

HIGH-FREQUENCY CROWD INSIGHTS FOR PUBLIC SAFETY AND CONGESTION CONTROL

Karthik. Nandakumar, Sebastien Blandin, Laura Wynter

IBM Research

Singapore

{nkarthik, sblandin, lwynter}@sg.ibm.com

ABSTRACT

We present results from several projects aimed at enabling the real-time understanding of crowds and their behaviour in the built environment. We make use of CCTV video cameras that are ubiquitous throughout the developed and developing world and as such are able to play the role of a reliable sensing mechanism. We outline the novel methods developed for our crowd insights engine, and illustrate examples of its use in different contexts in the urban landscape. Applications of the technology range from maintaining security in public spaces to quantifying the adequacy of public transport level of service.

1 INTRODUCTION

Increasing urbanization in the developing and developed world puts significant pressure on public infrastructure. High urban population leads to large crowds in public spaces which in turn can lead to a degradation of efficiency and public safety, if not well managed. The management of crowds at events and at critical infrastructures such as airports and public transport facilities is thus an important component in sustainable urban planning, especially under increasing demand pressure.

As of 2018, 33 cities in the world have population exceeding 10 million. According to United Nations (2018), this number is expected to grow to 43 in 2030. This urban population growth affects on a daily basis the public transport infrastructure. In Singapore, for example, average daily ridership increased from 5.01 million in 2002 to 7.67 million in 2014, see Land Transport Authority (2019). In Hong-Kong, the monthly patronage increased from 66.07 million in January 2000 to 148.22 million in January 2019, see Mass Transit Railway (2019). Large scale gatherings and events are also affected, seeing increasing crowds and increasing chances of crowd-related incidents. An extreme example occurred in 2015 during the annual Muslim pilgrimage near the holy city of Mecca, in Saudi Arabia, when thousands of pedestrians were trampled.

The challenges of crowd monitoring and management in the above settings are significant, Blandin et al. (2019). There is a dramatic variation in usage, requiring fast feedback during times of large crowds and incidents. Crowding events must be precisely estimated in time, space, and magnitude as well as their nature (calm vs. angry), in order to support the following feedback mechanisms; (i) information to individuals within and outside the crowd, (ii) local prescriptive enforcement via ground personnel, (iii) regional adjustment of resources to manage the event. As the risk of incidents increases with the density of the crowd, along with crowd-level analysis, the precursors of such incidents are required so as to enable rapid mitigation.

This requirement on precision in both the nature of the behavior and the time-frame of changes to those parameters means that many technologies such as GPS-based crowd analysis are less effective than video, the latter enabling analysis of the full population rather than an unknown subsample, as well as important features such as movement and emotion. Much of video analysis however aims to account for each person individually; however, in crowded spaces, bodies and faces overlap.

Many of the requirements and challenges were detailed two decades ago in the classic reference Davies et al. (1995). Helbing et al. (2000) studied crowd-related disasters and emphasized the need for monitoring of crowd behavior in densely-populated spaces. Mathematical aspects of crowd modeling can be found in Cristiani et al. (2014).

We provide details of our system that both accurately estimates crowd flow and density and detects events of interest in the crowds, using existing CCTV cameras. Our approach has the following advantages; (i) the infrastructure to sense, collect and retrieve data exists across cities in both the developing and developed world, and (ii) accuracy is limited mainly by the algorithm, rather than the infrastructure, and can be improved as computing power increases, unlike geo-localized signals such as GPS or indoor positioning technologies.

2 CCTV-BASED CROWD MONITORING

The proposed system consists of two main modules: (i) Crowd flow and density estimation, and (ii) crowd anomaly detection, and uses both static features extracted from individual frames of the video and spatio-temporal features extracted from contiguous sequence of frames.

2.1 CROWD FLOW AND DENSITY ESTIMATION

Our crowd level estimation module uses adaptive fusion combining people detection with crowd-level image-feature based regression. The counting-by-detection approach scans the image frame for individual pedestrians and works well for sparsely crowded scenes (Figure 1 (a)), where there is less occlusion. Regression-based counting directly estimates the people count in densely crowded scenes, where occlusion impedes detection of individuals (Figure 1 (b)).



Figure 1: People counting by detection for (a) sparse crowd and (b) dense crowd. Occlusion leads to severe underestimation in (b). These images are extracted from the ShanghaiTech dataset, Zhang et al. (2016). Adaptive fusion used for video-based people counting, (c) foreground segmentation and (d) pedestrian detection. The fusion model combines the people detections with the regression method after de-emphasizing the regions with pedestrian detection (marked by yellow boxes in (c)).

People Detection and Flow: The people detector, appropriate for low crowd levels, detects features of individuals. Our algorithm is based on full-body, head and shoulder, and head characteristics. To improve robustness to occlusion, the algorithm extracts classical features such as histogram of Gradients (HoG), gradient magnitude, and texture, over multiple channels. Boosted decision trees are used as the classifier. Extending the algorithm via tracking of pedestrian trajectories, we obtain estimates of the people flow. Filtering is then performed to remove trajectories that are either very short or do not contain any strong directional movement. Finally, a two-class classifier is applied to categorize the remaining trajectories into incoming and outgoing pedestrian flows.

Regression-based Counting: This component first extracts simple features (e.g., shape and texture) from the image, weighted using a perspective normalization. We divide the image into 12 horizontal blocks and compute the ratio of foreground and edge pixels in each block. Thus, we obtain a 24-dimensional feature vector, which is used as the input to a linear regression algorithm to obtain an estimate of the crowd density. When computational resources are sufficient, we use deep learning for feature extraction. Specifically, we use the basic AlexNet convolutional neural network Krizhevsky et al. (2012) and the 4,096-dimensional embedding obtained at the output of the last fully convolutional layer. A support vector regression method is used to obtain the crowd density estimate.

Adaptive Fusion: Crowd levels at a specific location are highly variable and relying on a single approach leads to large errors when the conditions change. We thus developed an adaptive fusion scheme that robustly combines the above approaches to improve the accuracy of crowd density estimates across all conditions. An illustration is shown in Figures 1 (c) and (d). The key element

that leads to its success is to extract features required for crowd-based regression only from regions where there are no pedestrian detections. As such, the method automatically assigns a higher importance to the detection-based count in sparsely crowded scenes and to the regression-based count in densely crowded scenes. We set a stringent threshold to minimize the number of false detections.

2.2 CROWD INCIDENT DETECTION

We consider crowd incidents to be characterized by a rapid movement of people towards each other with abrupt directional changes. The following two flow features are thus crucial. In particular, in order to reduce false positives, we discard features directly extracted from individual trajectories.

Aggregate flow patterns: Assuming that the background is stationary and motion is primarily restricted to the people in the scene, optical flow-based features characterize crowd movement patterns at a coarse level. Figure 2 shows examples of aggregate flow-based features extracted from three scenarios (a fight, a group of people dancing, and dense moving crowd without any incident). While it may be easy to filter out the third scenario, distinguishing between the first two scenarios based on this set of features is challenging as both exhibit rapid and abrupt motions.

The image is divided into many smaller blocks, histogram of the aggregate flow field is obtained for each block, and these histograms are concatenated to obtain the global feature vector, which is finally input to a Support Vector Machine (SVM) classifier to detect fight/violence.

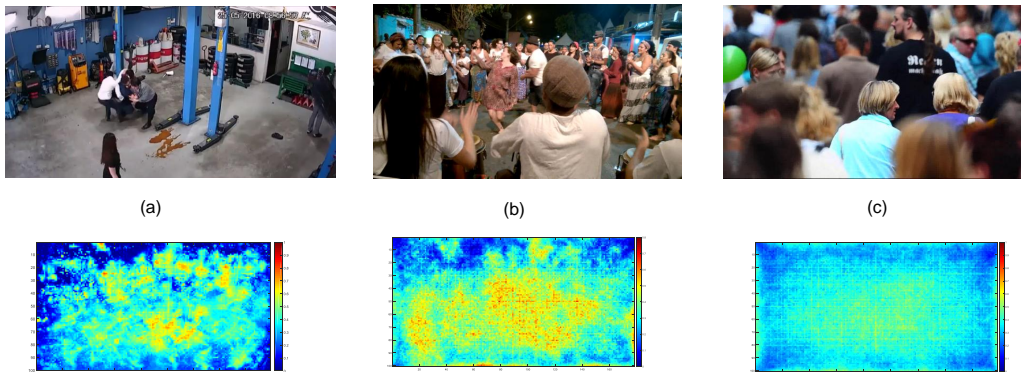


Figure 2: Flow-based features extracted from three different scenarios. (a) fight, (b) group of people dancing, and (c) dense moving crowd. Dark blue pixels indicate regions with low activity and the red pixels denote regions with rapid motion.

Features extracted from selected interest points: A subset of interest points is automatically selected (unlike the flow-estimator where the motion of all pixels is considered), and these specific interest points are tracked for short intervals (typically in the order of 0-5 seconds), generating what are known as tracklets. Spatio-temporal features are then extracted from tracklets.

Three types of crowd interaction patterns have been proposed in Shao et al. (2015). While collectiveness represents the degree to which individuals move together, stability indicates the degree of change in the topological crowd structure, and conflict occurs when individuals move towards different directions. Collectiveness and conflict patterns are especially useful for abnormal crowd behavior detection. These patterns can be extracted based on the velocity of the tracklets of selected interest points. We first segment the foreground region based on motion features and extract interest points only in the foreground region. The velocities of the individual tracklets are used for estimating collectiveness, conflict, and mean speed. An SVM classifier is applied to detect violence/fight based on these features.

Finally, the results of the two SVM classifiers are combined at the score level to derive an overall confidence measure (between 0 and 1) of abnormal crowd event.

3 RESULTS

3.1 IMAGE ANNOTATION

We annotated images from 3 urban train stations, indoor and outdoor, at 4 locations (platform, gateway, lobby, escalator) during periods of heavy crowds on 2 weekdays; one day includes a train incident during which the entire platform is fully crowded, for 420 minutes of video overall.

For each frame, we annotate each full body, head & shoulders, and head, occupying at least 30 pixels on the cameras. We observed that below that threshold, human annotation was not reliable. The native image resolution is 704×576 . We eventually create 63,114 boxes over 4,200 frames for low crowd scenarios, and 202,283 boxes over 6,527 frames for high crowd scenarios.

3.2 CROWD FLOW AND DENSITY ESTIMATION

We use 90% of the data for training and 10% for testing. Results are provided in Table 1.

Ground truth People count	Low crowd		High crowd					
	[0, 10]	[11, 25]	[26, 50]	[51, 75]				
Absolute error	Occurrences per crowd level (%)							
[0, 5]	85	100	71	99	71	92	61	91
[6, 10]	15	0	26	1	22	8	26	9
[11, 15]	0	0	3	0	7	0	9	0
[16, 20]	0	0	0	0	0	0	4	0

Table 1: Confusion matrix for absolute error per crowd level. The people detection model is used for low crowd, while the regression-based counting is used for high crowd. Results from the fusion model are indicated in bold.

While the error increases with the crowd level, the fusion model significantly improves upon the base models. These results also improve on the performance from off-the-shelf software that we tested, by 10% to 20%. In our applications, users were typically interested in count estimates within 5 persons of the ground-truth counts for low crowds, and within 25% of the ground-truth for dense crowds. The work presented in this article demonstrated the feasibility of estimating crowd level with high accuracy in practical urban monitoring contexts.

3.3 CROWD INCIDENT DETECTION

The incident detection module was evaluated on the Violent Flows dataset, see Hassner et al. (2012), which consists of 246 videos from YouTube. We followed the same protocol (with five-fold cross validation) as Hassner et al. (2012) for benchmarking. By varying the threshold for confidence value, the area under the receiver operating characteristic curve (AUC) of our system was found to be 87.9, improving on the 85 result from Hassner et al. (2012).

Since the Violent Flows database focuses only on crowd violence, we also collected a small test set of 10 videos, which were independently downloaded from YouTube. Among these 10 videos, 7 samples depict fights among individuals or small groups, while the remaining 3 samples depict other activities with rapid motion such as group dancing or playing games.

In all the 7 fight samples, our method detected an incident with a confidence value of greater than 0.4. For the remaining 3 samples, the confidence value of an incident did not exceed 0.2. The average time lag between the outbreak of violence and its detection was found to be 10 seconds.

4 CONCLUSION

This short article presented an overview of a set of video-based techniques designed for real-world crowd insights and the management of dense crowds in urban environments, and described results from actual deployments.

REFERENCES

- Sebastien Blandin, Laura Wynter, Hasan Poonawala, Sean Laguna, and Basile Dura. FASTER: Fusion AnalyticS for public Transport Event Response. In *Proceedings of the 18th Conference on Autonomous Agents and MultiAgent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- Emiliano Cristiani, Benedetto Piccoli, and Andrea Tosin. *Multiscale modeling of pedestrian dynamics*, volume 12. Springer, 2014.
- Anthony C Davies, Jia Hong Yin, and Sergio A Velastin. Crowd monitoring using image processing. *Electronics & Communication Engineering Journal*, 7(1):37–47, 1995.
- Tal Hassner, Yossi Itcher, and Orit Kliper-Gross. Violent flows: Real-time detection of violent crowd behavior. In *2012 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1–6. IEEE, 2012.
- Dirk Helbing, Illés Farkas, and Tamas Vicsek. Simulating dynamical features of escape panic. *Nature*, 407(6803):487, 2000.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Land Transport Authority. Singapore land transport statistics in brief. <https://www.lta.gov.sg/content/ltaweb/en/publications-and-research.html>, 2019.
- Mass Transit Railway. Mtr patronage figures. <http://www.mtr.com.hk/en/corporate/investor/patronage.php>, 2019.
- Jing Shao, Kai Kang, Chen Change Loy, and Xiaogang Wang. Deeply learned attributes for crowded scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4657–4666, 2015.
- United Nations. The world’s cities in 2018. https://www.un.org/en/events/citiesday/assets/pdf/the_worlds_cities_in_2018_data_booklet.pdf, 2018.
- Yingying Zhang, Desen Zhou, Siqin Chen, Shenghua Gao, and Yi Ma. Single-image crowd counting via multi-column convolutional neural network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 589–597, 2016.