Tracking Passengers and Baggage Items using Multiple Overhead Cameras at Security Checkpoints

Abubakar Siddique, Student Member, IEEE, Henry Medeiros, Senior Member, IEEE*†\$¶

Abstract—We introduce a novel framework to track multiple objects in overhead camera videos for airport checkpoint security scenarios where targets correspond to passengers and their baggage items. We propose a Self-Supervised Learning (SSL) technique to provide the model information about instance segmentation uncertainty from overhead images. Our SSL approach improves object detection by employing a test-time data augmentation and a regression-based, rotation-invariant pseudolabel refinement technique. Our pseudo-label generation method provides multiple geometrically-transformed images as inputs to a Convolutional Neural Network (CNN), regresses the augmented detections generated by the network to reduce localization errors, and then clusters them using the mean-shift algorithm. The selfsupervised detector model is used in a single-camera tracking algorithm to generate temporal identifiers for the targets. Our method also incorporates a multi-view trajectory association mechanism to maintain consistent temporal identifiers as passengers travel across camera views. An evaluation of detection, tracking, and association performances on videos obtained from multiple overhead cameras in a realistic airport checkpoint environment demonstrates the effectiveness of the proposed approach. Our results show that self-supervision improves object detection accuracy by up to 42% without increasing the inference time of the model. Our multi-camera association method achieves up to 89% multi-object tracking accuracy with an average computation time of less than 15 ms.

Index Terms—Self-supervised Learning, Detection, Tracking, Tracklet Association, Multi-camera Tracking, Surveillance.

I. INTRODUCTION

UTOMATED video surveillance requires the detection, tracking, and recognition of objects of interest in a scene. Accurate and precise surveillance in crowded scenes is one of the most challenging computer vision applications. To address the problem of visual surveillance in the domain of airport checkpoint security, the Department of Homeland Security

*Manuscript received August 22, 2022; accepted November 11, 2022. Date of publication December 14, 2022.

[†]This material is based upon work supported by the U.S. Department of Homeland Security, Science and Technology Directorate, Office of University Programs, under Award Number 2013-ST-061-E0001-04. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

[‡]Abubakar Siddique is with the Department of Electrical and Computer Engineering, Marquette University, Milwaukee, USA, e-mail: abubakar.siddique@marquette.edu

[§]Henry Medeiros is with the Department of Agricultural and Biological Engineering, University of Florida, Gainesville, USA, e-mail: hmedeiros@ufl.edu

[¶]Digital Object Identifier (DOI): 10.1109/TSMC.2022.3225252

(DHS) ALERT (Awareness and Localization of Explosives-Related Threats) center of excellence at Northeastern University initiated the CLASP (Correlating Luggage and Specific Passengers) project. This initiative aims to help the Transportation Security Administration (TSA) detect security incidents, such as theft of items and abandoned bags.

Current approaches for detecting and tracking passengers and luggage in airport checkpoints divide the image area within each camera's field of view into regions of interest where certain passenger behaviors are expected (e.g., passengers divest their items near the roller conveyor) [1], [2]. While these approaches are effective within individual regions of interest, they cannot detect and track passengers and their belongings throughout an entire checkpoint. Moreover, most recent detection algorithms [3]–[6] are unable to detect multiple objects in realistic overhead camera scenarios due to the unavailability of large-scale datasets obtained using unconventional camera perspectives.

Fine-tuning pre-trained models using human annotated labels is a common approach in computer vision methods. However, this strategy hinders the applicability of state-ofthe-art algorithms in scenarios where images are obtained from perspectives that are not commonly observed in existing publicly available datasets. The dramatic variability of video surveillance systems used in airport checkpoints would require deployment-specific fine-tuning of models, and in some scenarios, even camera-specific adjustments. To overcome this challenge, we leverage the fact that models pre-trained on large-scale datasets can build upon their initial predictions to adapt to new scenarios using SSL strategies. Our proposed SSL framework obviates the tedious and expensive human annotation procedure by automatically generating pseudolabels to update the model.

To generate pseudo-labels, we cluster multiple detections obtained from geometrically transformed images using the mean-shift algorithm [7]. Each cluster corresponds to the detection of one object observed at different orientations on several augmented input images. The cluster modes with the corresponding bounding boxes, segmentation masks, and confidence scores are used to update the model. Thus, our model learns from rotation-invariant pseudo-labels and can be integrated with a tracking-by-detection algorithm [8] to generate accurate target tracklets from overhead perspectives.

Our SSL algorithm is inspired by the methods described in [9]–[13]. However, unlike [9], instead of resorting to multi-task



Fig. 1. Proposed SSL framework. The augmented proposal generation stage uses multiple rotated versions of the unlabeled input images to generate augmented detections from an instance segmentation model and then remaps these predictions into their original coordinates. The clustering algorithm leverages the model's regression ability to reduce localization errors using the augmented predictions as region proposals. The regressed cluster modes are then used to generate augmented pseudo-labels to update the model.

strategies to guide the learning process, we employ a multiinference approach similar in spirit to the self-consistency method based on equivariant transformations proposed in [10]. Our method differs from [10] in that, rather than using the uncertainties from multiple model predictions to select image patches for additional training, it aggregates multiple inferences into accurate pseudo-labels that are used to refine the model. Our method departs significantly from unsupervised model adaptation [11] and knowledge distillation approaches [14] in that we only use automatically generated labels and avoid human annotations altogether during model update.

We also propose a Multi-Camera Tracklet Association (MCTA) algorithm to maintain the temporal identifiers of passengers across cameras. We leverage the fact that our system is comprised of overhead cameras with partially overlapping fields of view to employ a simple but effective geometry-based trajectory association method. Our algorithm compares the projected centroids of target detections on neighboring cameras using the homographies between their image planes. We track passengers and bags across multiple views and generate global tracks by combining pairwise associations from the partially overlapping camera views.

We evaluate the detection and tracking performance of our algorithms on videos from a simulated airport checkpoint and demonstrate that our approach performs on par with a model trained in an entirely supervised manner and substantially outperforms the pre-trained detection model. Our multi-camera evaluation shows that our MCTA method effectively handles the problem of passenger identity hand-off across cameras.

In summary, the key contributions of this work are:

- A novel self-supervised object detection algorithm that generates pseudo-labels based on instance segmentation uncertainties.
- A new data augmentation and regression-based clustering mechanism that substantially improves the quality of pseudo-labels for self-supervised training.
- A new recursive tracklet association algorithm to address the identity hand-off issue during transitions between

crowded overhead camera views.

- We provide an extensive evaluation of our methods on a dataset collected using multiple overhead cameras in a realistic airport checkpoint scenario.
- Our SSL models and the corresponding source code are available at https://github.com/siddiquemu/SCT_MCTA.

To our knowledge, this is the first approach to solve the overhead multi-view association problem in a network of cameras with partially overlapping fields of view using a selfsupervised detection strategy.

II. RELATED WORK

Multiple target tracking using camera networks is an active research topic with several potential applications [15]–[18]. Most works on camera networks focus on the multi-camera aspect of the problem and do not consider the challenges associated with camera perspectives. Although generic object tracking algorithms could be used in surveillance systems (e.g. [19]–[21]), when object categories are known, trackers based on specialized detectors are more accurate and less prone to model drift [22], [23]. This observation has led to the development of a variety of multiple target tracking algorithms that specialize in tracking humans [24]–[34] or vehicles [35]– [38]. However, in many scenarios, it is desirable to track additional objects of known categories. In these cases, more flexible detection algorithms are needed, but the effectiveness of modern object detection models is highly dependent upon the characteristics of the training datasets [3]-[5].

Previous works have used SSL techniques to improve visual feature learning [39]–[41], reducing dependency on human annotations for training backbone models. However, transferring knowledge from pre-trained backbones to downstream tasks is a far less explored topic. Unlike our proposed approach, SSL techniques for detection [9], [11] and semantic segmentation [10] rely on annotations to initialize the model before iterative learning can take place.

Data augmentation is an effective mechanism to improve the robustness of CNNs in scenarios not available during



Fig. 2. Visualization of our data augmentation approach. The first and second columns show the segmentation masks and detections at $\theta = 0^{\circ}$ and $\theta = 186^{\circ}$, respectively. The third column shows the remapped detections in the set $S^{\mathcal{C}}$ on the original image (using Alg. 1) with the best detections (blue) from Alg. 2.

training [14], [42], but little attention has been given so far to approaches for combining the response of the network to augmented samples. In multi-target tracking applications, multiple detections mapped to a common coordinate system can be interpreted as the probability of occupancy of the area observed by the cameras [43]. Although it is possible to use clustering techniques to map the modes of this distribution to unique target detections, bounding box alignment errors pose a challenge to the generation of high-quality pseudo-labels for SSL. Hence, we propose a test-time regression technique that leverages instance segmentation information for pseudo-label generation.

A systematic solution to the data association problem is another important component of multi-target tracking-bydetection methods [44]–[57]. Single-camera trackers [8], [58] use detectors trained on multiple datasets [59] to generate bounding boxes and form track hypotheses for all the targets in each frame. In this work, we employ a state-of-the-art single-camera tracker [8] using a detector based on our selfsupervised models, which achieves unprecedented tracking performance in previously unseen airport surveillance videos.

Finally, multi-camera tracking systems require sophisticated trajectory association mechanisms to maintain target identities across cameras [60]–[62]. Even within a single camera, occlusions must be handled using similar strategies [63], [64]. Most association approaches compute trajectory similarity scores based on a combination of appearance and motion features [60]–[64]. These features are learned using a large number of

continuous trajectories, which are difficult to obtain with typical ceiling-height overhead cameras due to their limited fields of view. Some methods use camera calibration information to project tracks onto a common plane and perform association using occlusion modeling [32] or re-identification techniques [33]–[37]. Dependency on camera calibration further limits the applicability of these methods to security systems since calibrating multiple cameras with partially overlapping fields of view is a complex task [65]–[68].

III. PROPOSED MODEL

Our system consists of two main components: i) a detection algorithm trained using SSL and ii) a multi-camera trackingby-detection mechanism. A single-camera tracking algorithm uses SSL detections to generate tracklets for passengers and baggage items. We then employ a novel multi-camera target trajectory association algorithm to uniquely identify passengers throughout the checkpoint.

A. Self-Supervised Learning

We use the PANet model [5] with a ResNet-50 backbone [69], [70] as the baseline detector. Since the categories of interest are persons and their belongings, we use a model pre-trained on the COCO dataset [71], which includes object classes related to these categories (i.e., person, handbag, backpack, and suitcase). Because the COCO dataset consists mostly of images captured at roughly eye-level, detectors trained using that dataset do not perform well on overhead perspectives.

To address this limitation, our SSL framework updates the baseline model using rotation-invariant pseudo-labels. As Fig. 1 shows, our SSL framework consists of three main steps: i) augmented region proposals generation, ii) pseudo-label generation and refinement through cluster regression, and iii) iterative model update.

Alg	orithm 1 Augmented Proposals Generation
1:	function AugmentedProposals $(I(t), r)$
2:	$S^{\mathcal{C}}(t) = \emptyset, \Theta = \{i \cdot \Delta \theta\}_{i=1}^{r}$
3:	for $ heta_i\in\Theta$ do
4:	$\Psi_{\theta_i}(t) = R_{\theta_i}(I(t))$
5:	$D_{\theta_i}^{\mathcal{C}}(t) = D_{PANet}(\Psi_{\theta_i}(t))$
6:	$S_{\theta_i}^{\mathcal{C}}(t) = R_{-\theta_i}(D_{\theta_i}^{\mathcal{C}}(t))$
7:	$S^{\mathcal{C}}(t) = S^{\mathcal{C}}(t) \cup S^{\mathcal{C}}_{\theta}(t)$
8:	end for
9:	return $S^C(t)$
10:	end function

1) Augmented Proposals Generation: Our data augmentation method, summarized in Alg. 1, uses the PANet model to detect and segment multiple instances of objects of interest. During the first iteration of SSL training, we retain only the outputs of the pre-trained model for the person, handbag, backpack, and suitcase classes. The person class corresponds to passengers and detections of handbag, backpack, and suitcase items are treated as baggage items. In subsequent iterations of SSL training, we modify the model to generate only the object categories $C \in \{pax, bag\}$, where pax corresponds to passengers and bag to baggage items. Let $D^{\mathcal{C}}(t)$ be the set of detections on image I(t) at time t. That is, $D^{\mathcal{C}}(t) = \{d_1, \dots, d_{n_r^c}\}$, where $d_j \in \mathbb{R}^5$ is the detection of the *j*-th object and $n_t^{\mathcal{C}}$ is the number of objects of class \mathcal{C} in frame I(t). Each detection d_i consists of the coordinates and dimensions of the target's bounding box, $b_i^{\mathcal{C}} \in \mathbb{R}^4$, as well as its detection confidence score $s_i \in [0, 1]$.

We noticed that the detector performs better when objects are observed at more commonly occurring angles (e.g., upright). Therefore, to reduce the negative effect of the overhead perspective, we generate multiple rotated copies of the input image $\Psi_{\theta_i}(t) = R_{\theta_i}(I(t))$ (line 4 in Alg. 1), where $R_{\theta_i}(\cdot)$ is the rotation operator, which rotates the image by an angle θ_i . The angle of rotation θ_i varies between 0 and 2π at intervals of $\Delta \theta = \left| \frac{2\pi}{r} \right|$, i.e., $\theta_i = \Delta \theta, \ldots, 2\pi$, where r determines



Fig. 3. Regression on test-time augmented bounding boxes (middle) and cluster modes (right) to generate pseudo-labels for SSL training.



Fig. 4. Probability of occupancy of passengers at one frame of our evaluation datasets (Fig. 2).

the rotation resolution. At each rotation step, we compute the detection set $D_{\theta_i}^{\mathcal{C}}(t)$ for both classes $\mathcal{C} \in \{pax, bag\}$ using a single call to the function $D_{\text{PANet}}(\cdot)$ (line 5). We then remap the resulting detections to the coordinate frame of the original image by applying the inverse rotation to each of the detections in $D_{\theta_i}^{\mathcal{C}}(t)$ (line 6). To avoid localization errors introduced by rotating axis-aligned bounding boxes, we apply the rotation operation to the binary segmentation masks produced by PANet and compute the corresponding bounding boxes using the rotated masks. At the end of Alg. 1, the set $S^{C}(t) = \bigcup_{i=1}^{r} S^{\mathcal{C}}_{\theta_{i}}(t)$ contains the detections at all the rotation angles θ_i . Fig. 2 illustrates the detections at two rotation angles and the result of mapping detections at 20 different orientations back to the original coordinate system.

Algorithm 2 Cluster Regression

1: function CLUSTERREGRESSION($S^{\overline{C}}(t)$)

 $\mathcal{D}^{\mathcal{C}}(t) = \emptyset$ 2: Refine the augmented detections using $S^{\mathcal{C}}(t)$ as region 3: proposals for the D_{PANet} model $O^{\mathcal{C}}(t) = \text{mean} - \text{shift}(S^{\mathcal{C}}(t))$ 4: for $Q \in O^{\mathcal{C}}(t)$ do 5: Compute the cluster score $\bar{\eta}_Q$ using Eq. 2

6:

- 7: if $\bar{\eta}_Q \geq \lambda$ then
- $d = \operatorname{argmax}(s_i)$ 8:

$$d_i \in Q$$

$$\mathcal{D}^{\mathcal{C}}(t) = \mathcal{D}^{\mathcal{C}}(t) \cup \{d\}$$

- end if 10:
- end for 11:

9:

return $\mathcal{D}^{C}(t)$ 12:

13: end function

2) Cluster Regression: Alg. 2 summarizes our approach to combine the set of augmented detections $S^{C}(t)$ into a set of refined target detections $\mathcal{D}^{C}(t)$. To reduce discrepancies among bounding boxes caused by segmentation errors, we leverage the pre-trained model to regress the set of augmented detections $S^{C}(t)$. As shown in Fig. 1, our cluster regression method uses the backbone features [72] with the augmented detections $S^{C}(t)$ as region proposals (instead of proposals generated using the region proposal network [73]) to the

Alg	orithm 3 Pseudo-Label Generation
1:	function PSEUDOLABELS($\mathcal{D}^{C}(t), r$)
2:	$\mathcal{P}^{\mathcal{C}}(t) = \emptyset, \Theta = \{i \cdot \Delta \theta\}_{i=1}^{r}$
3:	for $d_j \in D^{\mathcal{C}}(t)$ do
4:	for $ heta_i\in\Theta$ do
5:	Generate the augmented region proposals
	$d_{i,j} = R_{ heta_i}(d_j)$
6:	end for
7:	Generate the pseudo-label $(\hat{b}_i, \hat{m}_i, \alpha_i)$ using the
	region proposals $d_{i,j}$
8:	$\mathcal{P}^{C}(t) = \mathcal{P}^{C}(t) \cup \left\{ (\hat{b}_{i}, \hat{m}_{i}, \alpha_{i}) \right\}$
9:	end for
10:	return $\mathcal{P}^{C}(t)$
11:	end function

downstream box and mask heads. To avoid disregarding lowconfidence detections that might correspond to relevant region proposals, we do not apply non-maximum suppression to the model predictions. Fig. 3 shows that cluster regression significantly increases the accuracy of the bounding boxes generated using the augmented input, and the corresponding segmentation masks are consequently also more accurate.

a) Cluster Mode Detection: As Fig. 4 illustrates, detections and their corresponding confidence scores form a nonparametric distribution of the image's occupancy probability. We use the mean-shift algorithm [43] to identify the modes of that distribution and cluster detections corresponding to common targets. We cluster detections according to their bounding boxes b_j using a multivariate Gaussian kernel [43] with bandwidth h^c . We use the sample variances of the object bounding boxes at each frame to determine the kernel bandwidth, i.e.,

$$h^{\mathcal{C}} = \operatorname{diag}\left(\sum_{j=1}^{n_t^{\mathcal{C}}} (b_j^{\mathcal{C}} - \bar{b}_j^{\mathcal{C}})(b_j^{\mathcal{C}} - \bar{b}_j^{\mathcal{C}})^T\right),$$
(1)

where \bar{b}_j^C is the sample mean of b_j^C and diag (·) is the diagonal of the covariance matrix. The correlations among the elements of b_j are negligible and can be safely ignored. Each call to the mean-shift algorithm (line 4 in Alg. 2) produces a set of clusters $O^C(t)$ whose elements are sets of detections assigned to the same target. We consider the detections of passengers and baggage items separately. Hence, two separate invocations of the mean-shift procedure are required to produce the sets $O^{pax}(t)$ and $O^{bag}(t)$. The confidence score $\bar{\eta}_Q$ of cluster $Q \in O^C(t)$ is defined as the ratio between the total score of detections within that cluster and the number of rotation angles considered in the augmentation process, i.e.,

$$\bar{\eta}_Q = \frac{1}{r} \sum_{d_j \in Q} s_j. \tag{2}$$

Lines 6-10 of Alg. 2 show that we discard clusters with scores lower than a threshold λ to remove false positive detections.

3) Self-Supervised Model Update: Alg. 3 shows the procedure to generate the pseudo-labels used to update the model. Since our goal is to train the model using labels generated from multiple perspectives, we rotate both the original image and the corresponding predicted modes to generate pseudo-label proposals at each orientation. That is, for each mode $d_j \in D^{\mathcal{C}}(t)$, we generate the pseudo-label mask \hat{m}_j by using the rotated cluster modes $d_{i,j} = R_{\theta_i}(d_j), i = 1, \ldots, r$ as region proposals for the segmentation head, using the same approach described in Section III-A2. We then find the bounding box \hat{b}_j corresponding to \hat{m}_j . The confidence $\hat{\alpha}_j$ of the resulting pseudo-label is given by its corresponding cluster score. The set of pseudo-labels $\mathcal{P}^{\mathcal{C}}(t) = \left\{ (\hat{b}_j, \hat{m}_j, \hat{\alpha}_j) \mid d_j \in D^{\mathcal{C}}(t) \right\}$ thus contains accurate annotations even for targets that the model is unable to detect at certain orientations.

a) Rotation-Invariant Loss: To update the model using rotation-invariant pseudo-labels in a robust and efficient manner, we propose a novel uncertainty-aware, multi-task loss function given by

$$\mathcal{L} = \sum_{\hat{c} \in \mathcal{C}} \sum_{(\hat{b}_j, \hat{m}_j, \hat{\alpha}_j) \in \mathcal{P}^{\mathcal{C}}(t)} \hat{\alpha}_j (\mathcal{L}^c(\hat{c}, \tilde{c}) + \mathcal{L}^b(\hat{b}_j, \tilde{b}_j) + \mathcal{L}^m(\hat{m}_j, \tilde{m}_j)) + \mathcal{L}_{rpn}, \quad (3)$$

where \tilde{c} , b_j , and \tilde{m}_j are the object class, bounding box, and segmentation mask predicted by the network; \mathcal{L}^c , \mathcal{L}^b , and \mathcal{L}^m are the classification and bounding box regression losses defined in [73] and the pixel-wise binary cross entropy mask loss described in [3]; and \mathcal{L}_{rpn} is the region proposal network loss from [73]. In Eq. 3, the instance head losses are weighted by their corresponding cluster scores. This strategy ensures that instances with low cluster scores that might correspond to incorrect pseudo-labels have little impact on the update of the network parameters. As Alg. 4 indicates, a new set of pseudolabels is generated at each SSL iteration using the updated model from the previous iteration.

Algorithm 4 Self-Supervised Detection Model Update	
Input: Image sequence $I(t), t = 1,, T$	
Output: Updated detection model D_{PANet}	
1: repeat	
2: for $t = 1,, T$ do	
3: $S^{\mathcal{C}}(t) = \text{AugmentedProposals}(I(t))$	
4: $\mathcal{D}^{\mathcal{C}}(t) = \text{CLUSTERREGRESSION}(S^{\mathcal{C}}(t))$	
5: $\mathcal{P}^{\mathcal{C}}(t) = \text{PSEUDOLABELS}(\mathcal{D}^{\mathcal{C}}(t))$	
6: end for	
7: Fine-tune_the D_{PANet} model using the pseudo-labe	els
$\{\mathcal{P}^{\mathcal{C}}(t)\}_{t=1}^{T}$ according to the loss function in Eq. 3	

8: **until** Convergence criterion is met

B. Multi-View Passenger and Baggage Tracking

Our multi-camera tracking framework comprises two main steps: i) a single-camera, multiple-target tracking-by-detection algorithm, and ii) a multi-camera trajectory association mechanism. Our single-camera tracker uses the detections generated by our SSL framework and a Single-Camera Trajectory Association (SCA) method to keep track of the identities of individual passengers and baggage items within the field of view of each camera. Our MCTA strategy then projects the trajectories of passengers observed in cameras with overlapping fields of view onto a common image plane. These trajectories are then compared using the Fréchet distance and associated using a recursive graph-based mechanism.

1) Single-camera Tracking: We use the Tracktor algorithm [8] as our baseline single-camera tracker. The output of the algorithm at each image frame is a set $T^{\mathcal{C}}(t) = \{\omega_1, \dots, \omega_n c\},\$ where $\omega_i = [b_i, l_i]$, with l_i corresponding to a unique identifier label for each passenger and baggage item in the frame. These labels remain the same throughout the video sequence and hence perform temporal association among detections. The tracklet for the k-th object is thus given by the set of detections over the entire video sequence whose temporal identifier is $l_j = k$, i.e., $\tau_k = \bigcup_{t=1}^T \{ \omega_j \mid \omega_j \in T^{\mathcal{C}}(t), l_j = k \}.$ Temporary occlusions between passengers may lead to the fragmentation of trajectories within the field of view of a camera. Tracktor's simple re-identification strategy is unable to accommodate the longer occlusions, appearance variations, and somewhat erratic motion patterns commonly observed in airport checkpoints. Thus, we incorporate an SCA mechanism to resolve this issue. Our method associates new tracklets with recently terminated tracklets such that the Euclidean distance between the centroids of the last detection of the previous tracklet and the first detection of the new tracklet is minimized. That is, let τ_m and τ_n be two distinct tracklets, and $b_m^{t^*}$, $b_n^{t^f}$ be the first detection of τ_m and the last detection of τ_n , respectively. Defining $\delta_e = ||b_m^{t^i} - b_n^{t^f}||^2$, the association cost between τ_m and τ_n is given by

$$\mathcal{C}_{sc}(\tau_m, \tau_n) = \begin{cases} \delta_e & \text{if } 0 < t^i - t^f \leqslant t_{th}, \, \delta_e < \delta_{max} \\ \infty & \text{otherwise,} \end{cases}$$
(4)

where t_{th} , δ_{max} are the maximum temporal offset and maximum distance to consider two tracklets for association. We then compute the optimal tracklet assignment using the Hungarian algorithm based on the costs $C_{sc}(\tau_m, \tau_n)$.

2) Multi-Camera Tracklet Association: Since passengers may temporarily leave and later re-enter the fields of view of individual cameras, their corresponding trajectories may be fragmented into multiple segments. To associate tracklets across camera views, we consider the fact that two tracklets corresponding to the same target include temporally overlapping detections. Let the camera whose partial tracklets we wish to complete be our *primary* camera, and let the auxiliary camera be the one whose tracklets will be used to complement the tracklets observed by the primary camera. Further, let \mathcal{T}_p and \mathcal{T}_a be the sets of tracklets in the primary and auxiliary cameras, respectively. As Alg. 5 shows, we use the homography $H_{p,a}$ to project detections from the auxiliary camera onto the primary camera. However, due to projective distortions, the corresponding bounding boxes in the two cameras may not necessarily overlap. Hence, we compute the optimal association cost using the Fréchet distance [74] between the centroids of the detections in each tracklet as follows

$$\mathcal{C}_{mc}(\tau_a, \tau_p) = \begin{cases} \mathbf{f}(\tilde{\tau}_p, \tilde{\tau}_a) & \text{if } \tilde{\tau}_a \neq \emptyset, \ \tilde{\tau}_p \neq \emptyset, \ \mathbf{f} < f_{max} \\ \infty & \text{otherwise,} \end{cases}$$
(5)

Algorithm 5 Multi-Camera Tracklet Association Algorithm

Input: Set of tracklets from the primary camera T_p and the auxiliary camera \mathcal{T}_a , homography $H_{p,a}$ mapping the auxiliary camera image plane to that of the primary camera

Output: Updated set of primary tracklet labels

- 1: Project the detections of tracklets in \mathcal{T}_a onto the image plane of the primary camera using $H_{p,a}$
- 2: Compute the association costs $C_{mc}(\tau_a, \tau_p) \ \forall \tau_p \in \mathcal{T}_p$, $\forall \tau_a \in \mathcal{T}_a \text{ according to Eq. } 5 \\ \text{3: Initialize the graph } \mathcal{G}_{mc} = (V, E), E = \emptyset, V = \{\tau | \tau \in$
- $\mathcal{T}_p \cup \mathcal{T}_a$
- 4: while $\min_{\tau_p \in \mathcal{T}_p, \tau_a \in \mathcal{T}_a} (\mathcal{C}_{mc}(\tau_a, \tau_p)) < \infty$ do
- Associate tracklet segments using the Hungarian algo-5: rithm based on the costs C_{mc} Update the costs of the tracklets $\tau \in T_a$ and $\tau' \in T_p$
- 6: for which $\tau \cap \tau_a \notin \emptyset$ and $\tau' \cap \tau_p \notin \emptyset$ to $\mathcal{C}_{mc}(\tau, \tau_p) =$ $\mathcal{C}_{mc}(\tau_a, \tau') = \infty$ $E = E \cup (\tau_a, \tau_p)$
- 7:
- 8: end while
- 10:
- 9: for each $\tau_p \in \mathcal{T}_p$ do 10: $\mathcal{N}_p = \text{DFS}(\tau_p, \mathcal{G}_{mc})$ 11: Update the labels of tracklets in \mathcal{N}_p using Eq. 6 11:
- $E = E \{(\tau_i, \tau_j) | (\tau_i, \tau_j) \in \mathcal{N}_p\}$ 12:

where $\tilde{\tau}_p$ and $\tilde{\tau}_a$ are the temporally overlapping segments of tracklets $\tau_p \in \mathcal{T}_p$ and $\tau_a \in \mathcal{T}_a$, $\mathbf{f}(\tilde{\tau}_p, \tilde{\tau}_a)$ is the Fréchet distance of the centroids of the corresponding detections, and f_{max} is the maximum distance threshold that allows tracklet pairs to be considered for association.

We use the Hungarian algorithm again to determine optimal tracklet associations according to the costs $C_{mc}(\tau_a, \tau_p)$. However, since the trajectory of a passenger that re-enters the field of view of a camera multiple times consists of a sequence of tracklets, we iteratively update the association costs until no further associations are possible. We keep track of indirectly associated tracklets by constructing the reachability graph $\mathcal{G}_{mc} = (V, E)$, which contains one edge for each pair of associated tracklets. We then set the temporal identifiers of all the tracklets in \mathcal{T}_p associated with a common tracklet τ_a to the first identifier among them. That is, the temporal label of a tracklet τ is given by

$$l_{\tau} = \min_{(\tau_i, \tau_j) \in \mathcal{N}_p} (l_{\tau_i}), \tag{6}$$

where l_{τ_i} is the temporal label of tracklet τ_i , and \mathcal{N}_p is the set of tracklets that can be reached from tracklet τ_p on \mathcal{G}_{mc} , which we obtain through Depth-First Search (DFS).

IV. RESULTS AND DISCUSSION

In this section, we first discuss the datasets that we used to evaluate our algorithms. We then present an assessment of the proposed SSL approach in terms of passenger and baggage detection, followed by an evaluation of the single-camera tracking and multi-view tracklet association algorithms. Our evaluation is based on the Multi-Object Detection (MOD) and Tracking (MOT) metrics [59], [75]. Additional results are presented in the Supplementary Materials.



Fig. 5. Document checking station and divestiture area at the Kostas Research Institute simulated airport checkpoint.

A. Datasets

The video datasets used in this work were recorded at the Kostas Research Institute (KRI) video analytics laboratory at Northeastern University. As shown in Fig. 5, the laboratory is configured to emulate a realistic airport checkpoint. It is equipped with 14 standard IP surveillance cameras (Bosch NDN-832-V03P) with 1920×1080 resolution and focal lengths between 3 mm and 9 mm. The cameras are installed approximately three meters from the floor with partially overlapping fields of view. Fig. 6 shows a panoramic perspective of the fields of view of the cameras.

Several actors traverse the checkpoint with baggage items while performing a variety of activities commonly observed in real airports.¹ These activities range from simple scenarios in which just a few passengers pass through the checkpoint in sequential order to crowded scenes in which multiple passengers divest and retrieve their items in a more erratic manner. We collected two separate video datasets: CLASP1, which includes relatively simple scenarios, and CLASP2, which is more complex. Fig. 7 shows sample frames of videos from the two datasets. Of the 14 cameras in the laboratory, most passenger interactions take place on cameras 9 and 11. Camera 9 monitors the divestiture area and camera 11 observes the baggage retrieval area. Passengers place their belongings into bins or directly on the conveyor belt in the divestiture area. Then, after passing through the metal detector, they collect their belongings in the baggage retrieval area.

As Table I shows, a total of 146 passengers carrying 126 baggage items leave and re-enter the fields of view of the cameras several times. We manually annotate the videos with uniquely identified axis-aligned bounding boxes. Given the

 TABLE I

 DATASETS USED TO EVALUATE OUR ALGORITHMS. FOR EACH VIDEO

 SEQUENCE, THE TABLE SHOWS THE NUMBER OF PASSENGERS, BAGGAGE

 ITEMS, VIDEO FRAMES, ANNOTATED FRAMES, AND THE TOTAL NUMBER

 OF ANNOTATED BOUNDING BOXES.

Dataset	Video seq.	Pass- engers	Bag- gage	Video frames	Annotated frames/rate [fps]	Bounding boxes
	А	12	10	6,030	288 (1)	995
	В	12	10	6,180	564 (2)	1,720
CLASP1	С	8	9	6,030	491 (2)	853
	D	12	8	6,030	523 (2)	1,197
	Е	9	9	4,719	1,648 (10)	4,254
	F	20	20	12,910	179 (0.01)	737
CLASP2	G	38	31	10,390	1,346 (3)	4,826
	Н	35	29	11,200	198 (0.01)	900
Total	-	146	126	63,489	5,237	15,482

large number of video frames available in the datasets, the annotation rate for the video sequences varies between 0.01 and 10 frames per second (fps). We randomly partition each dataset into a training set containing 80% of the video frames and a test set with the remaining 20%. For a fair comparison, the Supervised Learning (SL) and SSL models are trained using only the frames from the training set, but the SSL models are fully self-supervised and do not use any manual annotations. However, disregarding every video sequence that includes annotated frames would substantially limit the amount of data available for the computation of tracking performance measures. Hence, to assess tracking performance, we consider all the annotations listed in Table I. The only method that uses the training set annotations is the SL approach. Although this evaluation strategy favors that method, it also more accurately reflects the generalization performance of the SSL approaches to unseen data. Due to space limitations, the results presented in this section were obtained using the aggregated CLASP1 and CLASP2 test sets. Dataset-specific results are given in the Supplementary Materials.

B. Self-Supervised Learning Detection Performance

During training, we freeze the network weights up to the region proposal network layer so that the pre-trained backbone features are effectively used in the downstream task. We use an initial learning rate of 5e–3, mini-batch size per image N = 256, r = 20 different orientations, and a cluster confidence threshold $\lambda = 0.1$. Similar to the baseline model, we use stochastic gradient descent with a momentum of 0.9, weight decay of 1e–4. At each SSL iteration, we fine-tune the model for 20k iterations, reducing the learning rate by a factor of 10 at every 5k iterations. In our evaluation, we use an IoU threshold of 0.5, and a non-maximum suppression threshold for region proposal generation is $\eta_{det} = 0.5$.

Fig. 8 shows the Multi-Object Detection Accuracy (MODA) of our model as a function of the number of SSL iterations. To illustrate the impact of the cluster confidence score, we also evaluate a model in which the samples are not weighed by their scores (SSL-wo- α). Instead, this model uses a hard threshold $\lambda \leq 0.4$ to discard noisy detections during training. The figure also shows the performance of the Multiple-Inference (MI)

¹The datasets are available upon request at alert-coe@northeastern.edu. Northeastern University's Institutional Review Board (IRB) and the Compliance Assurance Program Office (CAPO) within the DHS Science and Technology Directorate have reviewed the referenced human subjects research protocol and related research documentation. No compliance issues or concerns related to the use of human subjects in this protocol have been identified through the review, and the DHS policy requirements for human subjects research protocol review has been met.



Fig. 6. Panoramic overview of the camera views at the Kostas Research Institute simulated airport checkpoint.



Fig. 7. Sample images from the datasets collected at the simulated airport checkpoint (left: CLASP2 and right: CLASP1 in Table I). The images show the divestiture area (right: camera 9) and item retrieval area (left: camera 11).



Fig. 8. MODA measures for person (left) and baggage (right) classes during SSL training.

strategy used to generate the pseudo-labels, which reflects the quality of the pseudo-labels before SSL training. That is, in the MI model, the pseudo-labels themselves are used as model predictions. As the figure indicates, the SSL models gradually approach the performance of the MI strategy. The incorporation of cluster confidences not only increases the speed of convergence of the models but also leads to noticeable performance gains, particularly for baggage items.

Fig. 9 shows the precision-recall curves for passenger and baggage detection using four detector models: pre-trained PANet (baseline), PANet trained using SL, SSL-wo- α , and SSL. Even though the SSL models are trained without manual annotations, they perform on par with the SL model for passengers. For baggage items, the maximum average precision for the baseline model is less than half of the performance of the SSL models. As illustrated in Fig. 14, the performance



Fig. 9. Precision-recall curves for person (left) and baggage (right) detection. The legend shows the average precision of the models.

difference between the SL and SSL models is due to two main issues: i) appearance similarities among bags and certain garments/items placed inside security bins, and ii) baggage items that can only be partially observed before being placed on the conveyor belt.

Table VII demonstrates the benefits of incorporating cluster uncertainties in the SSL loss function (column α) and of the proposed cluster regression technique (column *reg.*). The method that incorporates both cluster uncertainty and regression is equivalent to the approach identified as SSL in Figs. 8 and 9 whereas the method that does not include cluster confidences corresponds to SSL-wo- α . The results in the table correspond to the point that maximizes the F_1 score of the curves in Fig. 9 at the best performing SSL iteration. The top-performing method in Table VII and in the remainder of this section is highlighted in boldface, the second-best is underlined, and ties are broken according to the MODA/MOTA results.

In comparison with the baseline model, our SSL algorithm substantially increases the recall (Rcll) and precision (Prcn) for passenger detection, which is a result of improvements in true positive (TP), false positive (FP), and false negative (FN) detections. The cluster confidence scores substantially reduce the contribution of low-confidence pseudo-labels, especially for baggage items, leading to a noticeable increase in the number of true positives. Cluster regression corrects pseudo-



Fig. 10. Sample results showing failure cases for baggage detection using the SSL model in the CLASP2 dataset. The magenta arrows indicate bag-like object detections that are not annotated (false positives), the red arrows indicate annotated baggage items the model fails to detect (false negatives), the green bounding boxes show passenger detections, and the red bounding boxes represent the manual annotations for both classes.

TABLE II PASSENGER AND BAGGAGE DETECTION EVALUATION.

Model	Me	ethod	↑Ro	ell	↑Pro	cn	↑T	Р	↓FI		↓FN	N I	↑MO	DA
	α	reg.	person	bag	person	bag	person	bag	person	bag	person	bag	person	bag
Baseline	X	X	73.8	37.1	87.0	82.9	1560	426	228	85	552	724	62.8	29.3
SSL	X	X	93.8	73.8	92.3	82.5	1989	858	155	194	123	291	86.0	57.6
SSL	\checkmark	X	93.5	75.9	93.5	86.1	1985	863	134	144	127	286	87.1	62.9
SSL	X	\checkmark	93.6	73.1	96.2	92.5	1985	844	73	71	127	305	90.1	67.1
SSL	\checkmark	\checkmark	95.7	78.6	96.0	91.8	2025	903	79	83	87	246	91.8	71.5
SL	X	X	95.6	91.4	96.4	92.8	2022	1048	70	83	90	101	92.1	84.4



Fig. 11. Qualitative detection results on the CLASP2 dataset using (a) Baseline, (b) SSL, and (c) SL models (the SL model only predicts bounding boxes).

label errors caused by inaccurate bounding boxes generated from poor segmentation results. As a result, the reduction in false positives for both classes is even more pronounced when cluster regression is incorporated. Overall, our SSL framework shows a relative MODA score improvement of 46% for passengers and 144% for baggage items with respect to the baseline model.

Fig. 11 shows qualitative results for the models under consideration. In comparison with the SL model, the SSL models not only improve the accuracy of the predicted bounding boxes but also generate improved segmentation masks since they are trained using instance segmentation pseudo-labels.

C. Single-Camera Tracking

We compare the performance of our single-camera tracking algorithm using the proposed SSL detectors with the pretrained baseline detector and the SL detector. We also evaluate the impact of our SCA algorithm, described in Section III-B1, where we use $t_{th} = 3$ seconds and $\delta_{max} = 200$ in Eq. 4. To preserve the entirely self-supervised nature of our pipeline, we refrain from fine-tuning the re-identification module of the baseline tracker, which is pre-trained on the MOT17 [59] dataset. To dissociate the evaluation of the tracking method from our MCTA approach, we use a modified version of the annotations in Table I where a passenger that re-enters the field of view of a camera receives a new identifier. Thus, the number of unique ground truth passenger identifiers (column GT in Table III) is much higher than those listed in Table I. We evaluate our system's ability to maintain consistent passenger identifiers across multiple perspectives in Section IV-D.

As Table III shows, the SSL-wo- α and SSL approaches outperform the tracker using the baseline detector by a large margin. The notable improvements in identity-based F_1 (IDF1), recall (IDR), and precision (IDP) [76] as well as in standard recall and precision are primarily a result of the reduction in false positives and false negatives generated by the SSL model. Self-supervision also improves the tracking-specific metrics of mostly tracked (MT), mostly lost (ML), identity switches (IDs), and fragmented (FM) trajectories [75]. As a result, our method produces substantial gains in MOTA. Again, both SSL models perform on par with the SL model for person tracking. For baggage items, we see similar performance improvements, but the challenges illustrated in Fig. 14 again preclude the SSL models from reaching the performance of the SL strategy. Finally, our SCA algorithm leads to further performance gains, particularly in terms of IDs.

D. Multi-Camera Tracklet Association

We evaluate the performance of our MCTA algorithm using the same experimental procedure described in the previous section, with the exception that passengers are now assigned unique identifiers as they leave and re-enter the fields of

 TABLE III

 SINGLE-CAMERA TRACKING EVALUATION FOR PERSON AND BAGGAGE CLASSES.

Class	Model	α	SCA	GT	†IDF1	↑IDR	↑IDP	↑Rcll	↑Prcn	↓FP	↓FN	↑MT	↓ML	↓IDs	${\downarrow}\text{FM}$	↑MOTA	↑MOTP
	Baseline	X	X	391	84.5	83.3	86.1	91.2	94.6	750	753	283	42	93	152	84.0	85.5
	SSL	X	X	391	87.8	87.2	88.3	95.1	96.4	350	554	319	31	80	123	90.1	85.2
Person	SSL	\checkmark	X	391	88.4	87.9	88.9	<u>95.6</u>	96.6	354	438	326	27	86	122	90.7	85.2
	SSL	\checkmark	\checkmark	391	88.5	88.1	88.9	95.6	96.5	358	435	326	26	76	123	90.8	85.2
	SL	X	X	391	87.0	86.4	87.7	95.2	96.9	357	457	332	24	<u>86</u>	121	90.5	85.2
	Baseline	X	X	255	67.5	57.0	86.4	61.1	92.3	431	1800	108	73	31	89	54.3	81.0
	SSL	X	X	255	78.9	74.2	85.4	81.0	93.7	308	1014	159	38	71	105	72.9	80.4
Bag	SSL	\checkmark	X	255	81.3	78.2	85.5	84.4	92.3	401	822	169	28	68	97	75.1	80.4
	SSL	\checkmark	\checkmark	255	<u>84.3</u>	<u>81.1</u>	88.5	84.4	92.3	401	822	169	<u>28</u>	<u>48</u>	97	<u>75.6</u>	80.4
	SL	X	×	255	85.3	86.6	84.1	94.7	91.7	<u>387</u>	339	226	10	103	<u>69</u>	83.2	80.0



Fig. 12. Sample results showing (a) cross-camera passenger association between cameras 5 and 11 using MCTA, and (b) tracking and association between passengers and baggage items where the top and bottom rows show image sequences from cameras 9 and 11 respectively. We associate passenger tracklets in cameras 9 and 11 by leveraging the associations between cameras 2 and 5 (passengers flow in Fig. 6: $C9 \rightarrow C2 \rightarrow C5 \rightarrow C11$). Baggage items are associated using temporally constrained distance-based matching when each item receives a unique identifier PiBj, representing the *j*-th item from the *i*-th passenger.

TABLE IVMCTA EVALUATION. THE COLUMN LABELED DIST. INDICATES WHETHER
WE EMPLOY THE HAUSDORFF (d_h) OR FRÉCHET (d_f) DISTANCE TO
EVALUATE TRACKLET SIMILARITY.

Dist.	SL	$\text{SSL-wo-}\alpha$	SSL	MCTA	†IDF1	↑IDR	↑IDP	↓IDs	↑MOTA
	X	\checkmark	X	X	82.1	83.2	81.0	157	88.2
-	X	×	\checkmark	X	82.0	83.1	80.9	170	88.5
	\checkmark	×	X	X	81.5	82.8	80.4	170	88.0
	X	\checkmark	X	\checkmark	87.4	88.9	86.4	115	88.8
d_h	X	×	\checkmark	\checkmark	84.8	88.6	86.2	140	89.0
	\checkmark	×	X	\checkmark	86.2	87.5	85.0	134	88.6
	X	\checkmark	X	\checkmark	88.0	89.3	86.8	115	88.9
d_f	X	×	\checkmark	\checkmark	88.2	89.5	87.0	132	89.1
	\checkmark	×	X	\checkmark	86.7	88.1	85.5	122	89.0

view of the cameras. Based on the overall flow of passengers through our simulated checkpoint, cameras 9 and 11 are the primary cameras for our tracklet association method (Alg. 5). Cameras 2 and 5, the cameras immediately below them in Fig. 6, are the respective auxiliary cameras. For a fair comparison among the detectors, we generate tracklets in the auxiliary cameras using the corresponding SL or SSL model used in the primary cameras (i.e., trained using only frames from the

primary camera). To provide a set of reference performance measures, we first evaluate our tracking algorithms in the absence of a MCTA mechanism. We then assess the performance of our association method when tracklet similarity is computed using the Fréchet distance and the more traditional Hausdorff distance [64] with $f_{max} = 0.25$ in Eq. 5.

Table VIII shows that tracklet association improves the IDF1 measure by up to 6.2%. This is mainly a consequence of the dramatic reduction in the number of identity switches. Using the Fréchet distance to determine tracklet similarity provides consistent performance improvements in all the metrics under consideration, especially for the SSL strategy. The more modest gains in MOTA (up to 1.0%) demonstrate the need for measures that focus specifically on the impact of identity switches on tracking performance.

Fig. 12(a) illustrates the tracklet association procedure between cameras 5 and 11. As the passengers with identities P2 and P3, whose trajectories are represented in green and yellow, move from the field of view of camera 5 to camera 11, their tracklets are projected from the former camera to the latter. The projected trajectories (red for P2 and pink for P3) are successfully associated with the tracklets from camera 11 based on the Fréchet distances among their temporally overlapping segments. In the instant shown in the figure, passenger P2 is re-entering the field of view of camera 5, and the corresponding tracklet is also correctly associated with that passenger's tracklet in camera 11. Hence, the passenger's identity is successfully handed off between the cameras. Fig. 12(b) demonstrates a potential application of the proposed system. Baggage items are associated with passengers when they are divested in camera 9 and their identifiers can be verified at retrieval time, which is observed in camera 11.

V. CONCLUSION

We propose a multistage tracking-by-detection framework to overcome performance limitations of object detection and tracking algorithms in overhead camera videos for which limited training data is available. Our framework is composed of an SSL mechanism to fine-tune object detection models to specific camera views without the need for manual annotations and an MCTA method that only requires the homographies among neighboring cameras. Our experiments show that the proposed framework can accurately detect and track passengers and baggage items across camera views in airport checkpoint scenarios. Our framework is flexible and scalable. It requires no training data, incurs no detection computational overhead at inference time, and is independent of the number of cameras in the network.

Our framework also allows seamless integration of additional data augmentation strategies and of manually annotated data when it is available. Our experiments show that these strategies further improve the selectivity of our detector, particularly for baggage items. For simplicity, our association methods are performed offline, i.e., after all the tracklets have been generated. However, it would be simple to implement online versions of the algorithms since the singleview association method can be executed whenever a new trajectory is initiated and each iteration of the MCTA algorithm can be performed once a trajectory in an auxiliary camera is terminated. The implementation of a real-time version of our tracklet association method is part of our future work.

REFERENCES

- Z. Wu and R. J. Radke, "Real-time airport security checkpoint surveillance using a camera network," in *Comp. Vis. and Pattern Recog. Workshops*, 2011.
- [2] A. Islam, Y. Zhang, D. Yin, O. Camps, and R. J. Radke, "Correlating belongings with passengers in a simulated airport security checkpoint," in *Int. Conf. on Distributed Smart Cameras*, 2018.
- [3] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask R-CNN," in *IEEE Int. Conf. on Comp. Vis.*, 2017.
- [4] W. Liu, D. Anguelov et al., "SSD: single shot multibox detector," in European Conf. on Comp. Vis., 2016.
- [5] S. Liu, L. Qi, H. Qin, J. Shi, and J. Jia, "Path aggregation network for instance segmentation," in *IEEE/CVF Conf. on Comp. Vis. and Pattern Recog.*, 2018.
- [6] G. Lin, A. Milan, C. Shen, and I. Reid, "Refinenet: Multi-path refinement networks for high-resolution semantic segmentation," in *IEEE Conf. on Comp. Vis. and Pattern Recog.*, 2017.
- [7] D. Comaniciu and P. Meer, "Mean shift: a robust approach toward feature space analysis," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 5, pp. 603–619, 2002.
- [8] P. Bergmann, T. Meinhardt, and L. Leal-Taixé, "Tracking without bells and whistles," in *IEEE/CVF Int. Conf. on Comp. Vis.*, 2019.

- [9] W. Lee, J. Na, and G. Kim, "Multi-task self-supervised object detection via recycling of bounding box annotations," in *IEEE/CVF Conf. on Comp. Vis. and Pattern Recog.*, 2019.
- [10] S. A. Golestaneh and K. Kitani, "Importance of self-consistency in active learning for semantic segmentation," in *British Mach. Vis. Conf.*, 2020.
- [11] M. Cai, M. Luo, X. Zhong, and H. Chen, "Uncertainty-aware model adaptation for unsupervised cross-domain object detection," arXiv preprint arXiv:2108.12612, 2021.
- [12] Y. Wang, J. Peng, and Z. Zhang, "Uncertainty-aware pseudo label refinery for domain adaptive semantic segmentation," in *IEEE/CVF Int. Conf. on Comp. Vis.*, 2021.
- [13] J. Mao, Q. Yu, Y. Yamakata, and K. Aizawa, "Noisy annotation refinement for object detection," in *British Mach. Vis. Conf.*, 2021.
- [14] I. Radosavovic, P. Dollár, R. Girshick, G. Gkioxari, and K. He, "Data distillation: Towards omni-supervised learning," in *IEEE/CVF Conf. on Comp. Vis. and Pattern Recog.*, 2018.
- [15] R. Mazzon and A. Cavallaro, "Multi-camera tracking using a multi-goal social force model," *Neurocomputing*, vol. 100, pp. 41 – 50, 2013.
- [16] S. Zhang, Y. Zhu, and A. Roy-Chowdhury, "Tracking multiple interacting targets in a camera network," *Comp. Vis. and Image Understanding*, vol. 134, pp. 64 – 73, 2015.
- [17] K. Hong, H. Medeiros, P. J. Shin, and J. Park, "Resource-aware distributed particle filtering for cluster-based object tracking in wireless camera networks," *Int. J. of Sensor Networks*, vol. 21, no. 3, pp. 137– 156, 2016.
- [18] H. Medeiros, J. Park, and A. Kak, "Distributed object tracking using a cluster-based kalman filter in wireless camera networks," *IEEE J. Sel. Topics Signal Process.*, vol. 2, no. 4, pp. 448–463, 2008.
- [19] N. Wang, Y. Song et al., "Unsupervised deep tracking," in IEEE/CVF Conf. on Comp. Vis. and Pattern Recog., 2019.
- [20] R. Jalil Mozhdehi, Y. Reznichenko, A. Siddique, and H. Medeiros, "Deep convolutional particle filter with adaptive correlation maps for visual tracking," in 25th IEEE Int. Conf. on Image Processing, 2018.
- [21] R. Jalil Mozhdehi and H. Medeiros, "Deep convolutional correlation iterative particle filter for visual tracking," *Computer Vision and Image Understanding*, p. 103479, 2022.
- [22] M. Chuang, J. Hwang, J. Ye, S. Huang, and K. Williams, "Underwater fish tracking for moving cameras based on deformable multiple kernels," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 47, no. 9, pp. 2467–2477, 2017.
- [23] R. Ding, M. Yu, H. Oh, and W. Chen, "New multiple-target tracking strategy using domain knowledge and optimization," *IEEE Trans. Syst.*, *Man, Cybern., Syst.*, vol. 47, no. 4, pp. 605–616, 2017.
- [24] R. Henschel, Y. Zou, and B. Rosenhahn, "Multiple people tracking using body and joint detections," in *IEEE/CVF Conf. on Comp. Vis.* and Pattern Recog. Workshops, 2019.
- [25] A. Sadeghian, V. Kosaraju *et al.*, "Sophie: An attentive GAN for predicting paths compliant to social and physical constraints," in *IEEE/CVF Conf. on Comp. Vis. and Pattern Recog.*, 2019.
- [26] A. Alahi, K. Goel et al., "Social LSTM: Human trajectory prediction in crowded spaces," in *IEEE Conf. on Comp. Vis. and Pattern Recog.*, 2016.
- [27] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi, "Social GAN: Socially acceptable trajectories with generative adversarial networks," in *IEEE/CVF Conf. on Comp. Vis. and Pattern Recog.*, 2018.
- [28] M. Babaee, A. Athar, and G. Rigoll, "Multiple people tracking using hierarchical deep tracklet re-identification," *arXiv preprint* arXiv:1811.04091, 2018.
- [29] M. Keuper, S. Tang, B. Andres, T. Brox, and B. Schiele, "Motion segmentation & multiple object tracking by correlation co-clustering," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 1, pp. 140–153, 2020.
- [30] S. M. Khan and M. Shah, "A multiview approach to tracking people in crowded scenes using a planar homography constraint," in *European Conf. on Comp. Vis.*, 2006.
- [31] T. Chavdarova, P. Baqué et al., "Wildtrack: A multi-camera HD dataset for dense unscripted pedestrian detection," in *IEEE/CVF Conf. on Comp. Vis. and Pattern Recog.*, 2018.
- [32] P. Baqué, F. Fleuret, and P. Fua, "Deep occlusion reasoning for multicamera multi-target detection," in *IEEE Int. Conf. on Comp. Vis.*, 2017.
- [33] A. Zheng, X. Zhang, B. Jiang, B. Luo, and C. Li, "A subspace learning approach to multishot person reidentification," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 50, no. 1, pp. 149–158, 2020.
- [34] J. Si, H. Zhang, C.-G. Li, and J. Guo, "Spatial pyramid-based statistical features for person re-identification: A comprehensive evaluation," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 48, no. 7, pp. 1140–1154, 2018.

- [35] H.-M. Hsu, T.-W. Huang *et al.*, "Multi-camera tracking of vehicles based on deep features re-id and trajectory-based camera link models," in *IEEE/CVF Conf. on Comp. Vis. and Pattern Recog. Workshops*, 2019.
- [36] P. Li, J. Zhang et al., "State-aware re-identification feature for multitarget multi-camera tracking," in *IEEE/CVF Conf. on Comp. Vis. and Pattern Recog. Workshops*, 2019.
- [37] N. Peri, P. Khorramshahi et al., "Towards real-time systems for vehicle re-identification, multi-camera tracking, and anomaly detection," in IEEE/CVF Conf. on Comp. Vis. and Pattern Recog. Workshops, 2020.
- [38] K. Zhang, J. Chen, G. Yu, X. Zhang, and Z. Li, "Visual trajectory tracking of wheeled mobile robots with uncalibrated camera extrinsic parameters," *IEEE Trans. Syst., Man, Cybern., Syst.*, vol. 51, no. 11, pp. 7191–7200, 2021.
- [39] L. Jing and Y. Tian, "Self-supervised visual feature learning with deep neural networks: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 11, pp. 4037–4058, 2021.
- [40] M. Caron, P. Bojanowski, A. Joulin, and M. Douze, "Deep clustering for unsupervised learning of visual features," in *European Conf. on Comp. Vis.*, 2018.
- [41] P. Goyal, M. Caron *et al.*, "Self-supervised pretraining of visual features in the wild," *arXiv preprint arXiv:2103.01988*, 2021.
- [42] Z. Liu, J. Zhang, and L. Liu, "Upright orientation of 3D shapes with convolutional networks," *Graphical Models*, vol. 85, pp. 22 – 29, 2016.
- [43] M. Taj and A. Cavallaro, "Multi-camera track-before-detect," in Int. Conf. on Distributed Smart Cameras, 2009.
- [44] J. Zhu, H. Yang et al., "Online multi-object tracking with dual matching attention networks," in European Conf. on Comp. Vis., 2018.
- [45] J. Son, M. Baek, M. Cho, and B. Han, "Multi-object tracking with quadruplet convolutional neural networks," in *IEEE Conf. on Comp. Vis. and Pattern Recog.*, 2017.
- [46] S. Sun, N. Akhtar, H. Song, A. Mian, and M. Shah, "Deep affinity network for multiple object tracking," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 43, no. 1, pp. 104–119, 2021.
- [47] H. Sheng, J. Chen et al., "Iterative multiple hypothesis tracking with tracklet-level association," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 29, no. 12, pp. 3660–3672, 2019.
- [48] H. Sheng, Y. Zhang, J. Chen, Z. Xiong, and J. Zhang, "Heterogeneous association graph fusion for target association in multiple object tracking," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 29, no. 11, pp. 3269–3280, 2019.
- [49] C. Kim, F. Li, and J. M. Rehg, "Multi-object tracking with neural gating using bilinear LSTM," in *European Conf. on Comp. Vis.*, 2018.
- [50] K. Fang, Y. Xiang, X. Li, and S. Savarese, "Recurrent autoregressive networks for online multi-object tracking," in *IEEE Winter Conf. on Applications of Comp. Vis.*, 2018.
- [51] S. Schulter, P. Vernaza, W. Choi, and M. Chandraker, "Deep network flow for multi-object tracking," in *IEEE Conf. on Comp. Vis. and Pattern Recog.*, 2017.
- [52] S. Jin, W. Liu, W. Ouyang, and C. Qian, "Multi-person articulated tracking with spatial and temporal embeddings," in *IEEE/CVF Conf.* on Comp. Vis. and Pattern Recog., 2019.
- [53] P. Voigtlaender, M. Krause *et al.*, "MOTS: Multi-object tracking and segmentation," in *IEEE/CVF Conf. on Comp. Vis. and Pattern Recog.*, 2019.
- [54] C. Kim, F. Li, A. Ciptadi, and J. M. Rehg, "Multiple hypothesis tracking revisited," in *IEEE Int. Conf. on Comp. Vis.*, 2015.
 [55] A. Sadeghian, A. Alahi, and S. Savarese, "Tracking the untrackable:
- [55] A. Sadeghian, A. Alahi, and S. Savarese, "Tracking the untrackable: Learning to track multiple cues with long-term dependencies," in *IEEE Int. Conf. on Comp. Vis.*, 2017.
- [56] R. Ding, M. Yu, H. Oh, and W.-H. Chen, "New multiple-target tracking strategy using domain knowledge and optimization," *IEEE Trans. Syst.*, *Man, Cybern., Syst.*, vol. 47, no. 4, pp. 605–616, 2016.
- [57] A. Siddique, R. J. Mozhdehi, and H. Medeiros, "Unsupervised spatiotemporal latent feature clustering for multiple-object tracking and segmentation," in *British Mach. Vis. Conf.*, 2021.
- [58] D. Stadler and J. Beyerer, "Improving multiple pedestrian tracking by track management and occlusion handling," in *IEEE/CVF Conf. on Comp. Vis. and Pattern Recog.*, 2021.
- [59] P. Dendorfer, A. Ošep et al., "MOTChallenge: A benchmark for singlecamera multiple target tracking," Int. J. of Comp. Vis., vol. 129, p. 845–881, 2020.
- [60] Y. Xu, X. Liu, Y. Liu, and S.-C. Zhu, "Multi-view people tracking via hierarchical trajectory composition," in *IEEE Conf. on Comp. Vis. and Pattern Recog.*, 2016.
- [61] N. Anjum and A. Cavallaro, "Trajectory association and fusion across partially overlapping cameras," in *IEEE Int. Conf. on Advanced Video* and Signal Based Surveillance, 2009.

- [62] K. Nithin and F. Bremond, "Multi-camera tracklet association and fusion using ensemble of visual and geometric cues," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 27, no. 3, pp. 431–440, March 2017.
- [63] S. Bae and K. Yoon, "Robust online multi-object tracking based on tracklet confidence and online discriminative appearance learning," in *IEEE Conf. on Comp. Vis. and Pattern Recog.*, 2014.
- [64] A. S. Hassanein, M. E. Hussein, and W. Gomaa, "Semantic analysis of crowded scenes based on non-parametric tracklet clustering," in *Int. Joint Conf. on Artificial Intelligence*, 2016.
- [65] H. Germain, V. Lepetit, and G. Bourmaud, "Neural reprojection error: Merging feature learning and camera pose estimation," in *IEEE/CVF Conf. on Comp. Vis. and Pattern Recog.*, 2021.
- [66] J. Zhang, C. Wang et al., "Content-aware unsupervised deep homography estimation," in European Conf. on Comp. Vis., 2020.
- [67] Y. Bok, D.-G. Choi, P. Vasseur, and I. S. Kweon, "Extrinsic calibration of non-overlapping camera-laser system using structured environment," in *IEEE/RSJ Int. Conf. on Intell. Robots and Syst.*, 2014, pp. 436–443.
- [68] H. Medeiros, H. Iwaki, and J. Park, "Online distributed calibration of a large network of wireless cameras using dynamic clustering," in *Int. Conf. on Distributed Smart Cameras*, 2008.
- [69] S. Xie, R. Girshick, P. Dollár, Z. Tu, and K. He, "Aggregated residual transformations for deep neural networks," in *IEEE Conf. on Comp. Vis.* and Pattern Recog., 2017.
- [70] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conf. on Comp. Vis. and Pattern Recog.*, 2016.
- [71] T.-Y. Lin, M. Maire et al., "Microsoft COCO: Common objects in context," in European Conf. on Comp. Vis., 2014.
- [72] T. Y. Lin, P. Dollár et al., "Feature pyramid networks for object detection," in IEEE Conf. on Comp. Vis. and Pattern Recog., 2017.
- [73] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 39, no. 6, pp. 1137–1149, 2017.
- [74] K. Buchin, M. Buchin, and C. Wenk, "Computing the Fréchet distance between simple polygons," *Computational Geometry*, vol. 41, no. 1, pp. 2–20, 2008.
- [75] K. Bernardin and R. Stiefelhagen, "Evaluating multiple object tracking performance: The CLEAR MOT metrics." *EURASIP J. on Image and Video Processing*, vol. 2008, 2008.
- [76] E. Ristani, F. Solera, R. Zou, R. Cucchiara, and C. Tomasi, "Performance measures and a data set for multi-target, multi-camera tracking," in *European Conf. on Comp. Vis. Workshops*, 2016.
- [77] M. Xu, Z. Zhang et al., "End-to-end semi-supervised object detection with soft teacher," In IEEE/CVF Int. Conf. on Comp. Vis., 2021.

Supplementary Materials: Tracking Passengers and Baggage Items using Multiple Overhead Cameras at Security Checkpoints

Abubakar Siddique, Student Member, IEEE, Henry Medeiros, Senior Member, IEEE

Abstract—This document supplements our main paper with additional experimental results on the CLASP1 and CLASP2 datasets. We extend our Self-Supervised Learning (SSL) approach into a Semi-Supervised Learning (Semi-SL) mechanism to further improve target detection performance, especially for baggage items. We also investigate the impact of additional data augmentation strategies, rotation resolution, and the computational requirements of our proposed technique. These additional evaluation results show that our algorithm outperforms the baseline as well as state-of-the-art supervised and semi-supervised approaches.

Index Terms—Detection, Tracking, Association, Homography, Tracklet, Multi-camera, Surveillance.

VI. SINGLE CAMERA DETECTIONS

This section presents a breakdown on the performance of our SSL detector for individual cameras in the CLASP1 and CLASP2 datasets. It also evaluates the performance impact of additional data augmentation strategies, number of rotation angles used for data augmentation, and incorporation of labeled data in a semi-supervised approach.

A. Self-Supervised Learning

Fig. 13 shows a detailed breakdown of the performance of our SSL detection model for individual camera views in the CLASP1 and CLASP2 datasets. The high recall, precision, and MODA values indicate that our SSL approach detects most passengers correctly in these video sequences. The average precision (AP) for passenger detection is slightly higher for camera 11 in both datasets. The main factor contributing to this performance difference is that in camera 9, passengers are only partially visible most of the time, whereas camera 11 has a better view of the region where the passengers stand next to the conveyor belt. On the other hand, this also contributes to the lower baggage detection performance in

*Manuscript received August 22, 2022; accepted November 11, 2022. Date of publication December 14, 2022.

[†]This material is based upon work supported by the U.S. Department of Homeland Security, Science and Technology Directorate, Office of University Programs, under Award Number 2013-ST-061-E0001-04. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security.

[‡]Abubakar Siddique is with the Department of Electrical and Computer Engineering, Marquette University, Milwaukee, USA, e-mail: abubakar.siddique@marquette.edu

[§]Henry Medeiros is with the Department of Agricultural and Biological Engineering, University of Florida, Gainesville, USA, e-mail: hmedeiros@ufl.edu

[¶]Digital Object Identifier (DOI): 10.1109/TSMC.2022.3225252

camera 11. That is, in camera 11, partially observed baggage items being carried by passengers (see Fig. 14) are much more common than in camera 9. As with passenger detection, we observed similar baggage detection improvements in the camera-specific performance comparisons. This performance could be further improved by using additional unlabelled video frames available in the CLASP1 and CLASP2 datasets to train the SSL models.

B. Additional Data Augmentation Strategies

We investigate the impact of other data augmentation strategies during SSL training, including color jittering and motion blur along with multiple rotations. For color jittering, we increase/decrease image brightness, contrast, saturation, and hue by a factor sampled uniformly from the range $[0, \max_{iit}]$, where max_{iit} is 0.4 for brightness, 0.5 for contrast, 0.2 for saturation, and 0.05 for hue. To emulate motion blur, we use Gaussian blur with kernel size uniformly sampled from the set $\{5, \ldots, 9\}$ and standard deviation sampled from the interval [0.1, 5]. We observe that applying color jittering and motion blur on the pseudo-label augmentation further improves MODA scores by up to 2.9% and 4.8% for passengers and baggage items, respectively. For a fair comparison, we reduced the number of rotation angles used for augmentation such that the total number of augmented images remains the same in both scenarios. Maintaining the original number of rotations would further increase performance gains.

TABLE V Performance impact of additional data augmentation strategies in the SSL iterations.

Dataset		Met	hod	↑A	Р	$\uparrow F$	71	↑MODA		
	Rot.	C-Jit.	MotBlur	person	bag	person	bag	person	bag	
CLASD1	\checkmark	X	×	89.2	43.4	92.0	59.7	83.9	41.6	
CLASFI	\checkmark	\checkmark	\checkmark	91.5	48.3	92.3	64.5	84.3	46.4	
CLASD	\checkmark	X	×	79.4	47.4	86.2	62.5	73.6	42.2	
CLASP2	\checkmark	\checkmark	\checkmark	84.2	47.8	88.0	63.0	76.5	43.6	

C. Impact of Rotation Resolution

Table VI shows the impact of rotation resolution r on the generation of pseudo-labels. One SSL iteration with r = 20 improves the MODA scores by up to 3.1% for passengers and 5.6% for baggage items. The inference time for a single frame increases linearly with the number of rotations, contributing to longer SSL training iterations. If training time is a concern, r = 10 offers a reasonable speed vs. performance trade-off. We



Fig. 13. Passenger and baggage detection performance in cameras 9 and 11 for the CLASP1 (CL1) and CLASP2 (CL2) datasets. Here, P stands for passenger and B for baggage. A description of the methods under consideration is given in Section IV.B of the main paper.



Fig. 14. Additional illustrative failure cases for baggage detection using the SSL model in the CLASP1 dataset (see Fig. 10 in the main paper for failures in the more challenging CLASP2 dataset). The magenta arrows indicate bag-like object detections that are not annotated (false positives), the red arrows indicate annotated baggage items the model fails to detect (false negatives), the green bounding boxes show passenger detections, and the red bounding boxes represent manual annotations for both classes.

use r = 20 for all the SSL models to demonstrate the potential performance of our framework. As Table VI indicates, further increasing the value of r would likely lead to minor additional performance gains.

TABLE VI PERFORMANCE IMPACT OF THE NUMBER OF ROTATION ANGLES USED IN THE SSL ITERATIONS.

Dataset	r	↓Infer. Time	$\uparrow F$	1	↑MODA		
		(secs)	person	bag	person	bag	
	1	0.3	94.5	69.5	89.0	51.7	
CLACD1	5	2.2	95.0	70.4	90.1	52.8	
CLASPI	10	4.5	95.3	70.8	90.6	53.6	
	20	9.1	95.8	70.8	91.5	53.7	
	1	0.3	91.0	74.5	82.3	56.8	
CLACDO	5	2.6	92.1	76.4	84.6	59.6	
CLASP2	10	4.0	92.1	76.5	84.5	59.8	
	20	11.7	92.2	76.5	84.9	60.0	

D. Semi-Supervised Learning

As Table VII indicates, the performance of our SSL algorithm is limited by the initial accuracy of the baseline model. Thus, we extend our method to a semi-supervised approach where we use a certain amount of manual annotations to initialize our model before initiating SSL training. For the labeled frames, we employ the same data augmentation procedure used to generate augmented labels. Fig. 15 shows that training the SSL model using 10% of the manual labels leads to a performance comparable to the SL model, outperforming SoftTeacher [77], a state-of-the-art Semi-SL technique. Our method is particularly effective when small amounts of annotations are used. For example, using only 1% of the manual labels, our Semi-SL approach outperforms SoftTeacher by 104% and is only 1.6% behind the SL method (Table VII) for baggage items. Furthermore, we observe a 5.7% MODA improvement over the SL method when we use all the manual annotations during training.

VII. SINGLE-CAMERA TRACKING

Fig. 16 shows the Single-Camera Tracking (SCT) performance of our algorithm for passengers and baggage items in the individual cameras of the CLASP1 and CLASP2 datasets.

Dataset	Model	Me	ethod	↑Ro	:11	↑Pro	cn	↑TI	P	↓FI	2	↓FN	N	↑MO	DA
		α	reg.	person	bag	person	bag	person	bag	person	bag	person	bag	person	bag
	Baseline	X	X	73.8	36.2	89.0	87.9	886	233	110	32	314	411	64.7	31.2
	SSL	X	X	96.4	80.4	95.8	78.6	1157	518	51	141	43	126	92.2	58.5
CLASD1	SSL	\checkmark	X	96.9	70.4	94.2	<u>90.9</u>	1163	451	71	45	37	193	91.0	63.0
CLASFI	SSL	X	\checkmark	96.0	76.1	97.7	90.7	1152	490	27	50	48	154	93.8	68.3
	SSL	\checkmark	\checkmark	96.8	78.6	97.3	90.2	1162	506	32	55	38	138	94.2	70.0
	SL	X	X	96.7	89.4	98.1	91.4	1160	576	23	54	40	68	94.8	81.1
	Baseline	X	X	73.9	38.0	85.1	78.0	674	192	118	53	238	313	61.0	27.5
	SSL	X	X	91.2	67.3	88.9	86.5	832	340	104	53	80	165	79.8	56.8
CLASD	SSL	\checkmark	X	90.1	81.6	92.9	81.3	822	<u>412</u>	63	95	90	<u>93</u>	83.2	62.8
CLASP2	SSL	X	\checkmark	91.3	70.1	94.8	94.4	833	354	46	21	79	151	86.5	65.9
	SSL	\checkmark	\checkmark	94.6	78.6	94.8	93.4	863	397	<u>47</u>	28	49	108	89.5	73.1
	SL	X	X	94.5	93.5	94.8	94.2	862	472	47	29	50	33	89.4	87.7

 TABLE VII

 PASSENGER AND BAGGAGE DETECTION EVALUATION MEASURES ON THE CLASP1 AND CLASP2 TEST SETS.



Fig. 15. Semi-SL model performance on CLASP2 using a semi-supervised extension of our proposed SSL method versus SoftTeacher (ST) [77]. Here, P and B stand for the passenger and baggage categories. The SSL model uses no labeled data and the SL model is trained with 100% of the samples.

For passenger tracking, the SSL methods outperform the SL approach in terms of IDF1, IDP, and IDR in all the scenarios under consideration. In both datasets, the SL approach shows slightly higher MT results for camera 9, largely due to the partial passenger detection problem. Since CLASP1 has lower object density, we observe more consistent performance among different methods for both cameras in that dataset. While all the methods perform better on the CLASP1 dataset, the benefits of SSL training compared to the baseline detector are particularly evident in the MT results on the CLASP2 dataset.

Regarding baggage items, although the SSL models lead to a moderate increase in the number of IDs, these switches are offset by substantial gains in MT. As a matter of fact, the SL model shows a much more significant degradation in IDs for the more complex CLASP2 dataset. This is particularly evident for camera 9, and it explains the lower IDP obtained by the SL method in that dataset. The most evident performance gains for baggage tracking are observed in camera 11 on the CLASP2 dataset because the of the difficulty of partially visible baggage items using the baseline model.

VIII. MULTI-CAMERA TRACKLET ASSOCIATION

Regarding our Multi-camera Tracklet Association (MCTA) method, Table VIII shows that the Fréchet distance metric is particularly useful in crowded scenarios. Although we obtain comparable results using the Hausdorff distance on the easier CLASP1 dataset, we achieve noticeable improvements in all the evaluation criteria on CLASP2 using the Fréchet distance.

The single-camera trackers in the auxiliary cameras are trained using frames from the primary cameras. Hence, in crowded scenarios they sometimes fails to keep alive trajectories of targets that are temporarily outside the field of view of the primary camera. This is the main reason behind the overall lower tracking performance on the CLASP2 dataset. Training camera-specific detectors using our SSL framework would mitigate this issue.

TABLE VIII MCTA EVALUATION. THE COLUMN LABELED DIST. INDICATES WHETHER WE EMPLOY THE HAUSDORFF (d_h) OR FRÉCHET (d_f) DISTANCE TO EVALUATE TRACKLET SIMILARITY.

Data.	Dist.	SL	SSL-wo-α	SSL	MCTA	↑IDF1	↑IDR	↑IDP	↓IDs	↑MOTA
	-	X	\checkmark	X	X	87.4	87.8	86.9	45	94.4
	-	X	X	\checkmark	X	87.0	87.4	86.6	48	95.1
	-	\checkmark	×	X	X	87.3	87.5	87.2	42	94.4
		X	\checkmark	X	\checkmark	94.0	94.5	93.5	21	94.9
CLASP1	d_h	X	×	\checkmark	\checkmark	92.7	93.1	92.2	30	<u>95.5</u>
		\checkmark	X	X	\checkmark	93.0	93.1	92.8	25	94.8
		X	\checkmark	X	\checkmark	93.8	94.3	93.3	<u>22</u>	94.9
	d_f	X	X	\checkmark	\checkmark	92.8	93.3	92.4	27	95.6
	5	\checkmark	×	X	\checkmark	93.4	93.6	93.3	23	94.8
	-	X	\checkmark	X	X	76.9	78.7	75.1	112	82.0
	-	X	X	\checkmark	X	77.0	78.8	75.3	122	81.9
	-	\checkmark	×	X	X	75.8	78.2	73.6	128	81.6
		X	\checkmark	X	\checkmark	81.2	83.3	79.3	<u>94</u>	82.8
CLASP2	d_h	X	X	\checkmark	\checkmark	82.1	84.2	80.3	110	82.5
		\checkmark	X	X	\checkmark	79.5	82.0	77.2	109	82.5
		X	\checkmark	X	\checkmark	82.3	84.4	80.4	93	83.0
	d_f	X	X	\checkmark	\checkmark	83.6	85.7	81.7	105	82.7
	5	\checkmark	×	X	\checkmark	80.1	82.7	77.8	99	83.2

IX. COMPUTATIONAL COMPLEXITY

In this section, we analyze the theoretical computational complexity of our SSL strategy and measure the computation time and memory utilization of each step of our algorithm. All our experiments were performed on a workstation equipped with two RTX-2090Ti GPUs and an Intel[®] Xeon[®] Silver 4112 CPU @2.6*G*Hz.

A. Self-Supervised Learning

The computational complexity of our approach increases linearly with the number of rotation angles used for augmentation in the pseudo-label generation step. That is, for a baseline detection algorithm with computational complexity $\Theta(f(I(t)))$,



Fig. 16. Comparison of SCT performance of person and baggage classes in individual cameras of the CLASP1 and CLASP2 datasets using the Baseline, SSL-wo- α , SSL, and SL detectors.

the complexity of our approach is $\Theta(r \cdot f(I(t)))$, where r is the number of rotation angles. For example, for r = 20, the runtime is 20 times that of a single iteration without augmentation. However, these operations are parallelizable as long as the hardware resources support the simultaneous processing of multiple frames. With our unoptimized implementation, the total time to complete one SSL iteration is approximately six hours for both model training and pseudo-label generation. However, we have observed that hardware resources are severely underutilized, which indicates substantial room for reduction in overall computation time.

B. Inference Performance

Table IX shows the computation time of the proposed tracking-by-detection algorithm, employing a PANet detector with a ResNet-50 backbone. The SCT uses the detector results and a ResNet-50-based Re-Identification (Re-ID) model trained on MOT17 to re-label tracklets lost due to short-term occlusions. Hence, the computation time and memory utilization for the SCT are similar to those for the detector model. Since we are processing single images individually instead of image batches, the inference time for the detector and the SCT are far from optimal. Preliminary experiments indicate that processing batches of 10 images simultaneously leads to an approximate six-fold reduction in detector inference time without exceeding the memory capacity of the GPUs. Reusing the backbone features from the detector in the Re-ID model should also lead to a dramatic reduction in SCT time, since feature generation is the most computationally demanding element of the tracking algorithm.

The execution time of the proposed MCTA algorithm depends on the average length of the overlapping tracklet segments in each camera pair. In the CLASP1 dataset, which contains fewer and shorter tracklets, the algorithm can be executed in real time. In CLASP2, it can run at approximately

TABLE IX Computation time of the proposed tracking-by-detection framework.

Data	Model	Infer. Time (ms)	Memory (MB)
	Detector	333.3	1,850
CLASP1	SCT	142.8	1,748
	MCTA	25.6	9.1
	Detector	333.3	1,850
CLASP2	SCT	166.6	1,750
	MCTA	83.3	22.7

12 fps. However, the current implementation of the proposed system uses the full life-span of a tracklet to compute the Fréchet association distance in the MCTA algorithm. It is possible to substantially reduce computation time by limiting the length of single-camera tracklets compared by the algorithm. Table X shows that if we limit the length of the tracklets to 240 frames (or eight seconds), it is possible to achieve real-time performance for both datasets without degrading the accuracy of the algorithm.

TABLE X COMPUTATION TIME OF THE PROPOSED MCTA.

Data.	Dist. Metric	Max. Size	Infer. Time (ms)	Memory (MB)	MOTA
CLASP1	Hausdorff	-	0.52	4.5	94.5
		240	0.50	4.5	94.5
	Fréchet	-	25.6	9.1	94.6
		240	8.20	4.3	94.6
CLASP2	Hausdorff	-	0.64	16.3	78.4
		240	0.61	16.3	78.4
	Fréchet	_	83.3	22.7	78.5
		240	19.6	16.3	78.7