Auto-SDA: Automated Video-based Social Distancing Analyzer

Mahshid Ghasemi
Columbia University
mahshid.ghasemi@columbia.edu

1. INTRODUCTION

Social distancing can reduce infection rates in respiratory pandemics such as COVID-19, especially in dense urban areas. To assess pedestrians' compliance with social distancing policies, we use the pilot site of the PAWR COSMOS wireless edge-cloud testbed in New York City to design and evaluate an **Auto**mated video-based **S**ocial **D**istancing **A**nalyzer (**Auto-SDA**) pipeline (shown in Fig. 1). This pipeline measures the distance between unaffiliated pedestrians (i.e., the pedestrians who do not walk together as a social group) and assesses if they maintain 6 ft distance.

The main goal of this work is to design a highly accurate social distancing analyzer pipeline, whose *performance is insensitive to the camera's viewpoint and scene dynamics*. As illustrated in Fig. 1, the pipeline includes an object detector model (YOLOv4 [4]) and a tracker model (Nvidia DCF-based tracker) that extracts the trajectory of pedestrians. These trajectories are eventually used to compute the proximity duration of each two unaffiliated pedestrians separately. While these are off-the-shelf components, achieving the goal, mentioned above, calls for the design of tailored components. Specifically, we achieved this goal by incorporating three modules in Auto-SDA, as outlined below and in Section 2:

- **Camera calibration module:** Calculates the on-ground distances between pedestrians with less than 10 cm error.
- **ID correction module:** Compensates for the inaccuracies of the object detector and tracking model caused by the camera's tilt angle and the obstacles on the road.
- **Group detection module:** Detects the pedestrians affiliated with a single social group.

We applied Auto-SDA to our dataset collected by a camera located on the 2$^{\text{nd}}$ floor of Columbia's Mudd building. It consists of videos recorded before the COVID-19 outbreak, soon after the lockdown, and after the vaccines became broadly available (see Table 2 for a summarized comparison between Auto-SDA and the previously proposed social distancing frameworks). The detailed results (described in Section 3) show, for example, that after the lockdown, the density of pedestrians seen at the intersection decreased by almost 50%. Moreover, after the lockdown, less than 55% of the pedestrians violated the social distancing protocols compared to 65% post-vaccine. The results also show that the fraction of pedestrians walking as a social group has grown from 0-20% (after the lockdown) to 10-45% (post-vaccine).

2. PIPELINE MODULES

2.1 Camera Calibration

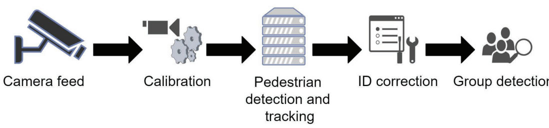

Figure 1: Different stages in the Auto-SDA pipeline.

The goal of the camera calibration module is to determine the intrinsic and extrinsic parameters of the camera to convert the 2D on-image coordinates, viewed by the camera, to the 3D on-ground coordinates. As part of this process, we used a checkerboard to compute the intrinsic parameters of the camera using OpenCV [5].

Moreover, we split the view of the intersection into 10 areas and for each area, using OpenCV, we determined the extrinsic parameters individually. This can further mitigate the impact of camera distortion and obtain the on-ground distances with less than 10 cm error. Auto-SDA uses these parameters to calculate the on-ground distance between the pedestrians by solving the corresponding photogrammetry equations (see [5, 9, 1]).

2.2 Pedestrians Detection and Tracking

Auto-SDA uses the YOLOv4 object detector to detect the pedestrians. It is also equipped with a tracker (NvDCF) that extracts the trajectory of each pedestrian and uses that to trace the number of pedestrians they are in contact with and the duration of each contact. Both models are set as building blocks inside the Deepstream pipeline which is an optimized architecture built using the Gstreamer framework [11].

2.3 Tracking IDs Correction

Due to the high altitude of the COSMOS cameras, their oblique view of the intersection, and the obstacles on the road, the object detector and tracker have degraded performance (i.e., it is likely that the tracker loses a pedestrian along the way or assigns multiple IDs to a single person).

The ID correction module detects the IDs that belong to a single pedestrian and extracts their entire trajectory. For each ID pair (id_1, id_2), the ID Correction algorithm verifies three conditions to determine whether they are associated with a single pedestrian or not. First, the gap between id_1 lost-time (t_1) and id_2 detected-time (t_2) must be less than a predefined threshold (e_1). Second, the distance between predicted location of id_1, at the time that id_2 was first detected (t_2) using the Linear Regression approximation for the tale of id_1 trajectory, and the location of id_2 at that time (t_2) has to be less than a specified threshold (e_2). Third, the angle

Table 2: A comparison of prior work to Auto-SDA

Framework	Object Detection	Tracking	Calibration Method	On-Ground Distance Computation Error	Correction	Group Detection	Real-World COVID-19 Pandemic Impact Analysis
[15]	✓	X	Homography trans.	$\gg 10$ cm	X	X	X
[12]	✓	✓	Pairwise L_2 norm	$\gg 10$ cm	X	X	X
[6, 3, 2, 7]	✓	X	Planar camera persp. trans.	$\gg 10$ cm	X	X	X
Auto-SDA	✓	✓	Multi-area calibration	< 10 cm	✓	✓	✓

between id_1's tail direction and id_2 head direction must be less than 90. If all three conditions hold, then id_1 and id_2 belong to a single person.

2.4 Group Detection

We enhance the social distancing analysis by excluding the pedestrians walking together as a social group (e.g., friends/family) from social distancing violation. There are several methods proposed for group-detection, e.g., see [8, 10, 14, 13]. All these group-detection methods require details such as velocity, body and head orientation, and exact trajectory. However, in our setting (and in many realistic deployments), because of the high altitude and oblique view of the cameras, that kind of detailed information cannot be obtained.

Therefore, we designed a group detection algorithm that can detect social groups with the limited data that we can derive from cameras such as the ones in the COSMOS pilot site. The Group Detection algorithm calculates the distance between each pair of pedestrians (id_1, id_2) on all the frames and then calculates the average distance $(\bar{d})$ and empirical standard deviation (σ). Two pedestrians are labeled as one social group, if their instantaneous distance (d) does not exceed $d_{\max}$ in more than $N_{\max}$ frames, and the mean and standard deviation of their distance are less than $\bar{d}_{\max}$ and $\sigma_{\max}$, respectively.

3. MEASUREMENTS AND EVALUATION

We applied Auto-SDA to videos recorded from a camera deployed on the 2$^{\text{nd}}$ floor of the Columbia Mudd building at the COSMOS pilot site. The camera is configured to record 180 sec (which is two times the signal timing cycle of the traffic lights at the intersection) videos, 5 times a day at 9 AM, 2 PM, 5:30 PM, 7:30 PM, and 10 PM. We deployed Auto-SDA in one of the COSMOS edge servers and applied it to the videos recorded between June 17 and July 20, 2020 (after the lockdown), and during May 2021 (after the vaccines became broadly available). We also used 16 sample videos recorded before the COVID-19 outbreak (in June 2019) to evaluate the impact of the pandemic on pedestrians' density. Below, we provide the analysis results (results corresponding to June-July, 2020 and May 2021 are labeled as *Pandemic* and *Post-vaccine*, respectively).

For each video, we calculated the percentage of pedestrians who violate social distancing and plot a normalized histogram of the results in Fig. 2. It can be seen that, after the lockdown, typically, less than 55% of the pedestrians violated social distancing, compared to 65% post-vaccine. Fig. 3 shows the fraction of recorded videos in which a certain percentage of pedestrians are walking as a group. One can see that the fraction of pedestrians walking as a social

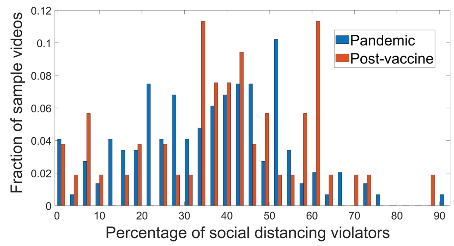

Figure 2: Normalized histogram of the percentage of pedestrians considered social distancing violators in the recorded videos.

group has grown from 0-20% (after the lockdown) to 10-45% (post-vaccine).

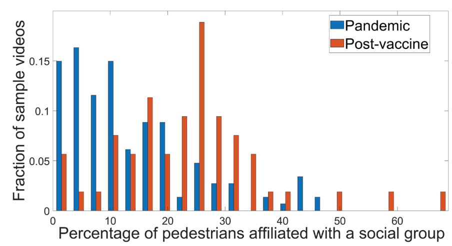

Figure 3: Normalized histogram of the percentage of pedestrians affiliated with a social group.

We compare the pre-pandemic, lockdown, and post-vaccine density of the crowd at the intersection in Fig. 4. One can observe that density of the pedestrians has decreased by almost 50% after the lockdown (compared to pre-pandemic), while it has slightly increased recently.

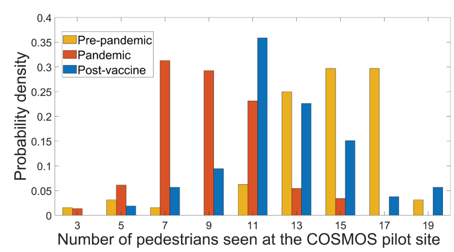

Figure 4: Comparison between the density of pedestrians walking at the COSMOS pilot site in different periods.

4. CONCLUSIONS

We presented the Auto-SDA pipeline that evaluates if unaffiliated pedestrians comply with the social distancing policies. It is equipped with a group detection module and is capable of calculating the on-ground distance between pedestrians with less than 10 cm error. We applied Auto-SDA to the videos recorded by a camera deployed at the COSMOS pilot site. The results demonstrate the impact of the COVID pandemic on pedestrians' behaviors.

References

[1] Peyman Alizadeh. "Object distance measurement using a single camera for robotic applications". PhD thesis. Laurentian University of Sudbury, 2015.

[2] John Betancourt. *Social Distancing Analyser*. `https://github.com/JohnBetaCode/Social-Distancing-Analyser`. 2020.

[3] Deepak Birla. *Social Distancing AI*. `https://github.com/deepak112/Social-Distancing-AI`. 2020.

[4] Alexey Bochkovskiy, Chien-Yao Wang, and Hong-Yuan Mark Liao. "YOLOv4: Optimal speed and accuracy of object detection". In: *arXiv preprint arXiv:2004.10934* (2020).

[5] Gary Bradski. "The OpenCV Library". In: *Dr. Dobb's Journal of Software Tools* (2000).

[6] Marco Cristani et al. "The visual social distancing problem". In: *arXiv preprint arXiv:2005.04813* (2020).

[7] Tom Farrand. *Social Distancing*. `https://github.com/FarrandTom/social-distancing`. 2020.

[8] Shoichi Inaba and Yoshimitsu Aoki. "Conversational group detection based on social context using graph clustering algorithm". In: *2016 12th Int. Conf. Signal-Image Technol. Internet-Based Sys. (SITIS)*. 2016.

[9] Matthias Jüngel, Heinrich Mellmann, and Michael Spranger. "Improving vision-based distance measurements using reference objects". In: *Robot Soccer World Cup*. 2007.

[10] Riccardo Mazzon, Fabio Poiesi, and Andrea Cavallaro. "Detection and tracking of groups in crowd". In: *2013 10th IEEE Int. Conf. Adv. Video Signal Based Surveill.* 2013.

[11] John Nickolls et al. "Scalable parallel programming with CUDA". In: *Queue* 6.2 (2008), pp. 40–53.

[12] Narinder Singh Punn, Sanjay Kumar Sonbhadra, and Sonali Agarwal. "Monitoring COVID-19 social distancing with person detection and tracking via fine-tuned YOLOv3 and DeepSort techniques". In: *arXiv preprint arXiv:2005.01385* (2020).

[13] Francesco Solera, Simone Calderara, and Rita Cucchiara. "Socially constrained structural learning for groups detection in crowd". In: *IEEE Trans. Pattern Anal. Mach. Intell.* 38.5 (2015), pp. 995–1008.

[14] Sebastiano Vascon and Loris Bazzani. "Group detection and tracking using sociological features". In: *Group and crowd behavior for computer vision*. Elsevier, 2017, pp. 29–66.

[15] Dongfang Yang et al. "A vision-based social distancing and critical density detection system for COVID-19". In: *arXiv preprint arXiv:2007.03578* (2020), pp. 24–25.