky region} \\ -\log(s) & \text{for risky region} \end{cases}$$
(14)

$$\mathcal{L}^{R}(\mathcal{S}) = \sum_{t=0}^{T} \sum_{i=0}^{N} L^{R}(s_{t}^{i}),$$
(15)

where $S = \{S_t\}_{t=0}^T$, and the *i*-th region is a risky region if the IoU between region location l^i and any ground truth box in ρ_t is over 0.4.

Agent location imagination. Inspired by [24], we employ a smooth ℓ_1 loss $L^P(\mathbf{c}_t, \mathbf{c}_t^*)$ for agent location imagination:

$$\mathcal{L}^{P}(C) = \sum_{t=0}^{T} L^{P}(\mathbf{c}_{t}, \mathbf{c}_{t}^{*}), \qquad (16)$$

where $C = {\{\mathbf{c}_t\}_{t=0}^T \text{ and } \mathbf{c}^* \text{ is ground truth transformation.}}$ Anticipation and Localization with Imagination loss. As described in Sec. 3.5, once we imagine the future agent location \hat{p} , we can update \hat{y} and \hat{s} iteratively. We incorporate these estimates by rewriting Eq. 11 as:

$$\mathcal{L}_{I}(Y, S, C) = \mathcal{L}^{P}(C) + \sum_{n=0}^{I} \lambda_{n} \left(\mathcal{L}^{A}(Y^{n}) + \mathcal{L}^{R}(S^{n}) \right).$$
(17)

where Y^n and S^n are the predictions at the *n*-th imagination iteration, I is the number of imagination steps, and λ_n are the same hyper-parameters in Eq. 10. Because our model is fully differentiable, the model can be end-to-end training.

Training with noisy agent info. Training and evaluating with ground truth agents does not reflect the challenges of handling noisy agent information as in real world. Hence, we apply online tracking-by-detection (TD) to obtain candidate agent tracks (see [37] for detail). At training, we utilize both candidate agent and ground truth tracks. The ground truth accident labels are shared among all tracks. At testing, we obtain candidate agent tracks using the same approach. For each candidate agent track, we apply our method to obtain per frame accident anticipation probability. At each frame, we take the maximum probability as the video-level anticipation probability. Then, we take the estimated riskiness of regions from the agent with maximum accident probability as the final per frame riskiness.

Summary of supervision. In training, temporal location of the accident within the video and bounding boxes of risky region at each frame in the positive examples. In testing, these information is only used for performance evaluation.

3.7. Implementation Details

We set I = 1, K = 5, $\lambda_0 = 0.6$ and $\lambda_1 = 0.4$. This are set empirically without heavily tuning. At each frame, we use Faster R-CNN [24] to propose 300 candidate risky regions. We find this setting to be effective, since the average recall at 0.4 IOU is 79.5%, 74.9% on Epic Fail dataset and Street Accident dataset, respectively. For each candidate risky region, we extract *pool*5 feature and utilize Global Average Pooling [17] to obtain an one dimensional representation. For an agent, we extract fc7 feature as the representation. All the feature extractor is using VGG16 model [27]. We use Adam [12] as optimizer with default hyperparameters and 0.0001 learning rate and set batch size by 5. The model selection is done by early stopping [7].

4. Dataset

In order to evaluate our method, We collect a large-scale Epic Fail (EF) dataset consisting of user-generated videos, where a large portion of them involves epic "human" accidents such as the parkour failure in Fig. 1. We also evaluate on the latest Street Accident (SA) dataset [4], where both humans and/or vehicles can involve in accidents. We further describe each dataset in detail.

4.1. Epic Fail (EF) dataset

The raw videos in EF dataset are harvested from YouTube channels and Zeng et al. [36, 35]. To build our new EF dataset, we first manually identify the time when accident occurs in a subset of raw videos. Then, we sample short videos of 3-4 seconds from the subset. In total, we sample 3K videos and slip them into 2K training and 1K testing videos. In the training set, there are 1K posi-



Figure 4. Examples in EF dataset. In each row, we show sampled frames form a video. For positve videos, we also show the annotated risky regions (orange boxes).

tive videos and 1K negative videos. In the testing set, there are 609 positive videos and 391 negative videos. For positive videos, we ensure accident happens at the end of each video. For negative videos, we ensure no sign of accident appears. Note the types of accidents in our dataset is very diverse which include all kinds of skateboarder failure, skier failure, parkour failure, etc.

In order to train and evaluate the risk assessment performance, we ask users to annotate the dataset with the following ground truth labels. Firstly, all videos (both positive and negative) are annotated with ground truth agent trajectory. Risky regions in all positive videos are also annotated, where we ask annotator to annotate the region causing the failure event. The tool that users use for annotating the agent and risky region is an interactive annotation and segmentation tool called iSeg developed by [26]. The agents and the risky regions are annotated by bounding boxes. Even with the help of the annotation tool, annotating bounding boxes are still time-consuming. Hence, training data is only annotated at every 15 frames. However, testing data is carefully annotate at every frame. In this dataset, there are not too many cases with multiple risky regions because the dataset is collected from the user-generated videos. The user-generated videos typically have a main agent and an apparent region causing the accident. More detail and the data for EF dataset can be found in [37].

4.2. Street Accident (SA) dataset

The SA dataset [4] is captured across six cities in Taiwan with high-quality dashcam (720p in resolution) and has diverse accidents occur in all videos consisting 100 frames. These accidents include 42.6% motorbike hits car, 19.7% car hits car, 15.6% motorbike hits motorbike, and 20% other types. The SA dataset also provides the annotation about the time when an accident occurs and the trajectories of objects involved in the accident. The dataset consists of 596 positive examples containing the moment of accident at the last 10 frames, and 1137 negative examples containing no accident. In the SA dataset, it contains 1266 training videos (446 positive and 820 negative examples) and 467 testing

videos (150 positive and 317 negative examples). In this dataset, many cases have multiple risky regions because a street accident usually involves multiple vehicles.

5. Experiments

We first describe the baseline methods and variants of our method. Then, we define the evaluation metrics. Finally, we show that our method achieves the best performance in both accident anticipation and risky region localization on both EF and SA datasets.

Baselines. We compare the following state-of-the-art methods with our method.

- DSA: Dynamic Soft-Attention [4].

- SP: Social Pooling [1]. In the agent-centric representation, we apply SP [1] to pool the nearby regions information (r, l). The agent information (a, p) and the SP pooled feature are concatenated and fed into a LSTM for anticipating accident probability at each frame.

- *R***CNN* [6]. We extend the the Contextual action classification method R*CNN [6] for accident anticipation. There are two extensions: (1) replacing classification loss with the same anticipation loss in Sec. 3.6, and (2) removing the original IOU constraint for the model to observe all candidate risky regions. Note that R*CNN uses hard-attention to select a region with maximum confidence, whereas our method uses soft-attention as in Eq. 4.

- *L*-*R***CNN*, an extended R*CNN to incorporate temporal modeling with LSTM. We add a LSTM to aggregate information across time similar to the RNN_{AA} in Sec. 3.4.

Ablation studies. We also evaluate the following four variants of our Risk Assessment (RA) model involving adding memory or not and applying imagination or not. Note that w denotes "with" and w/o denotes "without".

- *RA*. w/o memory, w/o imagining. This model observes a single frame without aggregating temporal information.

- *RAI*. w/o memory and w/ imagining. We add imagining layer (Sec. 3.5) to the RA model.

- *L-RA*. w/ memory and w/o imagining. We add LSTM cells (Sec. 3.4) to the RA model.

- *L-RAI*. w/ memory and w/ imagining. This is our full model which can handle the temporal information and imagine the future.

5.1. Evaluation Metrics

For accident anticipation, we are interested in not only the precision v.s. recall, but also the first time \hat{t} when the anticipation probability is above a threshold γ . Let's assume the accident occurs at time T. We follow [4] and define Time-to-Accident (TTA) as $T - \hat{t}$. Recall that given different γ , one can compute a precision and a recall. Similarly, we can compute TTA for each recalled positive video. This implies one can plot a TTA v.s. recall. We propose to report the "average TTA" across different recall to summarize the TTA (referred to as ATTA). If the ATTA value is higher,

Dataset	E	F	SA	
w/o memory	mAP (%)	ATTA (s)	mAP (%)	ATTA (s)
R*CNN	68.6	2.47	40.7	2.64
RA	72.2	2.10	47.8	2.55
RAI	72.4	2.13	48.8	2.62
w memory	mAP (%)	ATTA (s)	mAP (%)	ATTA (s)
DSA	45.7	1.16	48.1	1.34
SP	40.5	0.88	47.3	1.66
L-R*CNN	69.6	2.54	37.4	3.13
L-RA	74.2	1.84	49.1	3.04
L-RAI	75.1	2.23	51.4	3.01

Table 1. Quantitative results of accident anticipation. We evaluate accident anticipation by estimating mean average precision and average time-to-accident (ATTA) metrics. Bold-fonts indicate our best performance. Italics-fonts indicate best baseline performance.

the model can anticipate the accident earlier. We also report mean average precision (mAP) for all the videos. See detailed explanation of ATTA in [37].

For the risky region estimation, we use the object detection metric [5] with $IOU \ge 0.4$ as the positive detection criteria [4]. This is because annotating ground truth boxes in videos is very time-consuming. As a results, the quality of ground truth boxes on both EF and SA dataset is slightly worse than other object detection dataset. Note that each frame might contain more than one risky region. Moreover, risky region in a positive clip could appear and disappear due to occlusion or camera motion. Hence, the evaluation of risky regions is conducted per frame.

5.2. Accident Anticipation

Quantitative results using both mAP and ATTA are shown in Table. 1. For our model, adding memory improves mAP as well as ATTA on both datasets in general (i.e., L-RA outperforms RA and L-RAI outperforms RAI, except L-RA is worse than RA in ATTA on EF dataset.) Imagining future risk effectively improves both evaluation metrics on both datasets (i.e., RAI outperforms RA and L-RAI outperforms L-RA) On the other hand, RA/L-RA outperforms R*CNN/L-R*CNN significantly in mAP. This suggests that our soft-attention using dynamic parameter prediction outperforms hard-attention. Although L-R*CNN outperforms our method in ATTA on both datasets, this earlier anticipation comes with significant more false alarms since there is a significant $\sim 5\%$ drop in anticipation mAP. L-RA also outperforms both DSA and SP. This suggests that our dynamic parameter prediction layer is more effective than social pooling and dynamic soft-attention. Note that DSA and SP do not support risky region localization.

5.3. Risky Region Localization

Quantitative results using mAP of risky region localization are shown in Table. 2. Note that the mAP cannot be 100%, since our results depend on the detections performed by Faster R-CNN for proposing candidate risky regions. Hence, we report the oracle performance in the last row which is achieved by assuming all candidate regions are

Dataset	EF	SA	
w/o memory	mAP (%)	mAP (%)	
R*CNN	3.47	34.7	
RA	12.3	40.1	
RAI	14.1	43.1	
w memory	mAP (%)	mAP (%)	
L-R*CNN	3.5	35.6	
L-RA	14.0	43.8	
L-RAI	15.1	45.4	
Oracle	75.7	92.8	

Table 2. Quantitative results of risky region estimation. We evaluate risky region using traditional object detection metric and compute mean average precision over entire testing set.

classified correctly. This serves as the upper-bound performance. For our model, adding memory module improves mAP on both datasets (i.e., L-RA outperforms RA and L-RAI outperforms RAI). Imagining future risk effectively improves mAP on both datasets (i.e., RAI outperforms RA and L-RAI outperforms RAI). On the other hand, L-RA/RA significantly outperforms L-R*CNN/R*CNN. This suggests that our soft-attention using dynamic parameter prediction outperforms hard-attention.

5.4. Qualitative Results

We show qualitative results of accident anticipation and risky region localization in Fig. 5. From the positive and negative examples, our method shows great ability to differentiate them. These examples also demonstrate the ability to localize risky regions of different categories (e.g., car and bars). In the failure example, our system identifies the inflatable pool as potential risky, which is fairly reasonable. See more qualitative results in [37].

6. Conclusion

We introduce new risk assessment tasks including (1) accident anticipation and (2) risky region localization. To tackle these tasks, we propose a novel model with two main innovations: (1) dynamic parameter prediction to capture the relative spatial relation and appearance-wise coupling between agent and risky regions. Our proposed method outperforms baselines methods significantly on both accident anticipation and risky region estimation. In the future, we plan to extend our imagining layer for the environment. We believe that stimulate both agent and environment in the future simultaneously would enhance the model and give a way to explain how does the model anticipate the accident.

Acknowledgement. We thank MOST 104-2221-E-007-089, MOST 106-2633-E-002-001, National Taiwan University (NTU-106R104045), NOVATEK Fellowship, MediaTek, and Panasonic for their support. We thank Chi-Wen Huang, and Yu-Chieh Cheng for their collaboration. We thank Alexandra Alahi, Zelun Luo, and Shyamal Buch for helpful comments and discussion.



Figure 5. Qualitative Results. We set 0.9 as threshold of triggering accident anticipation to show the qualitative results. In each example, we show a typical example with accident anticipation probability (bottom row), the heat map (yellow for high risk and blue for low risk) for the risky regions (middle row) and ground truth risky region (orange box in top row). For risky heat map, we average risky confidences of covering boxes for each pixel and draw the map by using Matlab [19] imagesc tool. More detail for drawing heat map can be found in [37]. The first and second one are positive examples. The third and fourth one are negative examples. The last one is a failure case, where the model misunderstands risky regions so that it has higher accident anticipation probability at first. However, after long-term observation, the model correct the anticipation probability.

References

- A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese. Social LSTM: Human trajectory prediction in crowded spaces. In *CVPR*, 2016. 6
- [2] S. M. Arietta, A. A. Efros, R. Ramamoorthi, and M. Agrawala. City forensics: Using visual elements to predict non-visual city attributes. In *SciVis*, 2014. 2
- [3] J. Ba, K. Swersky, S. Fidler, and R. Salakhutdinov. Predicting deep zero-shot convolutional neural networks using textual descriptions. In *ICCV*, 2015. 2
- [4] F.-H. Chan, Y.-T. Chen, Y. Xiang, and M. Sun. Anticipating accidents in dashcam videos. In ACCV, 2016. 2, 5, 6, 7
- [5] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman. The pascal visual object classes (VOC) challenge. In *IJCV*, 2010. 7
- [6] G. Gkioxari, R. Girshick, and J. Malik. Contextual action recognition with R*CNN. In CVPR, 2016. 6
- [7] S. Haykin. Neural networks, a comprehensive foundation. 1994. 5
- [8] M. Hoai and F. De la Torre. Max-margin early event detectors. In CVPR, 2012. 2
- [9] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural Computation*, 1997. 4
- [10] A. Jain, A. Singh, H. S. Koppula, S. Soh, and A. Saxena. Recurrent neural networks for driver activity anticipation via sensory-fusion architecture. In *ICRA*, 2016. 2
- [11] A. Khosla, B. An, J. J. Lim, and A. Torralba. Looking beyond the visible scene. In CVPR, 2014. 2
- [12] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [13] K. M. Kitani, B. D. Ziebart, J. A. D. Bagnell, and M. Hebert Activity forecasting. In ECCV, 2012. 2
- [14] H. S. Koppula, A. Jain, and A. Saxena. Anticipatory planning for human-robot teams. In *ISER*, 2014. 2
- [15] H. S. Koppula and A. Saxena. Anticipating human activities using object affordances for reactive robotic response. In *TPAMI*, 2016. 2
- [16] T. Lan, T.-C. Chen, and S. Savarese. A hierarchical representation for future action prediction. In ECCV, 2014. 2
- [17] M. Lin, Q. Chen, and S. Yan. Network in network. In *ICLR*, 2014. 5
- [18] J. Mainprice and D. Berenson. Human-robot collaborative manipulation planning using early prediction of human motion. In *IROS*, 2013. 2
- [19] MATLAB. version 7.10.0 (R2010a). The MathWorks Inc., Natick, Massachusetts, 2010. 8
- [20] M. Naber, M. Hilger, and W. Einhäuser. Animal detection and identification in natural scenes: image statistics and emotional valence. In *Journal of vision*, 2012. 1
- [21] H. Noh, P. H. Seo, and B. Han. Image question answering using convolutional neural network with dynamic parameter prediction. In *CVPR*, 2016. 2
- [22] O. C. Okonkwo, M. Crowe, V. G. Wadley, and K. Ball. Visual attention and self-regulation of driving among older adults. In *International Psychogeriatrics*, 2008. 1

- [23] L. Pessoa and R. Adolphs. Emotion processing and the amygdala: from a "low road" to "many roads" of evaluating biological significance. In *Nature reviews neuroscience*, 2010. 1
- [24] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, 2015. 5
- [25] M. S. Ryoo. Human activity prediction: Early recognition of ongoing activities from streaming videos. In *ICCV*, 2011. 2
- [26] J. Schöning, P. Faion, and G. Heidemann. Pixel-wise ground truth annotation in videos. In *ICPRAM*, 2016. 6
- [27] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
 5
- [28] V. V. Valenzuela, R. D. Lins, and H. M. de Oliveira. Application of enhanced-2d-cwt in topographic images for mapping landslide risk areas. In *ICIAR*, 2013. 1, 2
- [29] C. Vondrick, H. Pirsiavash, and A. Torralba. Anticipating visual representations from unlabeled video. In *CVPR*, 2016.
 2
- [30] J. Walker, C. Doersch, A. Gupta, and M. Hebert. An uncertain future: Forecasting from static images using variational autoencoders. In *ECCV*, 2016. 2
- [31] J. Walker, A. Gupta, and M. Hebert. Patch to the future: Unsupervised visual prediction. In *CVPR*, 2014. 2
- [32] Z. Wang, M. Deisenroth, H. Ben Amor, D. Vogt, B. Schölkopf, and J. Peters. Probabilistic modeling of human movements for intention inference. In *RSS*, 2012. 2
- [33] B. Yao and L. Fei-Fei. Modeling mutual context of object and human pose in human-object interaction activities. In *CVPR*, 2010. 2
- [34] J. Yuen and A. Torralba. A data-driven approach for event prediction. In ECCV, 2010. 2
- [35] K.-H. Zeng, T.-H. Chen, C.-Y. Chuang, Y.-H. Liao, J. C. Niebles, and M. Sun. Leveraging video descriptions to learn video question answering. In AAAI, 2017. 5
- [36] K.-H. Zeng, T.-H. Chen, J. C. Niebles, and M. Sun. Title generation for user generated videos. In ECCV, 2016. 5
- [37] K.-H. Zeng, S.-H. Chou, F.-H. Chan, J. C. Niebles, and M. Sun. Technical report of agent-centric risk assessment: accident anticipation and risky region localization. http://aliensunmin.github.io/project/ video-Forecasting/. 5, 6, 7, 8