Accident Forecasting in CCTV Traffic Camera Videos

Ankit Shah[1,*], Jean Baptiste Lamare[1,*], Tuan Nguyen Anh[1,2,*], Alexander Hauptmann[1]
[1]Carnegie Mellon University, Pittsburgh, USA [2]University of Tokyo, Japan
[*] indicates equal contributions.
{aps1, jlamare}@andrew.cmu.edu, tuannguyenanh@g.ecc.u-tokyo.ac.jp, alex@cs.cmu.edu

Abstract

This paper presents a novel dataset for traffic accidents analysis. Our goal is to resolve the lack of public data for research about automatic spatio-temporal annotations for traffic safety in the roads. Our *Car Accident Detection and Prediction (CADP) dataset* consists of 1,416 video segments collected from YouTube, with 205 video segments having full spatio-temporal annotations. To the best of our knowledge, our dataset is largest in terms of number of traffic accidents, compared to related datasets. Through the analysis of the proposed dataset, we observed a significant degradation of object detection in pedestrian category in our dataset, due to the object sizes and complexity of the scenes. To this end, we propose to integrate contextual information into conventional Faster R-CNN using *Context Mining (CM)* and *Augmented Context Mining (ACM)* to complement the accuracy for small pedestrian detection. Our experiments indicate a considerable improvement in object detection accuracy: +8.51% for CM and +6.20% for ACM. For person (pedestrian) category, we observed significant improvements: +46.45% for CM and +45.22% for ACM, compared to Faster R-CNN. Finally, we demonstrate the performance of accident forecasting in our dataset using Faster R-CNN and an Accident LSTM architecture. We achieved an average of 1.359 seconds in terms of Time-To-Accident measure with an Average Precision of 47.36%. Our Webpage for the paper is `https://goo.gl/cqK2wE`

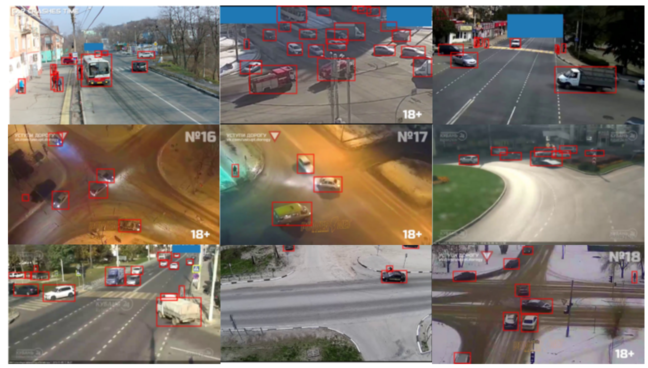

(a)

(b)

Figure 1: *(a) Can you depict* ***where*** *the accidents happen in the image plane?; (b) Can you* ***identify*** *and* ***forecast*** *the sequences containing accidents?* Best viewed in color.

Introduction

According to the National Safety Council, an estimated 40,200 people died on the nation's road in 2016, making motor vehicle crashes the second leading causes of unintentional deaths in the United States (The National Safety Council 2017). Vulnerable road users (jaywalkers), alcohol impaired driving, speeding, seat belts, drowsy and fatigue driving are major risk factors of injury fatalities on the road. A majority of the deaths came from low and middle-income countries where there are many risks for pedestrians, cyclists and motor vehicles on the road. With the desire to improve the road safety, many efforts have been made by research communities to improve the technology to forecast and detect traffic accidents. For instance, the European Research Program PROMETHEUS (Williams 1992) and the ADVISOR project (Naylor and Attwood 2003) have been established to involve research institutes in both academic and industry to detect abnormalities and to reduce traffic fatalities, as well as to enhance the road safety. Our work is devoted toward this objective with neutral views of the accidents from traffic cameras which are installed highly on a corner of the road. The advantage of third-person views is twofold: (i) Unlike the personal views, third-person views have a fixed and wider view because they are mounted higher; and (ii) Unlike the personal views, traffic camera views can be used in the public for a vast amount of vehicles daily, thus, the cost per vehicle per day is lower. While the former enhances the views of traffic accidents, the latter enhances the trade-offs between cost and safety: higher quality (HD 720p-1080p) and better featured cameras, such as Palt-Tilt-Zoom HD cameras, can be used with low cost to monitor the public crowds. Although the exploitation of traffic camera views is promising, the number of datasets aimed at learning to detect and predict the accidents on those views is limited due to several unaffordable factors: (i) traffic accidents are rare events, thus, acquiring enough data by recording in a road intersection is infeasible because one may have to wait endlessly for the an accident to happen; and (ii) the access to traffic camera data is legally difficult to obtain in practice. To this end, we propose an effective data collection process to exploit the edge-case data: YouTube videos of traffic accidents have been uploaded by users over the world. We exploited the search engine of YouTube, and added our annotations processes using both internal annotators and outside workers to build a novel dataset, the Car Accident Detection and Prediction (CADP) dataset for multiple purposes: temporal segmentation, object detection, tracking, vehicle collision, accident detection and prediction. Our dataset contains 230 videos, each video containing at least one accident captured from fixed traffic camera views and 1,416 segments of traffic accidents. Moreover, we selected 205 segments with HD quality to annotate spatio-temporal data for object detection, tracking and collision detection.

Target tasks The release of CADP dataset contains the fully annotated 205 video segments using VATIC tool (Vondrick, Patterson, and Ramanan 2013). It contains the bounding box annotation for each frames of 6 object categories: "Person", "Car", "Bus", "Two-wheeler", "Three-wheeler", and "Others". It also contains the tracking annotations with attribute information of timing that collisions between objects happen. Therefore, we consider three tasks: *object detection, accident detection and forecasting*[1].

Contributions Our contributions are as follows:

- We introduce a new spatio-temporally annotated dataset, the CADP dataset, for accident forecasting using traffic camera views. Our dataset provides a novel view for traffic accident learning, and we hope to contribute to the enhancement of re search on driving education as well as road safety.
- We apply state-of-the-art object detection models such as Faster R-CNN and accident forecasting models to our dataset and show their results.
- We exploit the contextual information around the object bounding box and test the impact of Context Mining and Augmented Context Mining within Faster R-CNN to improve the detection of small objects such as person and improve the Faster R-CNN baseline scores.

The remainder of this paper is organized as follows. In the next section we describe the related work in the field of car accident forecasting and advancement in recent tasks such as object detection, pedestrian detection. The Car Accident Dataset section describes our data collection pipeline and annotation strategy and provides basic dataset statistics for the released dataset. Finally, we describe our technical framework for object detection and accident forecasting, with a set of extensive experiments.

[1]Although tracking is possible with our dataset, it is not a simple task, and we will leave it for future exploration. Furthermore, 1211 unlabeled video segments are also available for future annotation efforts as well as for unsupervised, semi-supervised learning purposes.

Related Work

Dataset for Car Modelling and Accidents With the development of the concepts of smart cities and autonomous driving, there are recent works concerning traffic safety monitoring using computer vision techniques. (Datondji et al. 2016) provides information about relevant datasets for traffic monitoring at road intersections: the MIT dataset for traffic camera events (a 19-min video), NGSIM dataset for road traffic modeling, CBSR dataset for single views at complex intersections, CVRR dataset which simulated videos generated for traffic modeling, QMUL dataset that contains a one-hour recording at a busy intersection and KIT dataset which consists of videos with fog, rain and snow to model traffic car behaviour near intersections. (He and Zeng 2017) performs experimentation using Faster R-CNN (Ren et al. 2015) to show the detection performance on the INRIA dataset. For the traffic accident videos, a recent UCF-Crimes dataset (Sultani, Chen, and Shah 2018) has 13 real-world anomalies such as Abuse, Accidents, Shooting and is focused on understanding of violent scenes in video. Dashcam Accident Dataset (DAD) (Chan et al. 2016) uses Dashboard Camera captured videos to perform accident forecasting with 2.4 hours of video data. We believe that both Dashboard camera views and Traffic camera views could provide critical information for predicting accidents. However, traffic cameras give an overview of the complete road and thus will be able to track more vehicles as compared to dashboard camera views.

Object Detection In recent years, there have been many works on the object detection task (Girshick et al. 2014; Girshick 2015; Ren et al. 2015; He et al. 2017; Liu et al. 2016) which utilize the strength of deep learning (LeCun, Bengio, and Hinton 2015) in common benchmarks such as PASCAL VOC (Everingham et al.) and Microsoft COCO (Lin et al. 2014). R-CNN (Girshick et al. 2014) uses a region proposal algorithm as a pre-processing step prior to CNN architecture feature extraction. These proposals are generated using Edge Boxes (Zitnick and Dollár 2014) or Selective Search (Uijlings et al. 2013) and are independent of CNN. SPP-Net (He et al. 2015) were proposed to improve the R-CNN speed by sharing computation. Fast R-CNN (Girshick 2015) reduces the run time exposing the region proposal computation as bottleneck whereas Faster R-CNN implements region proposal mechanism using CNN and thus integrating region proposal as part of the CNN training and prediction (Ren et al. 2015). Mask R-CNN (He et al. 2017), Single-Shot Detector (SSD) (Liu et al. 2016) and FPN (Lin et al. 2017) combine multiple feature maps with different resolutions to handle multiple object sizes.

Pedestrian Detection predicts information about the pedestrian position based on the detection in current frame. (Dollar et al. 2012) provides a comprehensive overview and arguments to replace continuous detection by pedestrian tracking and thus achieve real-time performance for pedestrian detection. (Benenson et al. 2014) shows adding extra features, flow information and context information are complementary additions resulting in significant gains over other strong detectors. (Wang et al. 2018) uses body-part semantic and contextual information. (Li et al. 2017) proposes a

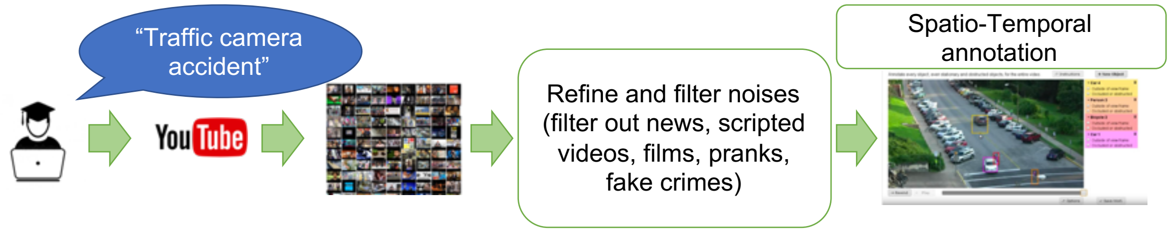

Figure 2: Data collection and annotation for traffic CCTV videos.

Haar-like cascade classifier design for fast pedestrian detection. (Zhang et al. 2016) reviews Region Proposal Network in Faster R-CNN to work better as a standalone classifier whereas downstream classifier degrades the result. (Kong et al. 2018) proposes an extension to Faster R-CNN using contextual information with multi-level features to detect pedestrians in cluttered background obtaining embedding pooling information from a larger area around original area of interest.

Accident Detection and Forecasting In recent years, there have been a few works focusing on the use of cameras for accident forecasting. For example, (Chan et al. 2016) uses Dashboard Cameras for accident forecasting. We believe that there is a strong requirement for those datasets to improve the reaction time of autonomous vehicles such as self-driving cars, and help the road surveillance.

Car Accidents Dataset

Data collection

The major challenge in collecting data for traffic accidents is two-fold: (i) *Abnormality*: because the accidents are rare, although there are live-streams from traffic cameras mounted on the corner of road intersections, this is infeasible to wait for an accident to happen; and (ii) *Access*: access to traffic camera data is often limited. Due to this challenge, the data of traffic accidents from fixed third-person views is often not available for public uses. To this end, in this work, we attempted to exploit an edge case, the traffic accidents captured from traffic camera views available on video sharing websites such as YouTube. The whole pipeline for data collection and annotation can be seen in Figure 2.

Keyword search To collect the data for traffic accidents, we exploited the search engine and resources available in YouTube. We used keywords like "car accidents traffic camera" to search for relevant videos. This step returned 582 YouTube videos.

Refinement However, the collected videos from these queries contain many irrelevant items. To collect only relevant items, we employ three annotators to manually watch and report items as follows. All annotators are instructed to know that our objective is to collect only videos which *contain at least one accident scene which is captured from a traffic CCTV footage*. The annotators then watched all collected videos one by one, and answered a survey about the videos. Besides basic questions to identify whether the annotators want to download the videos based on explained objectives, there are three follow-up questions to filter noisy responses. The first question asks the annotators to justify their concrete reasons for downloading the videos. The second and third questions ask annotators about side aspects of the videos to discover inconsistency in their responses. Videos with inconsistent responses will be removed.

Annotations

After the refinement step, there are 230 videos that were found to be strongly relevant to our objectives. However, for each video, there is only a portion relevant to traffic CCTV footage. Therefore, we employed a two-stage annotation process to get these relevant segments: first we asked human annotators to extract the starting and ending time-stamps for CCTV traffic camera segments from each videos, then we collected the segments and perform the spatio-temporal annotation using the VATIC tool (Vondrick, Patterson, and Ramanan 2013) (see Figure 2).

Stage 1: Temporal segmentation Most of the YouTube videos have a duration of several minutes but contain only several seconds with accidents from traffic CCTV footage. Using the BeaverDam tool[2], human annotators reported the starting and ending timestamps of each relevant segment. Based on the reported results, we extracted the frames of relevant segments using OpenCV[3].

Stage 2: Dense Spatio-Temporal annotation After Stage 1, we have 1416 video segments of positive events. The total duration is 5.24 hours with an average number of frames of 366 frames per video (see Table 1). About 80% of videos have a length from 100 to 600 frames. From short videos (less than 600 frames), we choose 240 videos with HD quality to do dense spatio-temporal annotation. This stage involved four human annotators and has been done in about two months. From the 240 selected videos, the annotators identified 35 videos which are duplicated with one of the other videos. They are the videos with identical contents or cropped (resized) versions of another video. Finally, we have 205 videos with full annotations. The categories of objects are "Person", "Car" (including minivans), "Bus", "Two-wheeler" (including cyclists, motorbikes), "Three-wheeler" and "Others" (objects which are not classified in other categories). About temporal annotations, human annotators were asked to mark when a collision between vehicles/pedestrians on the road happens, and when it ends.

[2] https://github.com/antingshen/BeaverDam

[3] https://opencv.org/

Dataset Statistics

Statistics of our dataset can be found in Table 1 and Figure 2. Some key characteristics of our dataset are as follows:

- **Object size**: As shown in Figure 2(c), a major portion of the CADP dataset is occupied by small objects. Accurate detection of small objects has been a challenge in surveillance videos for a long time. The CADP dataset provides additional samples for these objects from traffic CCTV footage.
- **Video length**: The average length of the videos in the CADP dataset is 366 frames per video, which is 3.66x longer than the dataset from (Chan et al. 2016). The UCF-Crimes (Sultani, Chen, and Shah 2018) also has a category for road incidents with long videos, but only temporal annotations are provided. The CADP dataset provides a set of videos with full spatio-temporal annotations.
- **Number of positive videos** (1416 videos) in our dataset for only traffic accidents is much larger than that in UCF-Crimes (151 videos of road accidents) and DAD (about 600 videos). Note that, in CADP, there are videos with more than one accident. Our dataset is devoted to traffic accidents (positive events), and we did not collect videos of negative events. Negative events can be found easily in other datasets such as DETRAC (Lyu et al. 2017).
- **Time to first accident** is the duration from time 0 in the video to the onset of the first accident. In the fully annotated subset of 205 videos in CADP dataset, this measure is 3.69 seconds in average. Compared to DAD (Chan et al. 2016) (4.50 seconds), CADP has a shorter time-to-first-accident. This characteristic can affect the design of experimentation for accident forecasting.
- **Real-world data**: CADP contains videos collected from YouTube which are captured under various camera types and qualities, weather conditions (see Figure 1) and edited/resampled videos.

Improved Faster R-CNN and Accident Forecasting

Improved Faster R-CNN for Object Detection

Faster R-CNN (Ren et al. 2015) is a deep learning architecture for object detection in still images. It has been successfully applied to object detection in well-known benchmarks such as PASCAL VOC 2007/2012 (Everingham et al.) and Microsoft COCO (Lin et al. 2014), and recently in pedestrian detection domain (Zhang et al. 2016; Ren, Zhu, and Xiao 2018). Like its preceders (Girshick et al. 2014; Girshick 2015), it extracts deep features of each proposal regions using a deep learning backbone such as ResNet-50. However, Faster R-CNN is an end-to-end architecture, because the proposal generation step is done using an internal proposal generation mechanism, the Region Proposal Network (RPN), which reduces the need for dependence on external proposal algorithm such as Selective Search or Edge Boxes, with a sliding window fashion. An important designing aspect of Faster R-CNN is its two-stage design: after features are extracted for proposals, they are classified and regressed to match the anchor boxes. Learning in Faster R-CNN is done with objectives for bounding box regression and classification as follows: $\mathcal{L}_{reg} = \sum_i \text{smooth}_{L1}(t_i - v_i)$, $\mathcal{L}_{cls} = \sum_i -\log p_u$, where u and v are the true class and target bounding box for a groundtruth anchor, p and t are the predicted probability of class u and predicted bounding box. The smooth_{L1} loss function is defined as in (Girshick 2015).

Implementation details We rescale the image to 600 pixels size to smallest size of the image as well as use 3 sizes for the anchor boxes 128^2, 256^2, and 512^2 pixels. Further, the aspect ratios of the anchor boxes is fixed at 1:1, 2:1 and 1:2 pixels as in Faster R-CNN paper (Girshick 2015).

Training procedure The multi-task objective for learning Faster R-CNNs is $\mathcal{L} = \mathcal{L}_{cls} + \lambda\mathcal{L}_{reg}$. For negative mining, we use the standard approach: after the predicted boxes are filtered using non-maximum suppression (NMS) at the overlap threshold 0.7, the RoIs which have confidences in the range [0.1,0.5) are considered as "hard negative", and the RoIs which have confidence larger than 0.5 are considered as "positive". Finally, assuming that we need 32 candidates to contribute to the final loss, we randomly select positive RoIs first to fill at least 16 positions, then we randomly select from the negative RoIs to fill all 32 positions. Only these candidates contribute to the final loss. For data augmentation, we use horizontal and vertical flips.

Context Mining

Context Mining As noticed from Figure 3(c), our dataset consists of objects which are small objects (<100 pixels) in majority. Moreover, from preliminary results on CADP using ResNet-50 backbone Faster R-CNN, we found that there is a significant degradation of accuracy (mAP@0.5) for "Person" (pedestrian) category. We argue that, the reason is because, when captured from CCTV traffic camera footage, a person often looks smaller than other vehicle categories in our dataset. Therefore, the bounding boxes of the pedestrians often contains fewer pixels than other objects. To this end, we propose to mine the context information around the small objects in CADP dataset, by extracting the context information in the RoI pooling layer (Girshick 2015). Given the region of interest of a small object $\mathbf{x}$, a context region $\mathbf{c}$ in common sense (Figure 3(a)) contains $\mathbf{x}$. By extending the context regions, more information is involved into the deep features. Let $\mathcal{C} = \{\mathbf{c}_i\}_{i=1}^n$ be the pooled contextual features, we choose the best responses from them by using Maxout networks (Goodfellow et al. 2013). By applying the dropout process to only linear parts of the signals, the Maxout network is considered to have more generalization ability than traditional Dropout approach. The Maxout operator is applied to $\mathcal{C} \cap \{\mathbf{x}\}$ to obtain final pooled feature $\mathbf{f} = \text{Maxout}(\mathcal{C} \cap \{\mathbf{x}\})$.

Augmented Context Mining The bounding box annotations for small objects can be inaccurate due to human errors (because the object size is too small and difficult for humans to draw a tight bounding box, annotators often draw a larger box or a box which truncates a part of the body). Furthermore, by enlarging the boxes, due to occlusion, the context may involve a different object into the box. Thus, to

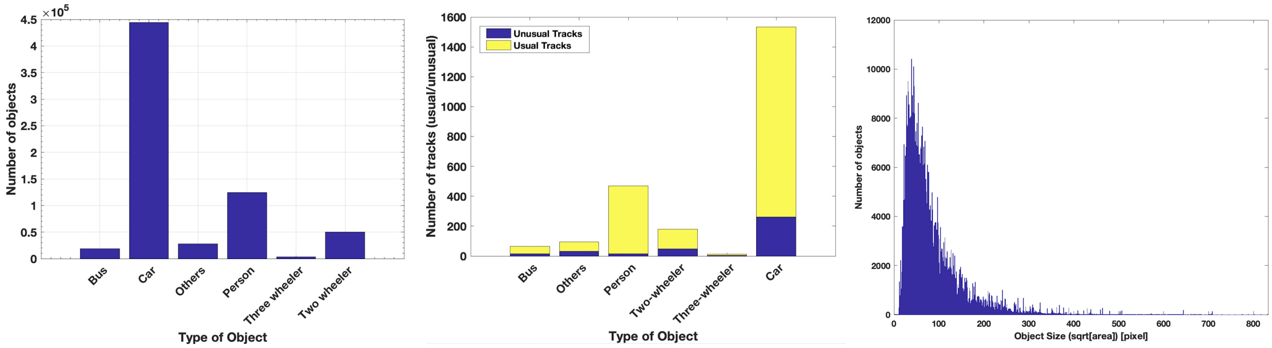

(a) Number of objects by categories (b) Number of usual/unusual tracks (c) Number of objects by sizes.

Figure 3: The statistics of the CADP dataset.

Table 1: Comparison between our dataset and related datasets. **T**: temporal annotation; **S**: spatial annotation (e. g. bounding boxes or pixel-level annotations); **A**: traffic accidents; **C**: videos were captured from a CCTV footage. The "# positives" refer to the number of videos which contain an accident. This statistics is computed from video-level labels (no accident/accident). Our dataset is not the largest in terms of the number of hours, but is the largest in terms of number of accidents (positive events).

Dataset name	# videos	# positives	Total duration	Avg. # frames	T	S	A	C
UCF-Crimes (Sultani, Chen, and Shah 2018)	**1900***	151	**128 hours***	**7247***	✓	✗	✓	✓
DAD (Chan et al. 2016)	1730	620	2.4 hours	100	✓	✓	✓	✗
Ours	1416	**1416**	5.2 hours	366	✓	✓	✓	✓

* These numbers from UCF-Crimes dataset are of 13 categories of crimes (not only for traffic accidents).

address these concerns, we also consider a different context mining, the Augmented Context Mining (ACM) to fully exploit all possible patterns of context around the small person boxes. Rather than gradually extending the small regions to obtain the contexts, we narrow down and extend the small boxes in both horizontal and vertical directions. Given a step stride s and the number of horizontal and vertical steps $m, n \in \{0, \pm 1, \pm 2, \ldots\}$, an *augmented context* $\mathbf{a} = \mathbf{x}_{m,n}$ is defined by extending (when $m, n > 0$) or narrowing down (when $m, n < 0$) $\mathbf{x}$. In the Results section, we compare the performance of these two mining strategies.

Implementation details To control the effects of CM/ACM on small objects, we introduce constraints based on the area ratio of the bounding box and the image. Given a bounding box with area B and image with area I, and a threshold $\alpha \in [0, 1])$. The context mining will be applied to a region if and only if $B \le \alpha S$. We choose $\alpha = 0.01$ in our experiments.

Accident Forecasting

Our framework for Accident Forecasting can be found in Figure 4(c). First, we extracted the features from the last fc layer (2048D) in Faster R-CNN. The features are then fed into the Dynamic-Spatial-Attention LSTM (DSA-LSTM) (Chan et al. 2016) to output accident scores over time. DSA-LSTM is built upon the famous Soft-Attention LSTM (Xu et al. 2015). However, instead of applying spatial attention to regular grid, DSA-LSTM distributes the attentional weights to spatial objects detected by a state-of-the-art detector (Ren et al. 2015). Furthermore, DSA-LSTM applies to "sequences" of frames dynamically (when Soft-Attention LSTM applies to a single frame for caption generation). The full-frame features are also exploited and exponential loss is applied for training with positive sequences. In our view, the exponential loss fits the nature of traffic accidents in CADP because accidents often happen *suddenly* and the damages grow exponentially in a short time. The exponential loss for positive events can be formulated as follows: $\mathcal{L}_p(\{\mathbf{a}\}) = \sum_t -e^{-\max(0, y-t)} \log(a_t)$, where $\mathbf{a}$ is the attended object, y is the time the accident happens, and a_t is the accident probability of $\mathbf{a}$ at time t. For the negative sequences (no accidents), we used cross-entropy loss: $\mathcal{L}_n(\{\mathbf{a}\}) = \sum_t -\log(a_t)$.

Exhaustive negative mining Negative examples are often critical for learning in various situations. However, CADP does not provide explicit negatives. A potential source to mine these examples is existing datasets such as the DETRAC dataset (Lyu et al. 2017). In this work, we exhaustively mine the negatives from positive sequences. Given an accident happens at time t, we mine a positive segment with length 100 frames from time $t - 90$ to time $t + 10$. We randomly mine a segment with length 100 frames which does not overlap with the positive event. Because our videos are longer than 100 frames in average, this mining scheme was possible. However, many accidents happen between the first 100 frames and the Time-to-First-Accident in our test set is only 3.69 seconds, therefore sampling from $t - 90$ may not be possible. To exhaustively mine the negative segments, we append the dummy frames before time 0 to have 90 frames.

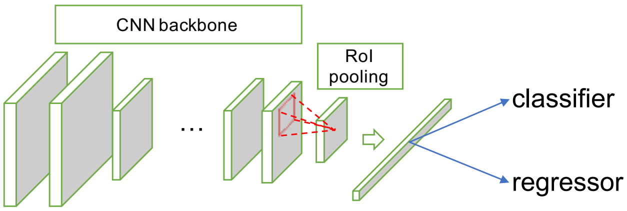

(a) Original Faster R-CNN architecture.

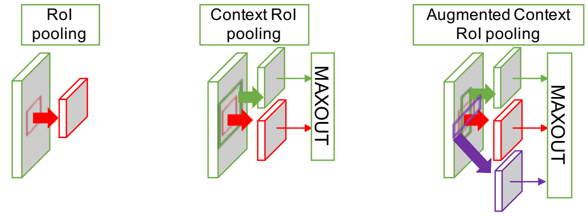

(b) Our proposed RoI pooling layers.

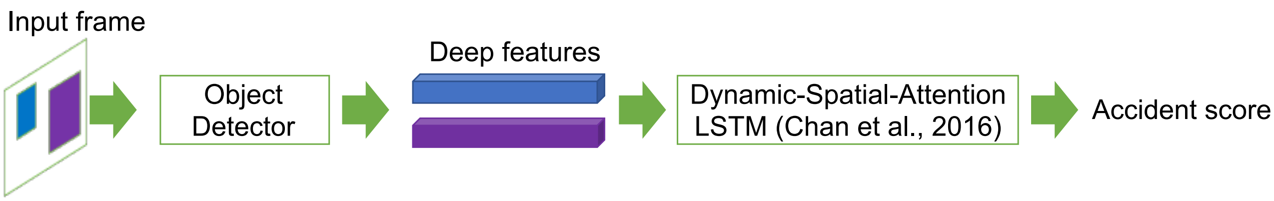

(c) Our workflow to solve the problem of accident forecasting.

Figure 4: Improved Faster R-CNN with Augmented Context Mining and a system for Accident Forecasting.

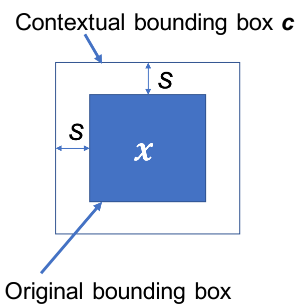

(a)

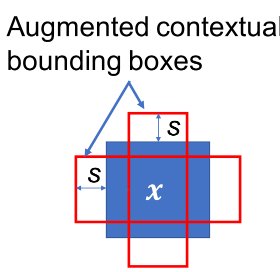

(b)

Figure 5: Contextual patterns. (a) The *contextual* bounding box is created by extending the small object region by s pixels in horizontal and vertical directions. (b) The *augmented contextual* bounding boxes are created by extending or narrowing down the horizontal and vertical sides by s pixels.

Results

Experimental setup

Cross-validation We sample a *trainval* set of 103 videos for training of object detectors and accident forecasters. The 102 remaining videos have been used to test the forecasters. Our choice was contingent on creating a robust model which we wanted to test on enough samples and thus split with a 50:50 ratio (train and test set) where each set has similar set statistics in terms of number of objects. For object detection, from the frames of the 103 videos in *trainval* set, we sample randomly three folds (train/test split) to compute the accuracy. After the cross-validation of object detectors in *trainval* set, we select the best performers as the feature extractor for training the accident forecaster (see Figure 4(c)).

Implementation details Our system is implemented using the Tensorflow framework[4]. During testing, we improve performance by detecting objects with different scales of images (multi-scale testing). For SSD, we use the implementation of (Liu et al. 2016). We fine-tune all object detectors in the CADP *trainval* set until convergence. For accident forecasting, we follow the details described in the previous section. The initial learning rate for Faster R-CNN was 10^{-5} and the Adam optimizer was used.

Evaluation measures For object detection, we use mean Average Precision at IoU=0.5 (mAP@0.5) (Everingham et al.) to assess the accuracy of the detectors. For accident forecasting, we follow (Chan et al. 2016) and use Time-to-Accident (ToA) and recall, precision and average precision (AP). Given the number of true-positives (TP), false-positives (FP) and false-negatives (FN), $Precision = \frac{TP}{TP+FP}$, and $Recall = \frac{TP}{TP+FN}$. To compute the AP, we sample various thresholds and compute ToA, recall and precision at each operating point. AP and mean ToA are computed from these data.

Object Detection

Baselines: SSD vs. Faster R-CNN The comparison between SSD and Faster R-CNN can be found in Table 2. Interestingly, we observed a large gap between the mAP@0.5 of SSD and Faster R-CNN (approx. 19.69%). From the observation about the performance of these two detectors, we choose Faster R-CNN as the baseline for further experimentation. The performance of Faster R-CNN over three sampled folds are reported in Table 3. We can observe stable performances of this detector in CADP *trainval* set. However, we can also observe that the performances degrade and become unstable in the "Person" category.

Context Mining We choose the third fold to perform ablation study on hyper-parameter of Context Mining (CM) and Augmented Context Mining (ACM). The results are

[4] https://www.tensorflow.org/

Table 2: Comparisons between state-of-the-art methods in our dataset. We choose SSD and Faster R-CNN because they are the popular choices for object detection in surveillance video literature. Please see the text for details.

Method	All	Person	Car	Bus	Two-wheeler	Three-wheeler	Others
SSD (Liu et al. 2016)	64.70	-	-	-	-	-	-
Faster R-CNN (Ren et al. 2015)	84.39	52.26	89.39	97.27	77.56	98.88	91.00

Table 3: Cross-validation results for Faster R-CNN. We used mAP@0.5 (0.5 is IOU score) as the measure. Except the "Person" category, Faster R-CNN performs stably across all categories of vehicles.

Fold	All	Person	Car	Bus	Two-wheeler	Three-wheeler	Others
1	82.32	36.89	81.04	97.24	76.00	98.71	94.58
2	84.39	52.26	89.39	97.27	77.56	98.88	91.00
3	84.33	47.22	85.40	98.57	82.32	98.30	94.58
Mean	**83.68**	**45.46**	**85.28**	**97.69**	**78.63**	**98.63**	**93.39**

Table 4: Ablation study on different object detectors. s is the step stride to extend or narrow down the width/height of a context, n_c is the number of contexts in Context Mining, and m, n are the parameters of ACM.

Method	Parameter	s	mAP@0.5
Faster R-CNN	-	-	84.33
Context Mining	$n_c = 2$	2	70.52
	$n_c = 4$		81.00
	$n_c = 8$		90.49
	$n_c = 16$		92.83
	$n_c = 2$	4	77.43
	$n_c = 4$		89.04
	$n_c = 8$		92.59
	$n_c = 16$		**92.84**
Augmented CM	$m = n = 8$	4	90.53

Table 5: Person detection results of the best methods.

Method	mAP@0.5	Improvement
Faster R-CNN	47.22	-
Context Mining	93.67	**+46.45**
Augmented CM	92.44	+45.22

shown in Table 4. With appropriate hyper-parameters ($n_c = 16, s = 4$ for CM), CM and ACM significantly outperform the baseline (+8.51% for CM and +6.20% for ACM). For pedestrian detection, results in Table 5 indicate significant improvements of CM and ACM over Faster R-CNN. Between CM and ACM, CM outperforms ACM by about two points in terms of mAP@0.5. It implies that mining by small number of contexts and by extending the original regions gradually can lead to a better performance.

Runtime analysis Increasing the number of contexts results in an increase of th inference time: for Faster R-CNN, it takes 0.56 seconds for inference of a single image, while CM ($n_c = 16, s = 4$) takes 1.04 seconds and ACM ($m = n = 8, s = 4$) takes 5.81 seconds, with single GPU. Mining in a large space of contexts requires time and resources.

Table 6: Performance comparison between different accident forecasters. ToA@0.8 is the ToA when Recall is 80.0%. The results are obtained after training each models for 40 epochs like in (Chan et al. 2016).

Method	AP	mToA	ToA@0.8
DSA (Chan et al. 2016)	**47.36**	1.359	1.798
ACM+DSA	47.09	**1.457**	**2.104**

Accident forecasting

The results for accident forecasting using DSA-LSTM (Chan et al. 2016) can be found in Table 6. For a dataset with average Time-to-First-Accident (ToA) is 3.84 seconds, DSA-LSTM with Faster R-CNN features can issue warning prior to the accidents at 1.359 seconds with highest AP is 47.36%. Moreover, when recall is 80%, the ToA is 1.798 seconds. DSA-LSTM with ACM features can issue warning prior to the accidents at 1.457 seconds with AP is 47.09%. ToA@0.8 is 2.104 seconds.

Conclusion

We introduced the Car Accident Detection and Prediction (CADP) Dataset from CCTV Traffic Camera videos. A detailed account of the challenges faced in creation of the dataset such as data collection, access to traffic camera footage were tackled in the paper. We presented the results of state-of-the-art object detection and accident forecasting models on our dataset. We highlighted the strengths and weaknesses of these baseline models, and outperformed the initial results by adding context mining or augmented context mining. We finally showed that augmented context mining does not improve the score obtained with a gradual context mining for object detection. We also demonstrate the final model for accident forecasting that can predict accidents about 2 seconds before they occur with 80% recall.

References

Benenson, R.; Omran, M.; Hosang, J.; and Schiele, B. 2014. Ten Years of Pedestrian Detection, What Have We Learned? In *ECCV 2014 Workshops*, 613–627.

Chan, F. H.; Chen, Y. T.; Xiang, Y.; and Sun, M. 2016. Anticipating accidents in dashcam videos. In *Asian Conference on Computer Vision*.

Datondji, S. R. E.; Dupuis, Y.; Subirats, P.; and Vasseur, P. 2016. A Survey of Vision-Based Traffic Monitoring of Road Intersections. *IEEE Transactions on Intelligent Transportation Systems* 17(10):2681–2698.

Dollar, P.; Wojek, C.; Schiele, B.; and Perona, P. 2012. Pedestrian Detection: An Evaluation of the State of the Art. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 34(4):743–761.

Everingham, M.; Van Gool, L.; Williams, C. K. I.; Winn, J.; and Zisserman, A. The PASCAL Visual Object Classes Challenge 2012 (VOC2012) Results. http://www.pascal-network.org/challenges/VOC/voc2012/workshop/index.html.

Girshick, R.; Donahue, J.; Darrell, T.; and Malik, J. 2014. Rich Feature Hierarchies for Accurate Object Detection and Semantic Segmentation. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 580–587. IEEE.

Girshick, R. 2015. Fast R-CNN. In *IEEE International Conference on Computer Vision*.

Goodfellow, I. J.; Warde-Farley, D.; Mirza, M.; Courville, A.; and Bengio, Y. 2013. Maxout Networks. In *International Conference on Machine Learning*.

He, X., and Zeng, D. 2017. Real-time pedestrian warning system on highway using deep learning methods. In *International Symposium on Intelligent Signal Processing and Communication Systems (ISPACS)*, 701–706. IEEE.

He, K.; Zhang, X.; Ren, S.; and Sun, J. 2015. Spatial Pyramid Pooling in Deep Convolutional Networks for Visual Recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

He, K.; Gkioxari, G.; Dollar, P.; and Girshick, R. 2017. Mask R-CNN. In *IEEE International Conference on Computer Vision*, 2980–2988. IEEE.

Kong, W.; Li, N.; Li, T. H.; and Li, G. 2018. Deep Pedestrian Detection Using Contextual Information and Multi-level Features. In *International Conference on Multimedia Modeling (MMM)*, 166–177.

LeCun, Y. A.; Bengio, Y.; and Hinton, G. E. 2015. Deep learning. *Nature*.

Li, J.; Liang, X.; Shen, S.; Xu, T.; Feng, J.; and Yan, S. 2017. Scale-aware Fast R-CNN for Pedestrian Detection. *IEEE Transactions on Multimedia* 1–1.

Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft COCO: Common Objects in Context. In *European Conference on Computer Vision*, 740–755.

Lin, T. Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.

Liu, W.; Anguelov, D.; Erhan, D.; Szegedy, C.; Reed, S.; Fu, C.-Y.; and Berg, A. C. 2016. SSD: Single Shot MultiBox Detector. In *European Conference on Computer Vision*, 21–37.

Lyu, S.; Chang, M.-C.; Du, D.; Wen, L.; Qi, H.; Li, Y.; Wei, Y.; Ke, L.; Hu, T.; Del Coco, M.; et al. 2017. Ua-detrac 2017: Report of avss2017 & iwt4s challenge on advanced traffic monitoring. In *Advanced Video and Signal Based Surveillance (AVSS), 2017 14th IEEE International Conference on*, 1–7. IEEE.

Naylor, M., and Attwood, C. I. 2003. Annotated digital video for intelligent surveillance and optimized retrieval: Final report. In *ADVISOR Consortium*.

Ren, S.; He, K.; Girshick, R.; and Sun, J. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems 28*, 91–99.

Ren, Y.; Zhu, C.; and Xiao, S. 2018. Small Object Detection in Optical Remote Sensing Images via Modified Faster R-CNN. *Applied Sciences* 8.

Sultani, W.; Chen, C.; and Shah, M. 2018. Real-world Anomaly Detection in Surveillance Videos. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*.

The National Safety Council. 2017. The state of safety - a state-by-state report.

Uijlings, J.; van de Sande, K.; Gevers, T.; and Smeulders, A. 2013. Selective search for object recognition. *International Journal of Computer Vision*.

Vondrick, C.; Patterson, D.; and Ramanan, D. 2013. Efficiently Scaling up Crowdsourced Video Annotation. *International Journal of Computer Vision* 101(1):184–204.

Wang, S.; Cheng, J.; Liu, H.; Wang, F.; and Zhou, H. 2018. Pedestrian Detection via Body-part Semantic and Contextual Information with DNN. *IEEE Transactions on Multimedia* 1–1.

Williams, M. 1992. The prometheus programme. In *IEE Colloquium on Towards Safer Road Transport - Engineering Solutions*, 4/1–4/3.

Xu, K.; Ba, J.; Kiros, R.; Cho, K.; Courville, A.; Salakhudinov, R.; Zemel, R.; and Bengio, Y. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *International Conference on Machine Learning*, 2048–2057.

Zhang, L.; Lin, L.; Liang, X.; and He, K. 2016. Is faster R-CNN doing well for pedestrian detection? In *European Conference on Computer Vision*.

Zitnick, C. L., and Dollár, P. 2014. Edge Boxes: Locating Object Proposals from Edges. In *European Conference on Computer Vision*, 391–405.