Dynamic character graph via online face clustering for movie analysis



Prakhar Kulshreshtha¹ · Tanaya Guha²

Received: 27 September 2019 / Revised: 16 July 2020 / Accepted: 28 July 2020 / Published online: 31 August 2020 © The Author(s) 2020

Abstract

An effective approach to automated movie content analysis involves building a network (graph) of its characters. Existing work usually builds a static character graph to summarize the content using metadata, scripts or manual annotations. We propose an unsupervised approach to building a *dynamic* character graph that captures the temporal evolution of character interaction. We refer to this as the *character interaction graph* (CIG). Our approach has two components: (i) an online face clustering algorithm that discovers the characters in the video stream as they appear, and (ii) simultaneous creation of a CIG using the temporal dynamics of the resulting clusters. We demonstrate the usefulness of the CIG for two movie analysis tasks: narrative structure (acts) segmentation and major character retrieval. Our evaluation on full-length movies containing more than 5000 face tracks shows that the proposed approach achieves superior performance for both the tasks.

Keywords Online clustering \cdot Face clustering \cdot Media content analysis \cdot Character graph \cdot Narrative structure

1 Introduction

Automated analysis of media content, such as movies has traditionally focused on extracting and using low level features from shots and scenes for analyzing narrative structures and key events [10, 11]. For humans, however, a movie is not just a collection of shots or scenes. It is the characters that usually play the most important role in storytelling [18]. More recently, character-centric representation of movies, such as character networks have emerged as

Prakhar Kulshreshtha pkulshre@andrew.cmu.edu

- ¹ Carnegie Mellon University, Pittsburgh, USA
- ² University of Warwick, Coventry, UK

A part of this work was done when both the authors were at IIT Kanpur.

Tanaya Guha tanaya.guha@warwick.ac.uk

an effective approach towards media content analysis [15, 16, 22]. A character network usually has the major characters as its nodes where the edges summarize the relationship between character pairs. Such networks have been shown to facilitate a number of movie analysis tasks including character analysis [16], story segmentation [22] and major character identification [15]. The existing methods build a single, static character network for the entire movie. While static graphs offer a convenient summary of the overall interactions among characters, they can not capture the evolution of a movie's dynamic narrative.

In this paper, we present an unsupervised approach to building a *dynamic* character network via online face clustering. We refer to this network as the *character interaction graph* (CIG), where each movie character is represented as a node, and an edge represents pairwise interaction between characters. The dynamic aspect of the CIG offers an effective way to capture the variations in character interactions over time - particularly helpful for story segmentation and event localization. Our approach (see Fig. 1 for an overview) has of two main components - online face clustering, and simultaneous creation of the CIG using the resulting clusters. Building on our previous work on online face clustering [8], we develop a new algorithm to create (and update) a CIG via clustering i.e., utilizing the information from the cluster dynamics. We demonstrate the usefulness of the CIGs for two important movie analysis tasks: (i) semantic segmentation of a movie into *acts*, and (ii) major character discovery. Performance is evaluated on a database of six full-length Hollywood movies containing more than 5000 face tracks. Results are compared with relevant past work and manual annotations, where our CIG-based approach shows superior performance.

In summary, the contribution of this work is two-fold: (i) We propose an unsupervised approach to building dynamic character graph via online face clustering. This is the first work on dynamic CIG construction. (ii) We demonstrate superior performance of our CIG-based approach for two important movie analysis tasks - three act segmentation and major character identification.



Fig. 1 Overview of the proposed approach: A movie is processed at shot level. For each shot, face tracks are created, and our online clustering algorithm either groups a face track with an existing cluster or creates a new one. In this example, at shot0, three face tracks are created and grouped into two clusters. A new cluster is added in shot1 as it belongs to a new character. The face track in shot2 is added to an existing cluster belonging to the same character. The CIG is updated after each shot is processed. Note that the CIG for the $(i - 1)^{th}$ shot is obtained after the $(i - 1)^{th}$ shot is processed

The rest of this paper is organized as follows. Section 2 discusses relevant literature for character network-based movie analysis and online face clustering. Section 3 describes our approach to dynamic CIG creation via online face clustering. Section 4 proposes the methodologies to apply CIG for two movie analysis tasks. Section 5 presents extensive results, and Section 6 concludes the article with summary and discussion on future work.

2 Related work

Following the two major components of our approach, we discuss past work related to character graph construction and face clustering in multimedia content.

2.1 Character network construction

Character networks are useful for multimedia content analysis due to its wide applicability in story summarization, segmentation, character identification, and character-based search and indexing [15, 16, 21, 22]. Character networks have been constructed using movie scripts [16], spoken dialogs [15], manually-labeled data [22], and supervised face recognition [22].

Ramakrishna et al. [16] used scripts to construct a character network, where an edge between two characters (nodes) is added if the characters have consecutive dialogs. This network is used to examine the character analytics based on gender, race and age [16]. Weng et al. [22] constructed a character network, called the RoleNet, that captures the co-occurrence statistics of movie characters via face recognition. This network is used to identify the lead characters and communities, and for story segmentation [22]. Park et al. [15] built a network aligning scripts and subtitles. This network is employed in classification of major and minor characters, community clustering and sequence detection [15]. Along the similar lines, Tran and Jung constructed a CoCharNet [21] using manual annotations to encode information regarding character co-occurrences.

The work most related to our work is that of Yeh and Wu [28], where character network is constructed using face clustering. This work clusters faces and constructs a character network in an iterative fashion. However, this requires prior knowledge of the number of clusters, and is an offline method. To the best of our knowledge, this is the only prior work that uses (offline) face clustering for constructing character graphs.

2.2 Face clustering in videos

Offline methods. The problem of unsupervised face clustering is relatively less studied as compared to its supervised counterpart, i.e., face recognition. The dominant approach to face clustering involves learning a suitable distance measure between face pairs [9, 17, 20, 30]. Several methods have proposed to use partial supervision to improve performance [3, 23]. While image-based clustering is more common, face clustering in videos can achieve significant improvement by exploiting the temporal information about the faces [1, 25–27]. Temporal constraints have been used in frameworks based on hidden Markov random field (HMRF) [26] and unsupervised logistic discriminative metric learning (ULDML) [2] with applications to face clustering in movies and TV series. A constrained multiview face clustering technique used constrained sparse subspace representation of faces with constrained spectral clustering [1]. Recent clustering approaches use convolutional neural networks (CNN) to learn robust face representations by using aggregated

deep features [19], deep features with pairwise constraints [30], and deep features with triplet loss [29].

Online methods. The approaches discussed above are all *offline* methods i.e., they assume the availability of the entire data at once. In an online setting, a clustering algorithm does not have the luxury of 'seeing' the entire data simultaneously. To the best of our knowledge, there is only one work on online face clustering in videos in the existing literature [14]. This work created small tracklets of faces from the video, and clustered them in an online fashion based on temporal coherence and the Chinese restaurant process (TCCRP) [14]. An extension of this work is Temporally Coherent Chinese Restaurant Franchise (TCCRF) [13], that jointly models short temporal segments. These online methods tend to create multiple clusters for the same person thereby degrading the completeness of the clusters [14].

3 Proposed approach

Overview. In our dynamic CIG construction approach, we process a movie stream at shotlevel, where a shot is a contiguously recorded sequence of frames. Our approach consists of two main components: (i) face track creation and clustering, and (ii) CIG formation and update. All the components are executed simultaneously in an *online* fashion by processing one shot at a time. As a shot appears, all faces are detected frame by frame and face tracks are created. Our online clustering algorithm then assigns the face tracks to either an existing cluster or to a new one. The information about the cluster updates, including formation of new clusters, are used to create a dynamic CIG. Figure 1 presents an overview of the proposed method. Below, we describe each component in detail.

3.1 Face track creation and clustering

Face track creation. Consider a movie \mathcal{M} comprising T frames: $\mathcal{M} = \{I_t\}_{t=1}^T$. We define the *i*th shot S_i as a sequence of consecutive frames $\{I_{t_{(i-1)}+1}, \ldots, I_{t_i}\}$, where t_i is the *i*th shot boundary. The shot boundary t_i corresponding to S_i is detected by computing the pixel differences between consecutive frames (as they appear) and by comparing the difference to a predefined threshold. The accuracy of shot boundary detection is not critical to the performance of our method, hence we stick to this simple frame differencing method.

Once we have detected the boundaries of S_i , a standard face detector [7] is employed to detect the faces in each frame in S_i . This frame level face detection can be done in parallel to searching for shot boundaries. The face detector returns the bounding boxes of each face detected in every frame. To build a robust representation of these faces, we use a pretrained CNN, called the FaceNet [17]. Each face \mathbf{f}^p is forward-passed through the FaceNet to obtain its corresponding *d* dimensional feature vector \mathbf{v}^p .

To create face tracks, we use a simple yet effective strategy to combine the faces detected in consecutive frames. Let us define two faces detected in two consecutive frames as \mathbf{f}^p and \mathbf{f}^q . The overlap $a(\cdot)$ between the two faces is defined as:

$$a(p,q) = \frac{\operatorname{area}(\mathbf{f}^p \cap \mathbf{f}^q)}{\max(\operatorname{area}(\mathbf{f}^p), \operatorname{area}(\mathbf{f}^q))} * 100$$
(1)

where area(**f**) is the area of the rectangular bounding box of **f**. The squared distance between the feature vectors \mathbf{v}^p and \mathbf{v}^q are defined as $\delta(p,q) = \|\mathbf{v}^p - \mathbf{v}^q\|_2^2$. If a(p,q) > 0.85 and $\delta(p,q) \le 1.0$ i.e., if the faces have more than 85% overlap and less than 1.0 feature distance



Fig. 2 Example of 85% spatial overlap between face pairs in two consecutive frames which are combined to create face tracks

in consecutive frames, they are considered to be of the same person (see Fig. 2). Detected faces that overlap this way in consecutive frames are combined to form a *face track*, and the sequence of features corresponding to each of these faces is defined as a *feature track*.

Online face clustering.¹ The next task is to cluster the face tracks as they appear in each shot. For this subtask, we use our recently developed online clustering algorithm [8]. We assume the availability of all face tracks in a single shot at a given time. Our goal is to assign a face track belonging to a person who has appeared earlier to the correct existing cluster, and to form a new cluster for a face track belonging to a person appearing for the first time.

Let us consider a shot S_i containing K face tracks $\{\mathcal{F}_k\}_{k=1}^K$. Each face track \mathcal{F}_k is associated with a feature track $\mathcal{V}_k = \{\mathbf{v}_k^1, \mathbf{v}_k^2, \dots, \mathbf{v}_k^{N_k}\}$, where N_k is the number of faces in \mathcal{F}_k . Also consider that we have already processed previous (i - 1) shots and have obtained L clusters corresponding to L unique characters. The clusters are represented by their corresponding cluster centers $\mathcal{C} = \{\mathbf{c}_l\}_{l=1}^L$, where $\mathbf{c}_l \in \mathbb{R}^d$, where \mathbf{c}_l is the feature vector obtained by averaging all features across all face tracks within the l^{th} cluster. Note that the number of clusters and the clusters themselves are dynamic and they evolve as each shot is processed. We now define two matrices as follows:

- A temporal constraint matrix $\mathbf{Q} \in \mathbb{R}^{K \times K}$ is defined as

$$\mathbf{Q}(p,q) = \begin{cases} 0 & \text{if } \mathcal{F}_p \text{ and } \mathcal{F}_q \text{ overlap in time} \\ 1 & \text{otherwise} \end{cases}$$
(2)

where, $p, q \in \{1, 2, ..., K\}$. The matrix **Q** enforces a temporal constraint on the face tracks such that if two face tracks have any overlap in time, they are considered to belong to two different characters, and hence, are assigned to different clusters.

- A similarity matrix $\mathbf{D} \in \mathbb{R}^{L \times K}$ that measures the similarity between a face track (represented by \mathcal{V}_k) and a cluster center c_l for a given shot.

$$\mathbf{D}(l,k) = d(\mathbf{c}_l, \mathcal{V}_k) = 4 - \frac{1}{N_k} \sum_{j=1}^{N_k} \|\mathbf{v}_k^j - \mathbf{c}_l\|_2^2$$
(3)

where l = 1, 2, ..., L, and k = 1, 2, ..., K. The second component is an average squared distance, the maximum value of which is 4 (since each feature is a unit vector). By subtracting the distance from 4 we obtain a similarity value between [0, 4].

¹https://github.com/ankuPRK/COFC

Algorithm 1 Facetrack clustering in a single shot.

Input: Face track features in the current shot: $\{\mathcal{V}_k\}_{k=1}^K$, Initial clusters: \mathcal{C} Output: Updated \mathcal{C} Initialize: ind = [1, 2, ..., K], W = all-ones matrix. Compute Q, D using (2) and (3) while *length*(ind) > 0 do if \mathcal{C} not empty && max_{l,k}(D \odot W) $\geq \tau$ then $(\hat{l}, \hat{k}) \leftarrow \operatorname{argmax}_{l,k}(D \odot$ W) $k^* \leftarrow \operatorname{ind}[\hat{k}]$ Update cluster center $\mathbf{c}_{\hat{j}}$ with \mathcal{V}_{k^*} else Add new cluster $(\hat{l}, \hat{k}) \leftarrow (L + 1, 1)$ $k^* \leftarrow \operatorname{ind}[\hat{k}]$ $\mathbf{c}_{new} \leftarrow \operatorname{mean}(\mathcal{V}_{k^*})$ $\mathcal{C} \leftarrow \mathcal{C} \cup \mathbf{c}_{new}$ end Recompute D for $\mathbf{c}_{\hat{l}}$ $W(\hat{l}, :) \leftarrow \mathbf{Q}(\hat{k}, :)$ Delete D[:, \hat{k}], W[:, \hat{k}], Q[\hat{k} , :], Q[:, \hat{k}], ind[\hat{k}] end

Given $\{\mathcal{V}_k\}_{k=1}^{K}$, our task is to assign them to either one of the *L* clusters or create new clusters, if required. This is done by simply computing the similarities between \mathcal{V}_k for all *k* and $\{\mathbf{c}_l\}_{l=1}^{L}$.

$$(\hat{l}, \hat{k}) = \underset{l,k}{\operatorname{argmax}} (\mathbf{D} \odot \mathbf{W})$$
(4)

where, $\mathbf{W} \in \mathbb{R}^{L \times K}$ is a weight matrix (initialized with all ones) and \odot denotes element wise product. If $\max_{l,k}(\mathbf{D} \odot \mathbf{W}) \ge \tau$, where τ is an user defined threshold $\mathcal{V}_{\hat{k}}$ is assigned to the \hat{l}^{th} cluster. Consequently, we update $\mathbf{c}_{\hat{l}}$ by averaging over the existing and the newly added face track. On the other hand, if $\max_{l,k}(\mathbf{D} \odot \mathbf{W}) < \tau$, a new cluster is created assuming a new character has appeared. We add a new cluster $\mathcal{C} \leftarrow \mathcal{C} \cup \mathbf{c}_{new}$. Note that since \mathbf{W} is initialized as a matrix of all ones, it has no effect on the clustering of the first face track. For the subsequent assignments \mathbf{W} is updated to add temporal constraints. After $\mathcal{V}_{\hat{k}}$ is assigned to a cluster, we update \mathbf{D} and \mathbf{W} as follows:

- Case I: $\mathcal{V}_{\hat{k}}$ is assigned to an existing cluster \hat{l}

$$\mathbf{W}(\hat{l},:) \leftarrow \mathbf{Q}(\hat{k},:) \tag{5}$$

This updated **W** will make **D** \odot **W** zero for all the face tracks having any temporal overlap with $\mathcal{V}_{\hat{k}}$ in the \hat{l}^{th} row.

$$\mathbf{D}(\hat{l}, k) = d(\mathbf{c}_{\hat{l}}, \mathcal{V}_k) \text{ for } k \in [1, |\mathbf{ind}|]$$
(6)

- Case II: $\mathcal{V}_{\hat{k}}$ is assigned to a new cluster

$$\hat{l} = |\mathcal{C}| + 1 \tag{7}$$

$$\mathcal{C} \leftarrow \mathcal{C} \cup \mathbf{c}_{new} \tag{8}$$

$$\mathbf{W}(\hat{l},:) \leftarrow \mathbf{Q}(\hat{k},:) \tag{9}$$

$$\mathbf{D}(\hat{l}, k) = d(\mathbf{c}_{new}, \mathcal{V}_k) \text{ for } k \in [1, |\mathbf{ind}|]$$
(10)

where $\mathbf{ind} = [1, 2, \dots K]$. As $\mathcal{V}_{\hat{k}}$ is processed and sent into a cluster, its id is removed i.e., \hat{k}^{th} element of \mathbf{ind} , \hat{k}^{th} column of \mathbf{D} and \mathbf{W} , and \hat{k}^{th} row and column of \mathbf{Q} are removed.

This process goes on until all tracks in S_i are processed, and then we move to the next shot. We also keep track of the clusters that are updated during each shot. This information is later used to create and update the CIG. Algorithm 1 summarizes our proposed online face clustering algorithm.

3.2 CIG construction

We now describe the method to construct and update the CIG based on the online face clustering results. Each node in the CIG represents a single cluster corresponding to a character, and each edge captures the interaction between the two characters it connects. In our approach, the CIG is created in parallel to the online face clustering process, where new nodes are added to the CIG and the edge weights are updated after each shot is processed.

We define the relationship between two characters p and q in terms of their temporal co-occurrence in the same or consecutive shots. Considering an adjacency matrix A the relationship between p and q is formally defined as follows.

$$\mathbf{A}(p,q) = \sum_{i} [\mathbb{I}(p \in \mathcal{S}_i \& q \in \mathcal{S}_{i-1}) + \mathbb{I}(p \in S_i \& q \in S_i) + \mathbb{I}(p \in S_i \& q \in S_{i+1})]$$
(11)

where $\mathbb{I}(.)$ is the indicator function. This count defines the strength of the edge between p and q nodes in the CIG, and is denoted by the element $\mathbf{A}(p, q)$. A diagonal element $\mathbf{A}(p, p)$ denotes the number of times a character p appears in two consecutive shots. To construct and update \mathbf{A} in an online fashion, we begin with an empty \mathbf{A} and keep populating it with new rows and columns (corresponding to newly added nodes and edges) as new shots are processed. The dimension of \mathbf{A} thus increases as new characters are discovered, and consequently, new nodes are added to the CIG. According to our definition of character relationship in (11), we need to look for the characters in the shot immediately before and after it. Since we can not peek into the future shot, at shot S_i (i > 2), we update \mathbf{A} for S_{i-1} .

Our clustering algorithm yields updated cluster ids \mathcal{U}_{i-2} , \mathcal{U}_{i-1} , and \mathcal{U}_i pertaining to the shots S_{i-2} , S_{i-1} , S_i . We append N_{new}^{i-1} rows and N_{new}^{i-1} columns to **A** (all new elements initialized to 0), where N_{new}^{i-1} is the number of new clusters added during $(i-1)^{th}$ shot. Then **A**(p,q) is updated as follows.

$$\mathbf{A}(p,q) \leftarrow \mathbf{A}(p,q) + \mathbb{I}(p \in \mathcal{U}_{i-1} \& q \in \mathcal{U}_{i-2}) + \mathbb{I}(p \in \mathcal{U}_{i-1} \& q \in \mathcal{U}_{i-1}) + \mathbb{I}(p \in \mathcal{U}_{i-1} \& q \in \mathcal{U}_{i})$$
(12)

Algorithm 2 summarizes the entire process of online clustering and CIG creation as they are performed in parallel. Figure 3 shows an example of a CIG created using the proposed approach for a movie called *Hope Springs*. The CIG has 6 pure clusters corresponding to



Fig. 3 Character clusters (left) and the constructed CIG (right) for the movie *Hope Springs*. The CIG shows 7 nodes corresponding to the 7 clusters discovered by our algorithm. The node marked 'X' denotes a noisy cluster. The numbers below each node in the CIG denote the importance ($\sigma(p)$) of the characters

the 6 characters discovered by our online clustering algorithm and a noisy cluster denoted by 'X'. The edges depict the relationship between the characters where thicker edges denote higher interaction. The numbers represent the character importance scores, later described in Section 4.2 in detail.

Algorithm 2 Character interaction graph (CIG) construction via online clustering

```
Input: Movie frames: I_1, ..., I_T
Output:
Initialize: i \leftarrow 1, T_1 \leftarrow 1, C \leftarrow \{\}, \mathbf{A}_{CIG} \leftarrow \phi
for t = 1 to T do
      if shot boundary detected then
             T_{i+1} \leftarrow t
             Shot S_i \leftarrow \{I_{T_i}, I_{T_i+1} \dots I_{T_{i+1}-1}\}
            Create facetracks [\mathcal{F}_1, \ldots, \mathcal{F}_K] from \mathcal{S}_i
             Extract features [\mathcal{V}_1, \ldots, \mathcal{V}_K]; Set of cluster centers \mathcal{C} = \{\mathbf{c}_1, \ldots, \mathbf{c}_L\}
             \mathcal{C}, \mathcal{U}_i \leftarrow ClusterTracks(\mathcal{V}, \mathcal{C}) (see Algo. 1)
             if i > 2 then
                   N_{i-2} \leftarrow No. of clusters at (i-2)th shot
                   N_{i-1} \leftarrow No. of clusters at (i - 1)th shot
                 \mathbf{A}_{CIG} \leftarrow UpdateCIG (\mathcal{U}_{i-2}, \mathcal{U}_{i-1}, \mathcal{U}_i, N_{i-2}, N_{i-1}) \text{ (Using (12))}
             end
      end
end
```

4 Applications to movie analysis

In this section, we demonstrate the usefulness of the CIGs for two important movie analysis tasks: (i) Three act segmentation: detecting high level semantic structures in a movie, and

(ii) Major character identification. Below, we describe in detail how CIG can facilitate these tasks.

4.1 Three act segmentation

Popular films and screenplays are known to follow a well defined storytelling paradigm. The majority of movies consist of three main segments or acts (see Fig. 4): *Act I* - introduces the main characters and presents a key incident or plot point that drives the story, *Act II* - consists of a series of events including a key event which prepares the audience for the climax, and *Act III* - includes the climax and the resolution of the story [5, 12]. Discovering these high-level semantic units automatically can help in movie summarization and detection of the key events [6].

Our objective is to segment a movie into its three acts by detecting the two act boundaries as shown in Fig. 4. Consider the CIGs $A_{S_{i-1}}$ and A_{S_i} obtained at shots $S_{(i-1)}$ and S_i respectively. The difference between two CIGs is computed using *graph edit distance* (GED) as follows:

$$\text{GED}(\mathbf{A}_{\mathcal{S}_{(i-1)}}, \mathbf{A}_{\mathcal{S}_i}) = \Delta \eta + \Delta e$$

where, $\Delta \eta$ is the number of new nodes added to \mathbf{A}_{S_i} , and Δe is the number of edges that are modified to obtain \mathbf{A}_{S_i} from $\mathbf{A}_{S_{i-1}}$.

Using this measure, we compute how the CIG for a given movie changes over time between consecutive shots. A window of length T_w is used to sum all the GED scores within the window to incorporate a longer context and get a measure of overall interaction around each shot. Let this CIG difference be denoted as y^{ged} , where y_i^{ged} represents the changes in interaction around shot S_i . We detect act boundary I as follows

$$t_{b1} = \frac{\sum_{i \in \mathcal{B}_1} t_i y_i^{ged}}{\sum_{i \in \mathcal{B}_1} t_i}$$
(13)

where, t_i is the time at the center of S_i , and B_1 is a predefined interval. This interval B_1 is chosen leveraging information from film grammar [5], which suggests that act boundary I lies within 25 to 30 minutes from the start of the movie. We thus set \mathcal{B} to have all the shots between an interval of 22 to 40 minutes from the start of the movie. The act boundary II, t_{b2} is detected in a similar fashion with an interval \mathcal{B}_2 being 14 to 34 minute before the end of the movie.

4.2 Major character identification

Another important task in movie analysis is to identify its major characters. Past work on major character discovery using character networks usually rely on betweenness, centrality and sum of the edge weights [16, 21, 22]. We compute a new measure called the *eigenvector centrality* for each character in our CIG.



Fig. 4 The three-act narrative structure of a movie

The eigenvector centrality, e_p of a character (node) p measures the influence of the node p has on the CIG, and is defined as follows:

$$e_p = \frac{1}{\zeta} \sum_{q=1}^{|\mathcal{C}|} e_q \mathbf{A}(p,q) \tag{14}$$

where, ζ is the largest eigenvalue of **A**, and **A**(*p*, *q*) denotes the weight of the edge between nodes *p* and *q*. We then define a *node importance measure* $\sigma(p)$ for node *p* as follows:

$$\sigma(p) = \frac{e_p}{\sum_q e_q}.$$

It is easy to see that the higher the value of $\sigma(p)$, the more important is the node (character). We use the values of $\sigma(p)$ to rank the movie characters in terms of their importance in the movie. For example, Fig. 3 shows these node importance measures for the characters in a movie.

5 Performance evaluation

In this section, we present results and performance comparisons for the different components of our proposed method. First, we present results on the performance of the online face clustering algorithm as it is a critical component of the CIG construction algorithm, and its accuracy determines the quality of the CIG. Direct evaluation of a CIG is not very meaningful, as CIGs may have different characteristics by construction. Hence, we evaluate the usefulness of the CIGs via two movie analysis tasks - act segmentation and major character discovery.

5.1 Evaluating clustering performance

Databases: We use two databases that are commonly used to benchmark face clustering algorithms: (i) *Buffy* database (BF2006) [4, 26] containing 229 face tracks of 6 characters (17,337 faces altogether) extracted from the episode 2, season 5 of the TV series *Buffy - the Vampire Slayer*. The database includes the frame number, bounding box coordinates, track ids, and the character names for each face. (ii) *Notting Hill* database (NH2016) [24] that contains 277 face tracks of 7 characters (19,278 faces altogether) from the movie *Notting Hill*. It contains the frame numbers, bounding boxes, track ids, features and character names for each face in the database.

Experimental details: For each video, we obtain shot boundaries, create face tracks, extract deep features and cluster the faces using our proposed algorithm (see Algorithm 1). We use the FaceNet [17] to extract features from each face in a face track. The value of the threshold parameter τ is set to 2.80 and 2.85 for the BF2006 and the NH2016 database. For BF2006, we get a cluster for each of the 6 characters, and for NH2016, we get a cluster for 6 out of the 7 characters.

Comparison with existing methods: We compare with two baselines (Gaussian mixture model (GMM) with FaceNet features, and Kmeans with FaceNet features), and several state-of-the-art face clustering methods: (i) ULDML [2],(ii) a recently proposed constrained clustering method - the coupled HMRF (cHMRF) [24], and (iii) an

Method	BF2006	NH2016
ULDML [2]	49.29	43.82
cHMRF [24]	61.87	47.94
Luo et al. [30]	92.13	-
FaceNet + aCNN [19]	89.90	90.17
FaceNet + GMM	84.21	73.46
FaceNet + Kmeans	82.92	71.66
Ours	82.12	93.84

Table 1 Comparison with the state-of-the-art (offline) clustering methods in terms of clustering accuracy (%)

aggregated CNN feature-based clustering (aCNN) [19]. Performances of all the methods are compared in Table 1 in terms of clustering accuracy (expressed in %) which compares the predicted cluster labels with the ground truth labels. Note that all the methods in Table 1 are offline methods, where the entire data, information about the face tracks and the cluster counts are provided as an input to the algorithms. For the online method, however, no information about the face tracks or cluster counts are available. The performance of our algorithm on the BF2006 database is superior to that of cHMRF and ULDML, and is comparable to Kmeans. On the NH2016 database, our algorithm outperforms all its offline counterparts, achieving a clustering accuracy of 93.8%.

We next compare with the only existing online face clustering algorithm TCCRP [14]. We combine TCCRP with FaceNet features, and used a tracklet length of 10. Comparison is made in terms of *homogeneity* score, *completeness* score and their harmonic-mean i.e., the *V* measure (see Table 2). Table 2 shows that TCCRP has higher cluster homogeneity, but this is achieved at the cost of over-clustering (note the large number of clusters created by TCCRP) and thereby degrading completeness. Our method achieves significantly higher completeness and V measure while discovering a more accurate number of clusters.

5.2 Evaluating CIGs for act segmentation

Database: We use a database of six full length Hollywood movies: *Good Deeds, Hope Springs, Joyful Noise, Resident Evil, Step Up Revolution, The Vow.* These movies are known to have a well defined three act structure [6]. The labels for the act boundaries for the movies were annotated by three film experts. Each expert independently marked the act boundaries

	BF200	6	NH2016		
	TCCRP[14]	Ours	TCCRP[14]	Ours	
Homogeneity	0.93	0.68	0.92	0.88	
Completeness	0.44	0.69	0.44	0.89	
V measure	0.60	0.68	0.58	0.89	
Clusters formed	57	7	61	7	

Table 2 Comparison with the existing online clustering method

Table 3 Results on act	boundary detection: perfor	mance measured in terms	of the distance from huma	n expert annotated labels (in seco	nds)	
Movie	Video [6]	Music [6]	Text [6]	Multimodal [6]	Baseline	Ours
			Act boundary I			
Good Deeds	309	36	622	34	329	11
Hope Springs	13	34	31	31	620	66
Joyful Noise	98	143	452	19	28	17
Resident Evil	48	134	275	56	351	27
Step Up Rev.	96	346	277	265	578	289
The Vow	198	125	25	316	73	58
All movies	127.0	136.3	280.3	120.1	329.8	83.5
			Act boundary II			
Good Deeds	612	309	904	300	207	252
Hope Springs	210	1672	828	205	393	85
Joyful Noise	528	320	263	146	163	177
Resident Evil	836	575	702	575	751	52
Step Up Rev.	522	507	480	466	407	323
The Vow	413	2053	43	43	160	294
All movies	520.1	906.0	536.6	289.1	346.8	197.1



Fig. 5 Act boundary detection result for the movie *Hope Springs*

for each movie, and then decided on a final time stamp (at the precision level of seconds) through discussions [6].

Experimental set up: We run the DLib face detector [7] on each frame of the movie, and create face tracks. For removing the false detections and very small tracks we set the featuredistance threshold δ for track creation at 1.0, and the spacial overlap threshold α is set to 0.95. The online clustering threshold τ is set to 3.0 for all the movies. Also, the face tracks of length less than 15 are discarded.

Results and discussion: We detect the two act boundaries (see Fig. 4) in each movie using the CIGs as described in Section 4.1 and compute the error in terms of the distance (in seconds) from the expert annotations. The parameter value of T_w is set to 60s. We compare the performance of our CIG-based approach with that of an existing multimodal approach proposed in an earlier work [6]. We also create a simple baseline for comparison. The *baseline* sets the first act boundary at the 25th minute mark of the movie, and the second act boundary is at the 25th minute mark from the end of the movie. Table 3 present all the results of act boundary detection for the proposed method along with the baseline and the multimodal approach [6]. Our CIG-based approach performs the best in terms of overall error, even though it uses information from only the visual stream. We also note that

Movie	Faces	Shots	Tracks	Avg track length	Noisy clusters	Char clusters	Total
Good Deeds	38766	910	1078	35.96	2	7	9
Hope Springs	22663	1359	669	33.87	1	6	7
Joyful Noise	38703	2023	1205	32.11	7	14	21
Resident Evil	17962	2520	537	33.45	4	8	12
Step Up Rev.	21986	2790	733	29.99	6	13	19
The Vow	43860	1419	1161	37.77	3	5	8

 Table 4
 Face track statistics and clusters formed using online face clustering

Movie		Charac	Precision				
Good Deeds	Ours True	Lindsey Wesley	Wesley Lindsey	X Natalie	Walter Walter	X John	0.6
Hope Springs	Ours True	Arnold Arnold	Kay Kay	Bernie Bernie	X Eileen	Brad Brad	0.8
Joyful Noise	Ours True	Rose Rose	Olivia Olivia	Earla Sparrow	X Earla	Randy Randy	0.8
Resident Evil	Ours True	Alice Alice	X Ada	Ada Rain	Wesker Wesker	Rain Jill	0.8
Step Up Rev	Ours True	Emily Emily	Sean Sean	Eddy Eddy	Penelope William	William Jason	0.8
The Vow	Ours True	Paige Paige	Leo Leo	X Jeremy	X Bill	Jeremy Rita	0.6

Table 5 Major character identification results on six full-length movies ('X' denotes noisy cluster)

detecting act boundary II is more difficult as it has higher variability across movies. Figure 5 presents an example of CIG distance plot and the detected act boundaries for the movie *Hope Springs*.

5.3 Evaluating CIGs for major character identification

For this task, we use the same six movies as the three act segmentation task described in the previous section. The experimental settings remain the same.

Results and discussion: We first run our online face clustering algorithm on each movie. Some of the clusters thus obtained may be noisy i.e., they may contain faces from multiple characters. Such noisy clusters are formed due to (i) the presence of the minor characters in movies who do not appear on-screen long enough, and (ii) some wrongly clustered faces of major characters. Since the ground truth face labels are not available for the movies, we sought manual validation of the clusters to evaluate the performance of our method. Two human annotators labeled all the clusters formed for each movie, and identified each cluster as either a valid *character cluster* or a *noisy cluster*. The results are presented in Table 4.

After the clusters are formed, we compute $\sigma(p)$ for each cluster of any given movie, and identify the top 5 clusters (using the $\sigma(p)$ values) as the five major characters in the movie. To validate the results, we again seek manual evaluation. Two human annotators watched each movie, and based on the internet movie database (IMDb)² major cast list and the storyline, identified the top 5 characters in each movie. Table 5 presents the corresponding results, where 'X' denotes a noisy cluster. The results show that the top two characters are always retrieved correctly, and for most of the cases, our CIG-based approach is able to retrieve four out of the top five characters.

²www.imdb.com/

6 Conclusion

We proposed an unsupervised approach to building a *dynamic* character network of movie characters through online face clustering. This is significantly different from the existing body of work that builds a single, static character network using supervision from text, metadata, or human annotation. We demonstrated that the dynamic CIGs can successfully detect high-level semantic structures (acts) in movie narratives, and can also identify the major characters in a movie with high precision. Apart from the applications presented in this paper, dynamic CIGs are also expected to be useful for extracting character-level analytics, movie summarization, indexing and navigation.

Future work will be directed towards expanding the database used for validation, and leveraging the subtitle and audio information available for the movies to achieve better clustering accuracy. A scheme of splitting and fusing the clusters formed online can be useful to improve the quality of clusters, and in turn, can improve the quality of the CIG.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommonshorg/licenses/by/4.0/.

References

- Cao X, Zhang C, Zhou C, Fu H, Foroosh H (2015) Constrained multi-view video face clustering. IEEE Trans Image Process 24(11):4381–4393
- Cinbis RG, Verbeek J, Schmid C (2011) Unsupervised metric learning for face identification in tv video Computer Vision (ICCV), 2011 IEEE International Conference on, IEEE, pp 1559–1566
- 3. Du M, Chellappa R (2012) Face association across unconstrained video frames using conditional random fields European Conference on Computer Vision, Springer, pp 167–180
- 4. Everingham M, Sivic J, Zisserman A (2006) Hello! my name is... buffy-automatic naming of characters in tv video
- 5. Field S (2007) Screenplay: The foundations of screenwriting, Delta
- Guha T, Kumar N, Narayanan SS, Smith SL (2015) Computationally deconstructing movie narratives: an informatics approach Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on, IEEE, pp 2264–2268
- 7. King DE (2009) Dlib-ml: A machine learning toolkit. J Mach Learn Res 10:1755–1758
- Kulshreshtha P, Guha T (2018) An online algorithm for constrained face clustering in videos 2018 25th IEEE International Conference on Image Processing (ICIP), IEEE, pp 2670–2674
- 9. Le QV (2013) Building high-level features using large scale unsupervised learning Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on, IEEE, pp 8595–8598
- Li Y, Lee S-H, Yeh C-H, Kuo C-CJ (2006) Techniques for movie content analysis and skimming: tutorial and overview on video abstraction techniques. IEEE Signal Proc Mag 23(2):79–89
- Li Y, Narayanan S, Kuo C-CJ (2004) Content-based movie analysis and indexing based on audiovisual cues. IEEE transactions on circuits and systems for video technology 14(8):1073–1085
- 12. McKee R (1997) Substance, structure, style, and the principles of screenwriting. New York: Harper-Collins, New York
- Mitra A, Biswas S, Bhattacharyya C (2014) Temporally coherent bayesian models for entity discovery in videos by tracklet clustering. arXiv preprint arXiv:1409.6080
- Mitra A, Biswas S, Bhattacharyya C (2017) Bayesian modeling of temporal coherence in videos for entity discovery and summarization. IEEE transactions on pattern analysis and machine intelligence 39(3):430–443

- Park S-B, Oh K-J, Jo G-S (2012) Social network analysis in a movie using character-net. Multimedia Tools and Applications 59(2):601–627
- Ramakrishna A, Martínez VR, Malandrakis N, Singla K, Narayanan S (2017) Linguistic analysis of differences in portrayal of movie characters Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), vol 1, pp 1669–1678
- Schroff F, Kalenichenko D, Philbin J (2015) Facenet: A unified embedding for face recognition and clustering Proceedings of the IEEE conference on computer vision and pattern recognition, pp 815–823
- 18. Sharff S (1982) The elements of cinema: toward a theory of cinesthetic impact. Columbia University Press, New York
- Sharma V, Sarfraz MS, Stiefelhagen R (2017) A simple and effective technique for face clustering in tv series
- Sun Y, Chen Y, Wang X, Tang X (2014) Deep learning face representation by joint identificationverification Advances in neural information processing systems, pp 1988–1996
- Tran QD, Jung JE (2015) Cocharnet: Extracting social networks using character co-occurrence in movies. J. UCS 21(6):796–815
- 22. Weng C-Y, Chu W-T, Wu J-L (2009) Rolenet: Movie analysis from the perspective of social networks. IEEE Transactions on Multimedia 11(2):256–271
- Wolf L, Hassner T, Maoz I (2011) Face recognition in unconstrained videos with matched background similarity Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on, IEEE, pp 529– 534
- Wu B, Hu B-G, Ji Q (2017) A coupled hidden markov random field model for simultaneous face clustering and tracking in videos. Pattern Recogn 64:361–373
- Wu B, Lyu S, Hu B-G, Ji Q (2013) Simultaneous clustering and tracklet linking for multi-face tracking in videos Computer Vision (ICCV), 2013 IEEE International Conference on, IEEE, pp 2856–2863
- Wu B, Zhang Y, Hu B-G, Ji Q (2013) Constrained clustering and its application to face clustering in videos Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on, IEEE, pp 3507– 3514
- Xiao S, Tan M, Xu D (2014) Weighted block-sparse low rank representation for face clustering in videos European Conference on Computer Vision, Springer, pp 123–138
- Yeh M-C, Wu W-P (2014) Clustering faces in movies using an automatically constructed social network. IEEE MultiMedia 21(2):22–31
- 29. Zhang S, Gong Y, Wang J (2016) Deep metric learning with improved triplet loss for face clustering in videos Pacific Rim Conference on Multimedia, Springer, pp 497–508
- Zhang Z, Luo P, Loy CC, Tang X (2016) Joint face representation adaptation and clustering in videos European conference on computer vision, Springer, pp 236–251

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.