

Unsupervised Spatio-temporal Latent Feature Clustering for Multiple-object Tracking and Segmentation

Abubakar Siddique
abubakar.siddique@marquette.edu

Reza Jalil Mozhdehi
reza.jalilmozhdehi@marquette.edu

Henry Medeiros
henry.medeiros@marquette.edu

Opus College of Engineering
Marquette University
Milwaukee, USA

Abstract

Assigning consistent temporal identifiers to multiple moving objects in a video sequence is a challenging problem. A solution to that problem would have immediate ramifications in multiple object tracking and segmentation problems. We propose a strategy that treats the temporal identification task as a spatio-temporal clustering problem. We propose an unsupervised learning approach using a convolutional and fully connected autoencoder, which we call deep heterogeneous autoencoder, to learn discriminative features from segmentation masks and detection bounding boxes. We extract masks and their corresponding bounding boxes from a pretrained instance segmentation network and train the autoencoders jointly using task-dependent uncertainty weights to generate common latent features. We then construct constraints graphs that encourage associations among objects that satisfy a set of known temporal conditions. The feature vectors and the constraints graphs are then provided to the kmeans clustering algorithm to separate the corresponding data points in the latent space. We evaluate the performance of our method using challenging synthetic and real-world multiple-object video datasets. Our results show that our technique outperforms several state-of-the-art methods. Code and models are available at https://bitbucket.org/Siddiquemu/usc_mots.

1 Introduction

The goal of Multiple Object Tracking and Segmentation (MOTS) algorithms is to establish temporally consistent associations among segmentation masks of multiple objects observed at different frames of a video sequence. To accomplish that goal, most state-of-the-art MOTS methods [30, 39] employ supervised learning approaches to generate discriminative embeddings and then apply feature association algorithms based on sophisticated target behavior models [14, 23, 26]. This paper proposes a novel perspective on the problem of temporal association of segmentation masks based on spatio-temporal clustering strategies.

Subspace clustering algorithms applied to sequential data can separate sequences of similar data points into disjoint groups. State-of-the-art subspace clustering methods have shown

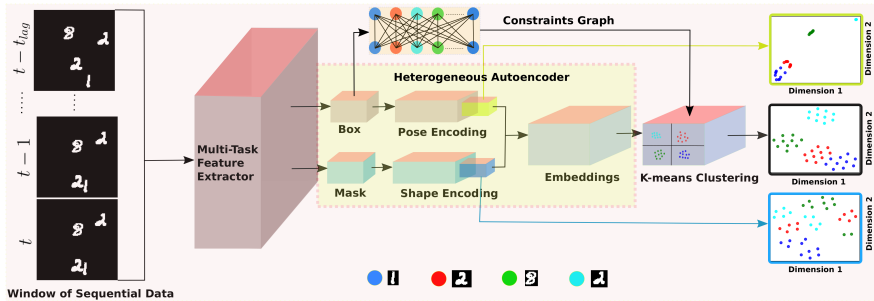


Figure 1: Proposed subspace clustering framework. The Multi-task Feature Extractor detects the bounding boxes and segmentation masks of multiple objects within a window of sequential data. The Deep Heterogeneous Autoencoder then uses these features to generate joint embedded representations of the objects. These embeddings are then clustered into target trajectories using constrained kmeans.

promising performance on single object patches [8], video sequences [22], and face tracking datasets [36]. For video sequences containing multiple objects, subspace clustering can be used as a data association strategy to assign a unique temporal identifier to each object. However, due to variations in the data distribution caused by changes in the appearance of the objects and by misdetections, occlusions, and fast motions, subspace clustering in video segments using only location [9, 12], shape [18, 29, 42, 47], or appearance [2, 36] features might not produce satisfactory results.

The goal of this work is to increase the discriminative capability of spatio-temporal latent representations. Traditional subspace clustering techniques are trained based either on appearance [18, 29, 47] or location information, generally in the form of bounding boxes [9]. Instead, in this work, we propose a novel approach that learns location and shape information jointly using a convolutional and fully connected autoencoder, which we call Deep Heterogeneous Autoencoder (DHAE). To learn a latent representation that leverages motion and appearance information in an unsupervised manner, we employ a multi-task loss function with task-dependent uncertainties [5]. Finally, we use constrained clustering techniques on the latent space to improve the robustness of spatio-temporal data association. In summary, we provide four main contributions: 1) We propose a novel unsupervised mechanism based on task-dependent uncertainties that learns to generate spatially and temporally distinctive latent features based on heterogeneous inputs; 2) We propose a new data partitioning algorithm that uses constrained clustering strategies to associate object detections over multiple frames with their corresponding temporal identifiers; 3) We evaluate our model on two synthetic and two real-world datasets that include most of the challenges commonly observed in MOTs problems, such as pose and appearance variations.

2 Related Work

Subspace clustering is an unsupervised learning technique in which data points are mapped to lower dimensionality subspaces where it is easier to make inferences about the relationships among different data points. Existing clustering methods employ two common strategies: i) extract low-dimensional discriminative features, ii) apply a robust clustering

approach to partition the subspaces. Earlier approaches employed methods based on factorization strategies [6, 23, 68] or kernels [0, 9, 61, 63] to separate the data points into their respective subspaces. More recent methods employ convolutional neural networks [0, 18, 42], or generative adversarial networks [28, 47], oftentimes in conjunction with self-expressive layers [18, 46, 47]. The autoencoder-based DSC-Net [18] uses fully connected layers to learn an affinity matrix that enhances the discriminative property of the embeddings. Some autoencoder-based techniques consider both subspace reconstruction error and cluster assignment error for better sample distribution [62, 42]. Although existing approaches may be able to find discriminative features to cluster static data, these features are not sufficiently distinctive to identify the subspaces corresponding to multiple objects in sequential data.

Clustering multi-object sequential data is an under-explored problem, particularly in real-world applications. While existing temporal clustering methods, such as ordered subspace clustering (OSC) [57] consider sequential data, they focus on clustering entire video frames, without taking into consideration the spatial aspect of the problem, which must be addressed when it is necessary to distinguish multiple objects in a video segment. Few methods [22, 66] address the problem of clustering objects over video sequences, which must take into account the fact that object features may change over time [16, 17, 27]. Our approach addresses this challenge using a simple yet effective unsupervised learning framework.

3 Subspace Clustering for Sequential Data

As Fig. 1 illustrates, our spatio-temporal clustering framework extracts features of interest from the objects in each video frame using a multi-task feature extractor module. Then, our proposed DHAE generates a discriminative latent representation of the pose and appearance of each object based on these features. To establish temporal coherence among targets, we adopt a graph-based method [44] to preclude the association of points that violate a set of constraints that are known to hold and enforce the association of points having common temporal identifiers within a temporal window. Finally, we use the constrained kmeans algorithm [40] to determine the labels of the targets by minimizing the dissimilarity of their latent representations while satisfying the association constraints. Alg. 1 summarizes the steps of the proposed method, which are described in detail below.

Algorithm 1 Subspace clustering

Input: Set of video frames $\{I^t\}_{t=1}^T$
Output: Subspace clusters C_K

- 1: **repeat**
- 2: $\mathcal{W}^t = \text{MTFE}(\{I^{t'} | t' \in \mathcal{T}^t\})$
- 3: $\mathcal{Z}^t = \text{DHAE}(\mathcal{W}^t)$
- 4: Compute \mathcal{G}^t using Eqs. (4)-(6)
- 5: $C_K = \emptyset$
- 6: $C_t = \text{kmeans}(\mathcal{Z}^t, \mathcal{G}^t)$
- 7: **for** $Q \in C_t$ **do**
- 8: $\bar{\tau} = 1/|Q| \sum_{d_i \in Q} (c_i)$
- 9: **if** $\bar{\tau} > \lambda$ **then**
- 10: $C_K = C_K \cup \{Q\}$
- 11: **end if**
- 12: **end for**
- 13: **until** end of the video sequence

3.1 Multi-Task Feature Extractor

The multi-task feature extractor (MTFE) module is responsible for generating segmentation masks and bounding boxes of objects of interest in each video frame. This task is independent of the proposed temporal clustering mechanism and can be performed by any supervised or unsupervised segmentation method such as [13, 20]. More specifically, let $x_{b,i}^t \in \mathbb{R}^N$ be the detected bounding box of the i -th target observed at time t , $x_{m,i}^t \in \mathbb{R}^{M \times M \times D}$ be a segmentation mask representing the appearance of that object, and $c_i^t \in [0, 1]$ be the corresponding

detection confidence. The MTFE takes as input a video frame I^t and generates the set

$$\mathcal{X}^t = \{[x_{m,i}^t, x_{b,i}^t, c_i^t]\}_{i=1}^{\mathcal{O}^t}, \quad (1)$$

where \mathcal{O}^t is the number of unique objects at time t . The bounding box $x_{b,i}^t$ is represented by the coordinates of its centroid and its dimensions, hence $N = 4$. Regarding appearance representation, we propose two closely related models. In the *shape* model, the number of channels of the mask $D = 1$ and $x_{m,i}^t$ corresponds to the binary segmentation mask of the object. In the *appearance* model, $D = 3$ and $x_{m,i}^t$ is given by the binary segmentation mask multiplied by the corresponding RGB contents of the image.

3.2 Deep Heterogeneous Autoencoder

Clustering methods that resort only to object appearance information do not perform well when multiple objects are observed simultaneously in a sequence of video frames. As the number of objects of a certain category (e.g., pedestrians) observed in a given frame increases, the average appearance difference among them becomes increasingly lower. At the same time, as the duration of the temporal segment increases, so does the variability in the appearance of any given target. Hence, to allow for sufficient temporal appearance variability while preventing incorrect associations among temporally proximal observations, we incorporate location information into the latent feature representation.

Fig. 2 shows our proposed DHAE architecture. To combine shape and location information, we design a network consisting of three parts: i) a pair of encoders that take as input the N -dimensional location vector¹ x_b and the $M \times M \times D$ mask x_m , ii) an uncertainty-aware module based on self-expressive layers [18] to reconstruct the concatenated feature f' and learn the latent feature \mathcal{Z} , iii) a pair of decoders to reconstruct the bounding box y_b and the mask y_m . The DHAE takes the extracted set of shapes and locations \mathcal{X}^t , which are generated by the MTFE, and reconstructs them by minimizing the combined reconstruction loss. To incorporate the location features x_b into our model, we employ a fully

connected auto-encoder (AE) with N inputs, which is represented by the yellow boxes in Fig. 2. The corresponding encoded feature vector is $f_b = h_b(x_b)$, where $h_b: \mathbb{R}^N \rightarrow \mathbb{R}^F$ is the encoding function. The shape information x_m is encoded by a convolutional auto-encoder (CAE) with an input size of $M \times M \times D$, which is represented by the blue boxes in Fig. 2. Let $f_m = h_m(x_m)$ be the latent feature of the CAE, where $h_m: \mathbb{R}^{M \times M \times D} \rightarrow \mathbb{R}^F$ is the encoding function. A function $h_a: \mathbb{R}^{2F} \rightarrow \mathbb{R}^F$ takes the concatenated feature vector $f = [f_m, f_b]$ and converts it into the latent representation $\mathcal{Z} \in \mathbb{R}^F$. A function $h'_a: \mathbb{R}^F \rightarrow \mathbb{R}^{2F}$ then takes the latent representation \mathcal{Z} and produces a new feature $f' = [f'_m, f'_b]$, which combines both shape and location information. We then process the two components of the feature

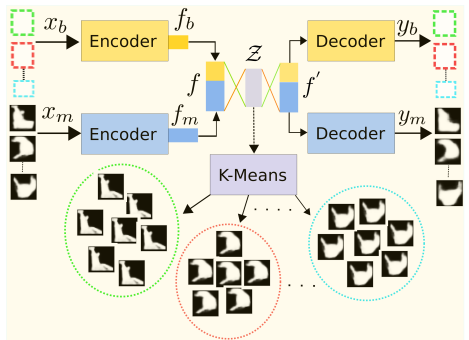


Figure 2: Proposed Deep Heterogeneous Autoencoder (DHAE) model architecture, which jointly learns shape and location information for sequential data clustering.

¹To simplify the notation, we henceforth drop the subscript i and the superscript t .

vector separately using the corresponding decoders. That is, the outputs produced by the decoding function $h'_m : \mathbb{R}^F \rightarrow \mathbb{R}^{M \times M \times D}$ and $h'_b : \mathbb{R}^F \rightarrow \mathbb{R}^N$ are $y_m \in \mathbb{R}^{M \times M \times D}$ and $y_b \in \mathbb{R}^N$.

3.2.1 Multi-task Likelihood

We train our DHAE using a multi-task loss function using maximum log-likelihoods with task-dependent uncertainties. Let $f^W(x) = [f_m^W(x), f_b^W(x)]$ be the output of the DHAE with weights W , input vector $x = [x_m, x_b]$, and predicted output $y = [y_m, y_b]$. The likelihood for the regression task is given by

$$p(y_m, y_b | f^W(x), \sigma_m, \sigma_b) = p_m(y_m | f_m^W(x), \sigma_m) \cdot p_b(y_b | f_b^W(x), \sigma_b), \quad (2)$$

where y_b and y_m are normally distributed with means $f_b^W(x)$, $f_m^W(x)$, and variances σ_b , σ_m , respectively. Thus, the cost function is given by

$$\begin{aligned} \mathcal{L}(W, \sigma_m, \sigma_b) &= -\log p(y_m, y_b | f^W(x), \sigma_b, \sigma_m) \\ &\propto \frac{1}{2\sigma_b^2} \|y_b - f_b^W(x)\|^2 + \frac{1}{2\sigma_m^2} \|y_m - f_m^W(x)\|^2 + \log \sigma_m + \log \sigma_b, \end{aligned} \quad (3)$$

where W , σ_m , and σ_b are trainable parameters. Unlike traditional multi-task learning approaches [9, 5], which focus on weighing the contributions of several *homogeneous outputs*, our method weighs the contribution of multiple *heterogeneous input* features. Thus, we learn the relative weights of each loss function term based on the uncertainty of distinct features.

3.2.2 Network Implementation and Training Details

The AE branch of our DHAE has an input size of $N = 4$ and one fully connected layer of size 128. The CAE branch uses 5 convolutional layers with kernel size 3×3 , ReLU activations, and a stride of 2×2 for downsampling and upsampling [4]. The size of the input layer is $M = 128$ and the subsequent layers have half the size of the previous layer. The number of convolutional channels in each layer is 16, 16, 32, 32, and 64 (the decoder mirrors the structure of the encoder). The functions $h_a(\cdot)$ and $h'_a(\cdot)$ are implemented using fully connected layers of size $F = 128$. We train the network using stochastic gradient descent with ADADELTA [45] learning rate adaptation. To account for the dimensionality of the feature vectors, we initialize $\log \sigma_b^2 = 1/N$ and $\log \sigma_m^2 = 1/(M^2)$. The network weights W are initialized using the Glorot method [46]. At inference time, the segmentation masks are isotropically scaled such that the largest dimension of the mask is M . The image is then centered along the smallest dimension and zero-padded.

3.3 Sequential Data Constraints

In spatio-temporal subspace clustering, constraints based on prior knowledge regarding the sequential data may reduce association mistakes by imposing penalties on unlikely pairwise matching. We use an undirected graph \mathcal{G}^t to encode constraints among pairs of detections in the temporal window \mathcal{T}^t . This graph determines which pairs of detections cannot belong to the same cluster. It also enforces the association of detections that were assigned the same temporal identifier in previous temporal windows.




Figure 3: Sequence of frames for (a) MNIST-MOT with $|O| = 5$ targets per frame and (c) MOTSChallenge (cropped for three pedestrians) sequences. (b) and (d) show the corresponding subspaces generated by our method.

3.3.1 Constraints Graph Formulation

To incorporate prior knowledge regarding object correspondences in a temporal window, we construct the graph $\mathcal{G}^t = (V^t, E^t)$ using *cannot link* and *must link* constraints. The vertices of the graph correspond to the set of segmentation masks in all the frames in the window \mathcal{T}^t , i.e.,

$$V^t = \{x_{m,i}^t | t \in \mathcal{T}^t, i \in \{1, \dots, \mathcal{O}^t\}\}. \quad (4)$$

The set of edges consists of pairs of nodes v_i and v_j that meet the spatial and temporal restrictions imposed by the *cannot link function* $f_{cl}(\cdot)$ and *must link function* $f_{ml}(\cdot)$, i.e.,

$$E_{cl}^t = \{(v_i, v_j) | v_i \in V^t, v_j \in V^t, f_{cl} = 1\}, \quad E_{ml}^t = \{(v_i, v_j) | v_i \in V^t, v_j \in V^t, f_{ml} = 1\}. \quad (5)$$

The function f_{cl} prevents the association of vertices from the same frame or vertices in close temporal proximity whose segmentation masks do not overlap. Conversely, f_{ml} enforces the association of detections with common temporal identifiers. That is, if a detection has been assigned an identifier at a previous temporal window, it is only allowed to be associated with detections that have the same identifier or that have not been assigned one, i.e.,

$$f_{cl}(v_i, v_j) = \begin{cases} 1 & \text{if } t_i = t_j \text{ or } \gamma_i \neq \gamma_j \\ 1 & \text{if } \text{iou}(v_i, v_j) = 0, t_i - t_j \leq \tau, \\ 0 & \text{otherwise} \end{cases}, \quad f_{ml}(v_i, v_j) = \begin{cases} 1 & \text{if } l_i = l_j, \gamma_i = \gamma_j, \\ & t_i \neq t_j, l \neq \emptyset \\ 0 & \text{otherwise} \end{cases}, \quad (6)$$

where t_i, t_j are the timestamps for the detections corresponding to v_i and v_j , the function $\text{iou}(\cdot)$ computes their mask intersection over union, γ_i, γ_j are the corresponding object classes, and l_i, l_j are the initialized cluster identities corresponding to nodes v_i and v_j .

3.3.2 Modified Constrained Kmeans

In the constrained kmeans algorithm [14], a node v_i may be assigned to the same cluster as v_j only if the edge $(v_i, v_j) \notin E_{cl}^t$ and the initialized detections maintain the same subspace clustering when $(v_i, v_j) \in E_{ml}^t$. We cluster objects that do not satisfy the *cannot link* constraints and whose tracklet identities are not yet initialized by minimizing the distance between their corresponding latent features z_i^t and the cluster centroids z_k .

The dimensionality of the subspace corresponding to window \mathcal{T}^t is given by the total number of objects observed during that period, which is unknown. We propose a novel

mechanism to determine the number of clusters $|\mathcal{K}|$ that leverages our temporal clustering constraints. Our approach consists of initially setting the number of clusters to the maximum number of targets observed in a single frame within the window \mathcal{T}^t , i.e., $|\mathcal{K}| = \max_{t' \in \mathcal{T}^t} \mathcal{O}^{t'}$. The centroids of the clusters are then initialized using the detections in $\mathcal{O}^{t'}$. Then, if the clustering constraints cannot be satisfied, the number of clusters is adjusted accordingly. That is, if a detection does not satisfy the constraints in Eq. 6, it is considered a tentative new target, a new cluster is created, and as new detections join that cluster in subsequent frames, a new tracklet is generated. If no subsequent detections join the cluster within a t_{lag} interval, the new detection is considered a mistake and the corresponding cluster is discarded.

3.4 Spatio-temporal Clustering

To assign unique temporal identifiers to multiple objects in a video sequence, at each time instant t , we cluster the observations present in the frames within the window $\mathcal{T}^t = \{t - t_{lag}, t - t_{lag} + 1, \dots, t\}$ (Fig. 3). Since the window is computed at each frame, there is an overlap of $t_{lag} - 1$ frames between subsequent windows, which ensures that most of the data points used to form the subspace cluster in each window of a video sequence are shared. As shown in Alg. 1, the input to our spatio-temporal clustering algorithm is the set of frames $\{I^t\}_{t=1}^T$ where T is the number of frames in the video. Our algorithm applies the MTFE to each video frame in the window \mathcal{T}^t to construct the set $\mathcal{W}^t = \{\mathcal{X}^{t'} | t' \in \mathcal{T}^t\}$, where $\mathcal{X}^{t'}$ is the output of the MTFE for frame $I^{t'}$. The algorithm then clusters the detections within each window using the embeddings $\mathcal{Z}^t = \{z_i^{t'} | t' \in \mathcal{T}^t, i \in \{1, \dots, \mathcal{O}^{t'}\}\}$ generated by the DHAe. Each call to the `kmeans`(\cdot) algorithm produces a set of clusters \mathcal{C}_t whose elements are detections assigned to the same object. For each cluster $Q \in \mathcal{C}_t$, we compute its normalized score $\bar{\tau}$ as the ratio $1/|Q| \sum_{d_i \in Q} (c_i)$, where c_i is the confidence score of detection d_i . Clusters with a score higher than a threshold λ are included in the cluster set \mathcal{C}_K .

4 Datasets and Experiments

We evaluate our algorithm on two synthetic and two real-world datasets. The MNIST-MOT and Sprites-MOT [14] synthetic datasets allow us to simulate challenging MOTS scenarios involving pose, scale, and shape variations. In addition, their bounding box and segmentation masks are readily available. Then, we use the recently published MOTS [59] benchmark, which includes video sequences from the MOTChallenge [6] and the KITTI [10] datasets annotated with segmentation masks, to evaluate our model in real-world videos.

Since traditional clustering measures [15, 64, 47] require each observation to be mapped to exactly one cluster, they are not suitable for real-world scenarios where incorrect detections may occur. Hence, we adopt the popular CLEAR-MOT [19] tracking performance assessment measures: Multi-object Tracking Accuracy (MOTA), Fragmentation (Frag), Identity Switches (IDs), Mostly Tracked (MT), and Mostly Lost (ML) targets [40]. Finally, we employ the MOTS [59] evaluation measures to quantify the effectiveness of our algorithm in maintaining the temporal consistency of target identities in real-world video sequences.

4.1 Synthetic Datasets

We generate synthetic MNIST-MOT and Sprites-MOT sequences using the procedure described in [24], which includes most of the common challenges observed in MOT problems. For the MNIST-MOT dataset, we generate 9 digit classes and for Sprites-MOT we generate 4 geometric shapes. For both datasets, the object density is $|O| = 3$, the target birth probability is 0.5, the size of each object is 28×28 at each frame of size 128×128 , and the average target velocity is 5.3 pixels per frame. We generate 20 video sequences with 500 frames for each dataset. In each sequence, the set of initial objects (digits or sprites) is chosen randomly at the first frame. In subsequent frames, the objects move in random directions.

To train the DHAE, we extract the bounding box and shape mask of each object from the synthetic video frames using a separate set of training sequences. Due to the lack of availability of methods that perform clustering based on location and shape features, we select one state-of-the-art MOT method for performance comparison [24]. We relax the IoU constraint by imposing a limit on the Euclidean distance between embeddings since the target displacement among consecutive frames may be relatively large with respect to the size of the targets due to the low resolution of the frames.

Table 1: Performance evaluation on MNIST-MOT and Sprites-MOT with $t_{lag} = 3$, $|O| = 3$.

Method	MNIST-MOT							Sprites-MOT						
	↑IDF1	↑MT	↓ML	↓FN	↓IDs	↓Frag	↑MOTA	↑IDF1	↑MT	↓ML	↓FN	↓IDs	↓Frag	↑MOTA
shape embed	89.1	944	0	304	5	132	98.6	86.7	906	4	853	13	275	96.1
loc embed	87.7	905	0	708	61	331	96.5	88.2	920	0	689	56	347	96.6
loc+ \mathcal{G}'	86.3	977	0	0	72	0	99.7	85.6	983	0	4	94	0	99.6
loc+shape	89.6	934	0	444	3	189	98.0	88.9	912	0	734	8	314	96.6
loc+shape+ \mathcal{G}'	100.0	977	0	0	0	0	100.0	99.5	983	0	45	17	0	99.7
TBA [24]	99.6	978	0	49	22	7	99.5	99.2	985	1	80	30	22	99.2

Table 1 summarizes the performance of our method on the MNIST-MOT and Sprites-MOT datasets according to the evaluation procedure described in [24, 29, 41]. Although location features play a critical role in clustering multiple moving targets, shape features also contribute significantly to the performance of our approach. As the t-SNE visualization in Fig. 1 illustrates, as the digits 1 and 2 approach each other, their embeddings remain separable. The shape-only model (*shape embed*) outperforms the location-only model (*loc embed*) on some of the evaluation criteria because the shapes remain unchanged until they leave the scene. This effect is more pronounced on the MNIST-MOT dataset because the appearance of the characters is more distinctive than the shapes in Sprites-MOT. Although the incorporation of the constraints graph into the location-only model (*loc+ \mathcal{G}'*) leads to slightly better performance on some evaluation measures than the joint embedding (*loc+shape*), the overall method (*loc+shape+ \mathcal{G}'*) achieves near-perfect results.

4.2 MOTs Dataset

The MOTsChallenge dataset consists of four fully annotated videos of crowded scenes and the KITTI MOTs dataset consists of 21 videos acquired from a moving vehicle. Both datasets contain objects that show substantial scale and shape variations over time. We use the segmentation masks and the corresponding RGB content from 12 KITTI MOTs sequences to train the DHAE and use the remaining sequences for testing.

In our evaluation, we use the publicly available instance segmentation masks and bounding boxes [25] from the benchmark validation set. We compare the performance of our

Table 2: Evaluation of person and car tracking on the KITTI MOTS validation set.

Method	\uparrow sMOTSA		\uparrow MOTSA		\uparrow MOTSP	
	car	ped	car	ped	car	ped
loc+shape+ \mathcal{G}_t	80.4	55.5	89.5	69.7	90.3	82.8
loc+app+\mathcal{G}_t	80.8	58.3	89.9	72.5	90.3	82.8
EagerMOT [20]	74.5	58.1	-	-	-	-
GMPHD [26]	76.9	48.8	-	-	87.1	76.4
MOTSFusion [43]	77.5	49.9	89.2	66.6	-	-
MOTSPNet [30]	78.1	54.6	87.2	69.3	89.6	79.7
TR-CNN [39]	76.2	46.8	87.8	65.1	87.2	75.7

Table 3: Evaluation of person tracking on the MOTChallenge training set.

Method	\uparrow sMOTSA	\uparrow MOTSA	\uparrow MOTSP	\uparrow MT
	loc+shape	51.3	60.1	86.3
loc+shape+ \mathcal{G}_t	65.3	76.3	86.3	53.1
loc+app	56.1	65.5	86.4	26.3
loc+app+ \mathcal{G}_t	65.5	76.5	86.3	53.1
GMPHD [26]	65.8	77.1	86.1	-
PointTrack [43]	58.1	70.6	-	-
MOTSPNet [30]	56.8	69.4	82.7	-
TR-CNN [39]	52.7	66.9	80.2	-




Figure 4: Qualitative results on the (a) KITTI MOTS and (b) MOTChallenge dataset.

method against the state-of-the-art approaches presented in [20, 25, 26, 30, 39, 43]. Table 2 shows that, without resorting to sophisticated mechanisms for target re-identification, trajectory interpolation, or entry/exit detection, our method outperforms all the baseline methods in the KITTI MOTS dataset even if only the binary segmentation masks are used in the joint embeddings ($loc+shape+\mathcal{G}_t$). Including the RGB information ($loc+app+\mathcal{G}_t$) leads to further performance gains, particularly for the pedestrian class. As Table 3 indicates, in the MOTChallenge sequences, our joint embeddings alone ($loc+app$) perform on par with [30, 39, 43], and the incorporation of the constraints graph leads to results comparable to [26] using the same set of detections (whereas [43] and [30] use privately refined detections). Figure 4 illustrates some of the results generated by our method. Both datasets are comprised of crowded scenes with significant amounts of temporary partial and full occlusions, particularly among pedestrians. Unlike [26], our method does not incorporate occlusion reasoning or motion modelling techniques, which contribute significantly to the performance.

To assess the impact of detection noise on the performance of our method, we evaluate the sMOTSA measure as a function of the minimum confidence score for a detection to be considered valid. Fig. 5 shows that performance increases until the detection threshold reaches approximately 0.65 and after 0.80 it starts to decrease again. In the experiments discussed above, we use a detection threshold of 0.70 for all the datasets. Further performance improvements would be achieved with dataset-specific thresholds. For our ablation studies, we use ground truth annotations instead to evaluate each step of the method independently from the performance of the underlying detector.




Figure 5: MOTS performance as a function of MTFE detection score threshold.

4.3 Ablation Study

Table 4 shows the impact of the uncertainty-aware multitask learning loss of Eq. 3, of the constraints graph \mathcal{G}^l , and of estimating the number of clusters $|\mathcal{K}|$ using the method described in Section 3.3.2 for three different window sizes t_{lag} . The table demonstrates the positive impact of multi-task learning using task uncertainties instead of assigning equal weights in Eq. 3. We observe that the constraints graph leads to consistent and substantial improvements in all the evaluation criteria. We also see that higher values of t_{lag} lead to an increase in the MT measure but also to an increased number of fragmentations. Since we do not model target motion, the overlap among detections reduces as t_{lag} increases, leading to violations of the must-link constraints. As a result, the optimal sMOTSA score is obtained with $t_{lag} = 3$. Finally, estimating $|\mathcal{K}|$ from the video segments does not degrade the performance of our algorithm. In summary, Table 4 shows that MTL and \mathcal{G}^l lead to improvements in the sMOTSA measure of 10.8% and 15.5% in the MOTSChallenge, 8.2% and 14.1% for the car category in the KITTI MOTS, and 11.1% and 26.2% for the person class, even when $|\mathcal{K}|$ is unknown.

Table 4: Ablation study for different components of our method. We evaluate the impact on tracking based on MOTS [69] oracle performance of: window size t_{lag} , constraints graph \mathcal{G}^l , multi-task learning (MTL), and number of subspaces $|\mathcal{K}|$ as a prior (\times) or estimated (\checkmark).

t_{lag}	\mathcal{G}^l	MTL	$ \mathcal{K} $	MOTSChallenge				KITTI MOTS (car)				KITTI MOTS (person)			
				\uparrow sMOTSA	\uparrow MT	\downarrow IDs	\downarrow Frag	\uparrow sMOTSA	\uparrow MT	\downarrow IDs	\downarrow Frag	\uparrow sMOTSA	\uparrow MT	\downarrow IDs	\downarrow Frag
3	\times	\times	\checkmark	74.0	43.9	962	1835	77.5	63.6	217	479	62.7	57.4	255	318
	\times	\checkmark	\checkmark	82.6	62.3	660	1416	84.4	68.2	115	348	70.5	54.4	138	258
	\checkmark	\times	\times	84.0	68.9	653	1367	84.7	70.9	113	343	71.7	52.9	138	248
	\checkmark	\checkmark	\times	97.5	100	636	639	98.1	98.7	127	127	95.6	98.7	119	115
5	\checkmark	\checkmark	\checkmark	97.8	99.6	548	548	98.3	98.7	119	121	95.6	98.5	124	124
	\times	\times	\checkmark	66.4	31.1	939	2000	75.5	58.9	191	474	53.0	52.9	226	282
	\times	\checkmark	\checkmark	78.9	50.9	606	1476	76.7	56.3	97	408	63.0	32.4	100	238
	\times	\checkmark	\times	80.9	58.3	610	1422	77.4	55.6	94	409	63.2	36.8	92	254
8	\checkmark	\checkmark	\times	97.2	100	692	733	97.7	98.7	160	171	94.1	100	186	199
	\checkmark	\checkmark	\checkmark	97.7	100	581	589	97.8	98.7	162	162	94.6	100	180	181
	\times	\times	\checkmark	56.9	19.7	897	1800	71.4	50.3	218	493	46.0	38.2	141	227
	\times	\checkmark	\checkmark	71.8	41.7	447	1499	73.0	47.7	79	376	58.2	30.9	83	246
8	\times	\checkmark	\times	73.6	46.9	528	1440	74.8	47.0	77	384	57.5	35.3	82	250
	\checkmark	\checkmark	\times	97.0	99.1	765	800	96.8	98.0	223	236	93.5	100	215	217
	\checkmark	\checkmark	\checkmark	97.5	100	657	662	97.2	98.0	210	210	94.1	100	194	196

5 Conclusions

Our proposed method uses task-dependent uncertainties to simultaneously learn the contribution of shape/appearance and location features from multi-object video datasets and further improves clustering performance by imposing simple constraints on acceptable sequential data patterns. Our experimental results show that our approach can accurately cluster multiple objects using embeddings generated by the DHAE. This method can be extended without significant modifications to include additional tasks of interest in similar scenarios such as object motion prediction. In the future, we intend to extend our method with target motion models and more robust entry/exit/occlusion detection techniques so that it can be employed as a robust, standalone data association mechanism to the problems of multiple object tracking [65] and video instance segmentation [24].

Acknowledgements

This work was 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.

References

- [1] M. Bäuml, M. Tapaswi, and R. Stiefelhagen. Semi-supervised learning with constraints for person identification in multimedia data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3602–3609, 2013.
- [2] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European Conference on Computer Vision*, pages 139–156, 2018.
- [3] Bin Cheng, Guangcan Liu, Jingdong Wang, Zhongyang Huang, and Shuicheng Yan. Multi-task low-rank affinity pursuit for image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2439–2446, 2011.
- [4] Yizong Cheng. Mean shift, mode seeking, and clustering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(8):790–799, 1995.
- [5] Roberto Cipolla, Yarin Gal, and Alex Kendall. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7482–7491, 2018.
- [6] Patrick Dendorfer, Aljoša Ošep, Anton Milan, Konrad Schindler, Daniel Cremers, Ian Reid, Stefan Roth, and Laura Leal-Taixé. MOTChallenge: A benchmark for single-camera multiple target tracking. *International Journal of Computer Vision*, 129: 845–881, 2020.
- [7] Vincent Dumoulin and Francesco Visin. A guide to convolution arithmetic for deep learning. *arXiv preprint arXiv:1603.07285*, 2018.
- [8] Ehsan Elhamifar and Rene Vidal. Sparse subspace clustering: Algorithm, theory, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(11): 2765–2781, 2013.
- [9] Chelsea Finn, Xin Yu Tan, Yan Duan, Trevor Darrell, Sergey Levine, and Pieter Abbeel. Deep spatial autoencoders for visuomotor learning. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 512–519, 2016.
- [10] Andreas Geiger. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012.

- [11] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feed-forward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 249–256, 2010.
- [12] Yi Guo, Junbin Gao, and Feng Li. Spatial subspace clustering for drill hope spectral data. *Journal of Applied Remote Sensing*, 8:083644, 2014.
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask R-CNN. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2980–2988, 2017.
- [14] Zhen He, Jian Li, Daxue Liu, Hangen He, and David Barber. Tracking by animation: Unsupervised learning of multi-object attentive trackers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1318–1327, 2019.
- [15] Lawrence Hubert and Phipps Arabie. Comparing partitions. *Journal of Classification*, 2(1):193–218, 1985.
- [16] Reza Jalil Mozhdehi and Henry Medeiros. Deep convolutional likelihood particle filter for visual tracking. In *The International Conference on Image Processing, Computer Vision, and Pattern Recognition*, pages 27–38, 2020.
- [17] Reza Jalil Mozhdehi, Yevgeniy Reznichenko, Abubakar Siddique, and Henry Medeiros. Deep convolutional particle filter with adaptive correlation maps for visual tracking. In *Proceedings of the 25th IEEE International Conference on Image Processing*, pages 798–802, 2018.
- [18] Pan Ji, Tong Zhang, Hongdong Li, Mathieu Salzmann, and Ian Reid. Deep subspace clustering networks. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 23–32, 2017.
- [19] Rainer Stiefelhagen Keni Bernardin. Evaluating multiple object tracking performance: The CLEAR MOT metrics. *EURASIP Journal on Image and Video Processing*, 2008, 2008.
- [20] Aleksandr Kim, Aljoša Ošep, and Laura Leal-Taixé. Eagermot: 3d multi-object tracking via sensor fusion. In *Proceedings of the IEEE International Conference on Robotics and Automation*, 2021.
- [21] Boah Kim and Jong Chul Ye. Mumford–Shah loss functional for image segmentation with deep learning. *IEEE Transactions on Image Processing*, 29:1856–1866, 2019.
- [22] Prakhar Kulshreshtha and Tanaya Guha. An online algorithm for constrained face clustering in videos. In *Proceedings of the 25th IEEE International Conference on Image Processing*, pages 2670–2674, 2018.
- [23] Guangcan Liu, Zhouchen Lin, and Yong Yu. Robust subspace segmentation by low-rank representation. In *International Conference on Machine Learning*, pages 663–670, 2010.

- [24] Jonathon Luiten, Philip Torr, and Bastian Leibe. Video instance segmentation 2019: A winning approach for combined detection, segmentation, classification and tracking. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019.
- [25] Jonathon Luiten, Tobias Fischer, and Bastian Leibe. Track to reconstruct and reconstruct to track. *IEEE Robotics and Automation Letters*, 5(2):1803–1810, 2020.
- [26] Young min Song and Moongu Jeon. Online multi-object tracking and segmentation with GMPHD filter and simple affinity fusion. *arXiv preprint arXiv:2009.00100*, 2020.
- [27] Reza Jalil Mozhdehi, Yevgeniy Reznichenko, Abubakar Siddique, and Henry Medeiros. Convolutional adaptive particle filter with multiple models for visual tracking. In *International Symposium on Visual Computing*, pages 474–486, 2018.
- [28] Sudipto Mukherjee, Himanshu Asnani, Eugene Lin, and Sreeram Kannan. Clustergan: Latent space clustering in generative adversarial networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4610–4617, 2019.
- [29] Christian Payer, Darko Štern, Marlies Feiner, Horst Bischof, and Martin Urschler. Segmenting and tracking cell instances with cosine embeddings and recurrent hourglass networks. *Medical Image Analysis*, 57:106–119, 2019.
- [30] Lorenzo Porzi, Markus Hofinger, Idoia Ruiz, Joan Serrat, Samuel Rota Bulò, and Peter Kotschieder. Learning multi-object tracking and segmentation from automatic annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6845–6854, 2020.
- [31] Marc’Aurelio Ranzato, Christopher Poultney, Sumit Chopra, and Yann LeCun. Efficient learning of sparse representations with an energy-based model. In *Proceedings of the 19th International Conference on Neural Information Processing Systems*, page 1137–1144, 2006.
- [32] Yevgeniy Reznichenko, Enrico Prampolini, Abubakar Siddique, Henry Medeiros, and Francesca Odone. Visual tracking with autoencoder-based maximum a posteriori data fusion. In *The IEEE Computer Society Computers, Software, and Applications Conference*, pages 501–506, 2019.
- [33] Peter J. Rousseeuw. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987.
- [34] Alexander Strehl and Joydeep Ghosh. Cluster ensembles — a knowledge reuse framework for combining multiple partitions. *The Journal of Machine Learning Research*, 3:583–617, 2003.
- [35] ShiJie Sun, Naveed Akhtar, HuanSheng Song, Ajmal Mian, and Mubarak Shah. Deep affinity network for multiple object tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(1):104–119, 2021.
- [36] Makarand Tapaswi, Marc Law, and Sanja Fidler. Video face clustering with unknown number of clusters. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5026–5035, 2019.

- [37] Stephen Tierney, Junbin Gao, and Yi Guo. Subspace clustering for sequential data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1019–1026, 2014.
- [38] René Vidal. Subspace clustering. *IEEE Signal Processing Magazine*, 28(2):52–68, 2011.
- [39] Paul Voigtlaender, Michael Krause, Aljosa Osep, Jonathon Luiten, Berin Balachandar Gnana Sekar, Andreas Geiger, and Bastian Leibe. Mots: Multi-object tracking and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7934–7943, 2019.
- [40] Kiri Wagstaff, Claire Cardie, Seth Rogers, and Stefan Schrödl. Constrained K-means clustering with background knowledge. In *International Conference on Machine Learning*, pages 577–584, 2001.
- [41] Bo Wu and Ram Nevatia. Tracking of multiple, partially occluded humans based on static body part detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 951–958, 2006.
- [42] Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *International Conference on Machine Learning*, pages 478–487, 2016.
- [43] Zhenbo Xu, Wei Zhang, Xiao Tan, Wei Yang, Huan Huang, Shilei Wen, Errui Ding, and Liusheng Huang. Segment as points for efficient online multi-object tracking and segmentation. In *Proceedings of the European Conference on Computer Vision*, pages 264–281, 2020.
- [44] Amir Roshan Zamir, Afshin Dehghan, and Mubarak Shah. GMCP-Tracker: Global multi-object tracking using generalized minimum clique graphs. In *Proceedings of the European Conference on Computer Vision*, pages 343–356, 2012.
- [45] Matthew D Zeiler. ADADELTA: an adaptive learning rate method. *arXiv preprint arXiv:1212.5701*, 2012.
- [46] Tong Zhang, Pan Ji, Mehrtash Harandi, Richard Hartley, and Ian Reid. Scalable deep k-subspace clustering. In *Proceedings of the Asian Conference on Computer Vision*, pages 466–481, 2018.
- [47] Pan Zhou, Yunqing Hou, and Jiashi Feng. Deep adversarial subspace clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1596–1604, 2018.