A Clustering-Based Method for Automatic Educational Video Recommendation Using Deep Face-Features of Lecturers

Paulo R. C. Mendes, Eduardo S. Vieira, Álan L. V. Guedes, Antonio J. G. Busson, and Sérgio Colcher
TeleMidia Lab - Department of Informatics
Pontifical Catholic University of Rio de Janeiro
Rio de Janeiro, Brazil
{paulo.mendes, eduardo, alan, busson}@telemidia.puc-rio.br, colcher@inf.puc-rio.br

Abstract—Discovering and accessing specific content within educational video bases is a challenging task, mainly because of the abundance of video content and its diversity. Recommender systems are often used to enhance the ability to find and select content. But, recommendation mechanisms, especially those based on textual information, exhibit some limitations, such as being error-prone to manually created keywords or due to imprecise speech recognition. This paper presents a method for generating educational video recommendation using deep face-features of lecturers without identifying them. More precisely, we use an unsupervised face clustering mechanism to create relations among the videos based on the lecturer's presence. Then, for a selected educational video taken as a reference, we recommend the ones where the presence of the same lecturers is detected. Moreover, we rank these recommended videos based on the amount of time the referenced lecturers were present. For this task, we achieved a mAP value of 99.165%.

Index Terms—Multimedia Retrieval, Deep Learning, Clustering, Educational Video, Video Analysis.

I. INTRODUCTION

The traditional paradigm of classroom courses centered on the physical presence of a teacher has been gradually giving space to online and hybrid courses, which enables the emergence of VTEs (Virtual Teaching Environment) and MOOCs (*Massive Open Online Courses*, such as Udacity,[1] Coursera,[2] and EdX).[3] For example, in 2018, a study [1] has shown that almost 59% of people aged 14 to 23 prefer YouTube as a learning tool rather than printed books, with 55% of them also saying that YouTube has contributed to their education. Recently, due to the covid-19 outbreak, the world has experienced an unprecedented usage of virtual education [2], and some say that this model of education came to stay.[4]

If, on the one hand, the abundance of educational videos can contribute to and facilitate learning, on the other hand, it also makes it challenging to discover and access the content of interest [3]. This issue is usually addressed by a proactive user search (using queries, for example), or by automatic recommendations made by specialized systems.

Recommendation mechanisms are usually based on two methods: *collaborative filtering* and *content-based filtering*. In collaborative filtering, the system groups users based on their common interest on items, using users' preferences, rates, purchases or accesses to those items. With this approach, knowledge about the item's content is not needed; the recommendation is purely based on the relationship between users and items. The content-based filtering, differently, requires items' description; similar items are the ones recommended to the user.

In general, the current video recommendation methods are heavily dependent on textual information from the video, such as labels (*i.e.* keywords) [4], [5], or automatically generated captions [6] from the lecturer speech. These systems face problems such as errors introduced by manually inserted labels and by imprecise speech recognition. In our research, we aim to investigate methods that are able to perform video recommendations that are not based on content nor on any error-prone textual descriptions, but solely on lecturers' presence. Notice that this approach does not necessarily have to completely replace textual-based recommendations; in fact, it can be easily used as an additional aid to enhance the ability to find content in any system.

Face detection methods have been attracting the attention of researchers for more than two decades [7]. Nowadays, it is used for surveillance, video analytics systems, smart shopping, automatic face tagging in photo collections, investigative tools that search for identities in social networks based on face images, and thousands of other applications in our daily lives. For instance, Facer [8] is the Facebook's face detection and recognition framework; given a photograph, it first detects all the faces, and then runs a deep model to determine the likelihood of that face belonging to one of the top-N user friends. This allows Facebook to suggest which friends the user might want to tag within the uploaded photographs.

This work aims at recommending educational video content based on lecturers' presence. To do that, we take advantage of face detection methods. More precisely, we detect lecturers

in a video taken as a reference and perform a clustering based on the face of these lecturers in different videos. Given these clusters, we extract their *centroids* (explained in Section III), and perform another clustering step for creating a relationship between videos that share the presence of the same lecturers. Finally, we rank the recommended videos based on the amount of time the referenced lecturers were present. A particular feature of this approach is that it can be done without supervision, allowing for new videos to be automatically analyzed. Moreover, our approach permits the creation of timelines based on lecturers' presence that can be used in the search for specific parts of a content where only specific lecturers' are present. To evaluate our recommendation ranking, we use the mAP (Mean Average Precision) metric, which is commonly used in information retrieval evaluation tasks [9].

The remainder of this paper is structured as follows. Section II discusses some related work. Then, we present our method in Section III. Section IV presents the used dataset, followed by Section V, that shows the experiments to validate the face clustering and the video recommendation ranking mechanisms. Finally, Section VI brings our final remarks.

II. RELATED WORK

We have organized the related work into two groups. In the first, we grouped works that share our goal of educational video recommendation but do not necessarily use face-embeddings (deep face-features). The second group is the one in which every work addresses the task of face recognition in videos.

Regarding *Educational Video Recommendation*, we cite works based on content-filtering. These works perform analyses and comparisons using the video textual description or speech recognition performed on them. Omisore *et. al.* [5], for example, propose combining *fuzzy* techniques to recommend books with content suitable for students based on their reading histories in a digital library, while Mahajan *et. al.* [4] propose, given a reference video, mining social media, and web for suggesting links for a student to visit. Moreover, Barrére *et. al.* [6] use texts from speech recognition to create recommendations. These works are only based on textual characteristics (or content converted to it) for performing recommendations. Our work focuses on using a visual part of the video, more precisely the presence of lecturers.

Works based on *Video Face Recognition* usually apply deep learning models for the task. DeepFace [10] and DeepID [11], for example, use a CNN (Convolutional Neural Network) with a fully-connected layer output to produce a representation of high-level features (face embeddings) from an input image, followed by a softmax layer to indicate the identity of classes. Other approaches, such as FaceNet [12], can directly measure the similarity among faces using euclidean space. Yang *et al.* [13] proposed a deep network for video face recognition called NAN (Neural Aggregation Network). They use a CNN to generate the embeddings, followed by an aggregation module that consists of two attention blocks which adaptively

aggregate the feature vectors to form a single feature inside the convex hull spanned by them. Rao *et al.* [14] proposed a method for video face recognition based on attention-aware deep reinforcement learning. They formulated the process of finding the attention of videos as a Markov decision process and training the attention model without using extra labels. Unlike existing attention models, their method takes information from both the image space and the feature space as the input to make use of face information that is discarded in the feature learning process. Sohn *et al.* [15] proposed an adaptative deep learning framework for image-based face recognition and video-based face recognition. Given an embedding generated by a CNN, their framework adaptation is achieved by (1) distilling knowledge from the network to a video adaptation network through feature matching, (2) performing feature restoration through synthetic data augmentation, and (3) learning a domain-invariant feature through an adversarial domain discriminator.

Like [13], [14], [15], our method uses a CNN to generate face embeddings from face images, with the difference that we use an unsupervised cluster-based method to compare the similarity among faces extracted from videos.

III. METHOD

Our method intends to recommend educational videos based on the lecturers that appear in each video, so that, when a person watches a video, other videos containing the same lecturers are recommended. For didactic purposes we decided to divide our exposition in two phases: (i) *video representation* and (ii) *video recommendation*, which are described in Sections III-A and III-B respectively.

A. Video Representation

The objective of this phase is to represent each video with vectors (centroids) of the lecturers that appear on it. Fig. 1 shows the pipeline we propose for this phase, described in the remainder of this subsection. It is divided into four steps: *Frames Extraction*, *Face Detection*, *Embeddings Generation* and *Clustering Representation*.

First, we perform the *Frames Extraction* by receiving a video file as input and extracting its frames according to a given frame rate. Next, for each of the frames, the *Face Detection* step uses an object detection model for detecting faces in each of them. The face detection model is responsible for returning the bounding boxes of the faces present in the image giving the x and y axes coordinates of the upper-left corner and of the lower-right corner of the rectangle that establishes the visual limits encapsulating each face. With these bounding boxes, we can isolate and extract the bounded images, obtaining a dataset composed of images of faces only.

The objective of the *Embeddings Generation* step is to represent each face image as a vector in $\mathbb{R}^n$. To achieve that, it processes each of the faces generated in the previous step through a CNN that generates their embeddings. An embedding is a representation of the input in a lower dimensionality space. Ideally, an embedding captures some semantics of the
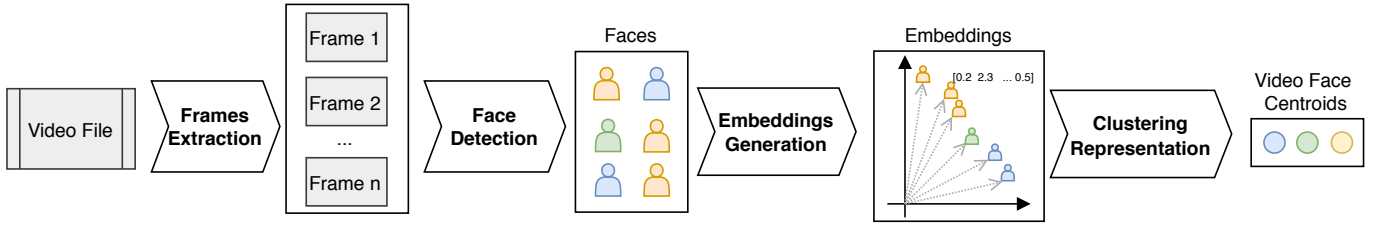
Fig. 1: Lecturers representation process in video. This process receives a video file and returns the centroids of the clusters that ideally represent each of the lecturers present in the video file.

input, e.g. by placing semantically similar inputs close together in an embedding space. Therefore, at the end of this step, we have all faces represented as embeddings.

In the *Clustering Representation* step, we group embeddings (e), and, consequently, faces that are close enough in the embedding space using a clustering algorithm. Clustering is the task of dividing a set of data points, embeddings in this case, into a number of groups (called *clusters*) such that data points in a given group are similar to other data points in the same group and dissimilar to the data points in other groups. The clustering process should produce a partition of the faces present in the frames, hopefully with each generated cluster representing a specific person; moreover, the union of all clusters covers the whole set of faces found in the video.

As most of the clustering algorithms require the number of clusters as parameter, we use a strategy (defined in Algorithm 1) based on the *Silhouette Score* (s) [16], that corresponds to the mean of the *Silhouette Coefficient* (σ) of all samples. This coefficient for each sample is

$$\sigma = \frac{b - a}{max(a, b)} \tag{1}$$

where a is the mean distance from a sample to all other samples in the same cluster, and b is the mean distance from a sample to all other samples in the closest cluster to that sample. In this way, the best value is 1 and the worst is -1. Values close to 0 indicate overlapping clusters, whereas negative values usually indicate that a sample has been assigned to the wrong cluster since a different cluster is more similar.

With the strategy defined in Algorithm 1, we increase the number of clusters until the maximum Silhouette Score decreases to more than t times in a row or until it reaches the maximum number of clusters (lines 5-18), which is the number of embeddings ($|e|$). The Clustering procedure (line 7) can be substituted by any clustering algorithm that requires the number of clusters in advance. When the iteration stops, we return the clustering configuration with the highest Silhouette Score. Since the Silhouette Coefficient requires at least two clusters, it would not be possible to compute the Silhouette Score for a clustering configuration with only one cluster (there are only faces of a single person). To overcome this problem, we start with 2 clusters consecutively increasing it as described above. Then, if the returned clustering configuration has a Silhouette Score smaller than a threshold ω, that

Algorithm 1 Iteratively finding the best clustering configuration for unknown number of clusters.

```
 1: procedure BLINDCLUSTERING(e, t, ω)
 2:     n_K ← 1
 3:     s_max ← −1
 4:     t_cur ← 0
 5:     while t_cur ≤ t & n_K < |e| do
 6:         n_K ← n_K + 1
 7:         K_cur ← Clustering(e, n_K)
 8:         s ← SilhouetteScore(K_cur)
 9:         if s < s_max then
10:             t_cur ← t_cur + 1
11:         else
12:             K ← K_cur
13:             t_cur ← 0
14:             if s > s_max then
15:                 s_max ← s
16:             end if
17:         end if
18:     end while
19:     if s_max < ω then
20:         K ← OneCluster(e)
21:     end if
22:     return K
23: end procedure
```

probably indicates overlapping, we say that all faces belong to one single cluster (lines 19-20).

Next, we compute the clusters' centroids for each of the clusters $k \in K$ where K is the best clustering configuration found with the *Silhouette Score*. A centroid c_k for each cluster is the mean of the elements present in the cluster, and can be defined as follows

$$c_k = \sum_{a \in k} \frac{a}{|k|} \tag{2}$$

where a represent each element of a cluster k.

By the end of this phase, we have each video in the dataset represented by its centroids where, ideally, each centroid represents a lecturer present in the video. We also record the frames where each lecturer is present.
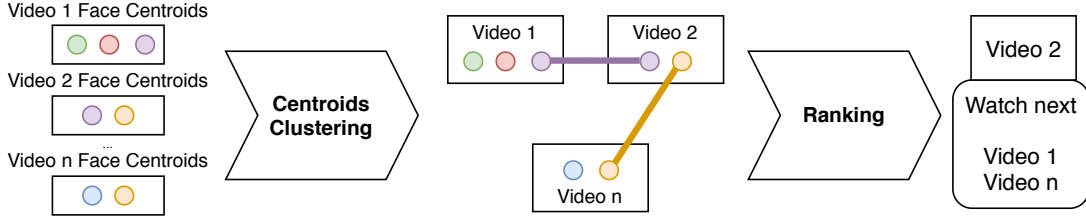
Fig. 2: Video Recommendation based on Lecturers Centroids Clustering. This pipeline receives the centroids of lecturers from all the videos in the dataset, then it creates relationships among videos that share the presence of the same lecturers. Finally, it performs ranking of recommended videos for each of the videos in the dataset. This ranking is based on the number of lecturers shared and their time presence.

B. Video Recommendation

This phase aims at recommending videos by the lecturers present in it and in the other videos. It is divided in two steps: *Centroids Clustering* and *Ranking*, as depicted in Fig. 2.

First, we gather the centroids from the videos of the dataset as one single set and perform the *Centroids Clustering*. For performing this clustering, we also use the strategy for an unknown number of clusters described in Algorithm 1. By doing that, we group centroids from the same lecturer that are in different videos. For instance, in Fig. 2, one can see that the *purple lecturer* is present in both Videos 1 and 2, while the *orange lecturer* is present in both Videos 2 and n. By the end of this step, we have the group L of lecturers present in the dataset of videos V, and we can also denote L_v as the group of lecturers present in video v.

Next, based on these relationships among different videos, we perform *Ranking*, by recommending videos in which lecturers of the current video are present. For doing that, we compute a similarity score using the presence of the lecturers in the current video and the presence of these same lecturers in the other video. Let $p_{l,v}$ denote the percentage of frames in which the lecturer $l \in L_v$ is present in video $v \in V$. For each video $v \in V$ and $u \in V - v$ we compute a score of similarity $S_{v,u}$.

$$S_{v,u} = \sum_{l \ \in \ L_v} p_{l,v} \cdot p_{l,u} \tag{3}$$

Finally, using this score, for each video v we compute a ranking R_v where $R_{v,i}$ denotes the *i-greatest* S_v and $R_{v,i} \geq R_{v,i+1}$ for all $i \ \in \ 1...n_v$, where n_v is the number of videos u in which $S_{v,u} > 0$. In this way, the more lecturers a video have in common with the reference video, and the more time these lecturers are present in both videos, the higher the video is positioned in the ranking of the reference video.

By the end of this phase, we have a ranking of recommended videos for each video in the dataset. It is important to notice that our method is unsupervised and does not require the information of the lecturers in advance. Consequently, we do not store any information regarding the identity of the lecturers, respecting their privacy.

IV. DATASET

The experiments were conducted using a dataset created in the context of this work. It is composed of 98 educational videos publicly available on YouTube.[5,6]

Each video contains at least one lecturer; moreover, some videos could have some special participation or collaboration. Thus, each video is annotated to contain between 1 to 5 people. In total, 16 people are present in the dataset. Each person has an average presence of 6.67% in the videos, and their identities are known and ready to be used to assist in the nominal dataset organization.

Regarding the duration, the videos vary between 00m:30s and 1h:49m:01s. The average duration of the videos is 23m:34s, with a standard deviation of 23m:05s. The high value of the standard deviation for the time estimates indicates that the videos are not in the same time range, and therefore have a wide duration variety.

V. EXPERIMENTS

First we compute the centroids that represent each lecturer in each video using the process described in Section III-A. Next we perform the video recommendation using the process described in Section III-B.

For representing the video files in the dataset, we start by performing *Frames Extraction* for each video file using a frame rate of 1 frame per second (fps). Next, in the *Face Detection* step, we use MTCNN [17] (Multitask Cascaded Convolutional Networks), which is widely used for the face detection task [18], [19], [20]. Once we have detected the faces of lecturers in the video frames, we perform *Embeddings Generation* using SE-ResNet-50 [21] (SeNet-50 for short) that generates embeddings on the $\mathbb{R}^{2048}$ feature space. We used the architecture and weights pre-trained on the VGGFace2 dataset [22], available on the *keras-vggface* library.[7] The VGGFace2 dataset contains 3.31 million images of 9,131 subjects and has large variations in pose, age, illumination, ethnicity, and profession. Finally, we use Algorithm 1 in the *Clustering Representation* step with the parameters $t = 5$, $\omega = 0.2$, and the Ward Agglomerative Clustering [23] as the `Clustering`

---

[5]https://www.youtube.com/channel/UCT0JugAtGmqiYkwxFZOwAtg
[6]https://www.youtube.com/user/deboraaladim
[7]https://github.com/rcmalli/keras-vggface

procedure using its implementation in the *scikit-learn* [24] library. The Ward Agglomerative Clustering algorithm merges pairs of clusters that minimize the *ward* criterion, which is the variance of the clustering being merged. By using this method, at each step, the algorithm finds the pair of clusters that lead to a minimum increase in total within-cluster variance after merging. Finally, we compute the centroids of each of the clusters generated.

For performing the video recommendation task, we gather the centroids (that represent each lecturer in the video) from all videos in the dataset. Next, we perform the process described in Section III-B. For *Centroids Clustering*, we also use Algorithm 1 with the same parameters $t = 5$, $\omega = 0.2$, and the Ward Aglommerative Clustering as `Clustering` procedure. Finally, based on the clusters generated, we perform the *Ranking* step.

The remainder of this Section describes the evaluation of the centroids clustering (Section V-A) and the evaluation of the video recommendation (Section V-B).

A. Centroids Clustering Evaluation

Our evaluation aims at discovering the precision achieved by the clustering over the centroids. More precisely, we want to evaluate how well our approach identified that the same lecturer is present in different video files. For this task, we require some human feedback. To receive that feedback, we developed an application, called *VideoFacesTool* consisting of a graphical web interface. The tool allows participants to import a file, which contains information after the *Centroids Clustering* step, so that we have the set of centroids and a sample face image of it. Inside the tool, a cluster is called a *group*. After being imported, faces are visually organized according to the group to which they belong and when a group is selected, all face centroids from that group are displayed, as shown in Fig. 3.

Each face centroid has a Boolean property, which indicates whether it is correctly grouped (it belongs to a group in which all the face centroids are from the same lecturer) or not. If there is an error, the participant can indicate it.

A total of 5 participants collaborated in the evaluation session. They were advised to mark a face centroid as wrong if it represents (a) an object, or (b) a part of the human body, or (c) a lecturer other than the lecturer in the group. Fig. 4 shows examples of these types of errors, and Table I provides an overview of the evaluations obtained from each participant. It is important to notice that these results do not reflect the recommendation of educational videos, they only evaluate the *Centroids Clustering* step. For instance, we could have a group of face centroids of people that appear for a few amount of time in the videos and are not lecturers. This case, of course, would reduce the precision of the *Centroids Clustering* step. However, our method for video recommendation and ranking is robust to these cases, as it considers the amount of time that a person appears for scoring the recommended videos. With the evaluation completed, it is possible to export the analysis information with the number of right and wrong face centroids.
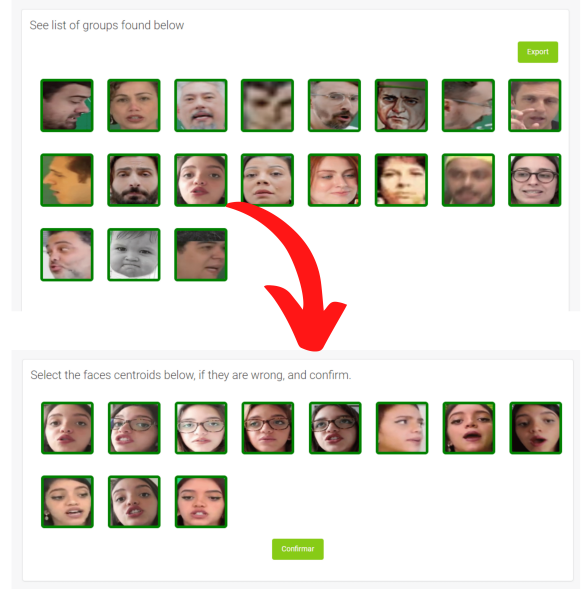

Fig. 3: Centroids images correction in *VideoFacesTool*. On the top, each image represents one lecturer. When a lecturer is selected, the tool displays all appearances (centroids) of that lecturer in different videos. The user can then mark each of these appearances as correct or wrong.
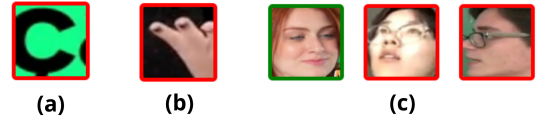

Fig. 4: Examples of wrong faces centroids. (a) a part of an icon (b) a hand and (c) face centroids that are not from the same lecturer

.

TABLE I: Results of visual evaluation of faces centroids showing the number of correct and wrong centroids classified by each participant.

Participant	#Correct	#Wrong	Precision
p1	163	62	72.44%
p2	160	65	71.11%
p3	164	61	72.89%
p4	164	61	72.89%
p5	162	63	72.00%
Avg	162.6	62.4	72.266%

B. Recommendation Evaluation

We evaluate our approach based on the relevance of the videos recommended. A video is considered relevant to another if they have at least one lecturer in common. To verify that, we use the information of the lecturers' presence available on our dataset.

To evaluate our ranking, for each video we compute the Average Precision (AP), that evaluates how well a ranking of

TABLE II: Results obtained with our approach with different thresholds of time presence for a lecturer to be considered as present in a video.

Thershold	MeanR	MinR	MeanP	MinP	MeanF1	MinF1	mAP	MinAP
0%	0,88851	0,45455	0,64681	0,20370	0,70971	0,33333	0,98641	0,58597
1%	0,88851	0,45455	0,64681	0,20370	0,70971	0,33333	0,98641	0,58597
2%	0,88851	0,45455	0,64885	0,20370	0,71171	0,33333	0,98641	0,58597
3%	0,88851	0,45455	0,67368	0,21569	0,73086	0,34921	0,98642	0,58597
4%	0,88851	0,45455	0,67368	0,21569	0,73086	0,34921	0,98642	0,58597
5%	0,88851	0,45455	0,69923	0,22449	0,74930	0,36066	0,98642	0,58597
6%	0,88851	0,45455	0,73615	0,25714	0,77648	0,40000	0,98642	0,58597
7%	0,88742	0,45455	0,77849	0,31429	0,80768	0,45833	0,98642	0,58597
8%	0,88742	0,45455	0,78408	0,31429	0,81111	0,45833	0,98643	0,58597
9%	0,88742	0,45455	0,80171	0,33333	0,82410	0,47826	0,98643	0,58597
10%	0,88742	0,45455	0,83165	0,42308	0,84306	0,51613	0,98643	0,58597
11%	0,88742	0,45455	0,85306	0,45714	0,85693	0,53333	0,98643	0,58597
12%	0,88616	0,45455	0,85956	0,44118	0,86018	0,50847	0,98662	0,58597
13%	0,88490	0,45455	0,88216	0,43750	0,87305	0,49123	0,98688	0,58597
14%	0,88289	0,45455	0,90265	0,47368	0,88450	0,51064	0,98884	0,58597
15%	0,88289	0,45455	0,90265	0,47368	0,88450	0,51064	0,98884	0,58597
16%	0,88163	0,44000	0,91327	0,47368	0,88980	0,47826	0,98908	0,58597
17%	0,87197	0,44000	0,91538	0,47368	0,88580	0,48889	0,98912	0,58597
18%	0,87197	0,44000	0,91538	0,47368	0,88580	0,48889	0,98912	0,58597
19%	0,87086	0,44000	0,93165	0,47368	0,89476	0,50000	0,98946	0,58597
20%	0,86130	0,35484	0,93645	0,44444	0,89218	0,46809	0,99000	0,58597
21%	0,86130	0,35484	0,93645	0,44444	0,89218	0,46809	0,99000	0,58597
22%	0,86130	0,35484	0,95054	0,61111	0,89886	0,46809	0,99000	0,58597
23%	0,86130	0,35484	0,95054	0,61111	0,89886	0,46809	0,99000	0,58597
24%	0,85805	0,32000	0,95718	0,64286	0,90046	0,43243	0,99165	0,58597
25%	0,85805	0,32000	0,95718	0,64286	0,90046	0,43243	0,99165	0,58597

recommendations is based on each element's relevancy. This metric penalizes more a ranking if a non-relevant element is recommended in the first positions than if it was in the last ones. Let P_k be the precision of the first k elements of a ranking, which is the percentage of videos that are relevant in the sub-ranking that starts at position 1 and ends at position k. Let α_k denote the relevancy of the video in position k, where $\alpha_k = 1$ if the video is relevant, and 0 otherwise. The AP of a given ranking is defined as follows

$$AP = \frac{1}{GTP} \sum_{k=1}^{n} P_k \cdot \alpha_k \qquad (4)$$

where GTP refers to the total number of ground truth positives in the ranking, which is the total number of videos that are considered relevant in a ranking. Fig. 5 shows an example of how the AP is computed for a given ranking. In this case, the $GTP = 3$ because the total number of relevant videos in the ranking is 3 (videos A, B and D).

In order to prevent outliers from having much influence in the recommendation (e.g. a person that is not a lecturer – and not relevant to the video – and appears for a short amount of time), we experimented different thresholds of presence intervals in a video for a person to be considered as "present" when computing the score for the ranking. In this way, a $p_{l,v}$ lesser than the threshold is considered as 0. Besides the Average Precision, we also compute the mean and minimum values of the recall (MeanR and MinR), precision (MeanP and MinP), and F1-Score (MeanF1 and Min F1) for the recommen-

AP = 1/3 (1/1 + 2/2 + 0 + 3/4 + 0) = 0.91667

Fig. 5: Example of how the Average Precision (AP) is computed for a reference video and its recommended videos.

dation generated for each of the videos, without considering the positioning of these videos in the rankings. The recall refers to the percentage of relevant videos that are present in the ranking. The F1-score represents an overall performance metric based on the harmonic mean of the precision and recall and is defined as follows.

$$F1 = \frac{2 \cdot P \cdot R}{P + R} \qquad (5)$$

Table II shows the thresholds used, the values of recall, precision, F1-score, and the mean and minimum Average Precision (mAP and MinAP),

One can observe from Table II that the precision clearly increased with the use of the threshold. Different from the precision, the recall decreased with the increase of the threshold. It means that with a greater threshold more videos that

should be recommended were not chosen by our method. It is important to notice that these two metrics (precision and recall) do not consider the ordering of the recommendations. Different from them, the Mean Average Precision (mAP) has high values for all thresholds, specially because the score for computing the ranking takes into consideration the percentage of time that a person appears in the reference and recommended videos. Then, we can conclude that our proposed approach for ordering the recommended videos tends to recommend more suitable videos first with a high mAP≈ 0.99. Moreover, despite the precision of the *Clustering Step* shown in Table I, our method for ranking was robust to outliers and was able to correctly recommend and rank relevant videos.

VI. FINAL REMARKS

In this paper, we present a method for educational video recommendation using deep-face-features of lecturers. More precisely, we use an unsupervised clustering-based method and an heuristic for ranking. It takes advantage of face detection mechanisms to perform educational video recommendation based on the lecturers' presence. Besides the face detection, we also perform face clustering of the lecturers in each video, and, given these clusters, we extract their centroids to perform another clustering step that creates a relationship of videos that share the presence of the same lecturers. Finally, we rank the recommended videos based on the amount of time each lecturer is present. It is worth mentioning that our method is completely automatic and does not require any information of the video files in advance. Moreover, our approach does not need to know or store the identity of the lecturers for performing recommendation, preserving their privacy.

A collateral contribution of our paper is video segmentation the by lecturer. As illustrated in Fig. 6, we can create a timeline based on lecturers' presence, which can be used to help students in finding moments where specific lecturers are present. With this segmentation, we could recommend specific parts of the video to the student.
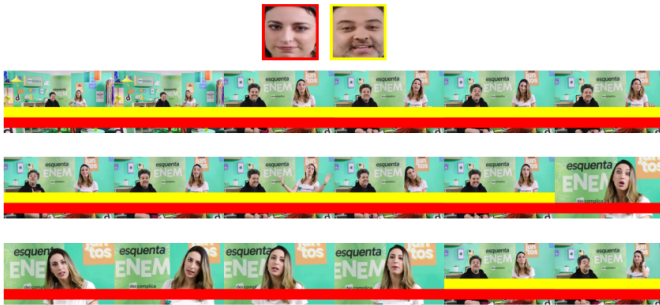

Fig. 6: Educational video timeline tagged by the lecturers presence. Notice that the frames with the lecturer on the left are tagged in red, while the frames with the lecturer on the right are tagged in yellow. This timeline is resulted from the *Video Representation* phase, described in Section III-A.

The main limitation of our work is that we can only recommend videos in which the lecturers are visually present.

As future work, we intend to use a hybrid recommendation approach, that combines both textual and audiovisual information from the video to create clusters. Video summarization is also a technique that can be explored to enhance video content searching and selection.

REFERENCES

[1] G. R. . Insights, "Beyond millennials: The next generation of learners," 2018. [Online]. Available: https://www.pearson.com/content/dam/one-dot-com/one-dot-com/global/Files/news/news-annoucements/2018/The-Next-Generation-of-Learners_final.pdf

[2] L. Sun, Y. Tang, and W. Zuo, "Coronavirus pushes education online," *Nature Materials*, vol. 19, no. 6, pp. 687–687, 2020.

[3] L. L. Dias, J. S. Barbosa, E. Barrére, and J. F. de Souza, "An approach to identify similarity among educational resources using external knowledge bases," *Brazilian Journal of Computers in Education*, vol. 25, no. 02, p. 18, 2017.

[4] R. Mahajan, J. Sodhi, and V. Mahajan, "Optimising web usage mining for building adaptive e-learning site: a case study," *International Journal of Innovation and Learning*, vol. 18, no. 4, pp. 471–486, 2015.

[5] M. Omisore and O. Samuel, "Personalized recommender system for digital libraries," *International Journal of Web-Based Learning and Teaching Technologies (IJWLTT)*, vol. 9, no. 1, pp. 18–32, 2014.

[6] E. Barrére, J. F. de Souza, M. A. Vitor, and M. A. de Almeida, "Utilização de enriquecimento semântico para a recomendação automática de videoaulas no moodle," *Revista Brasileira de Informática na Educação*, vol. 28, p. 319, 2020.

[7] I. Masi, Y. Wu, T. Hassner, and P. Natarajan, "Deep face recognition: A survey," in *2018 31st SIBGRAPI Conference on Graphics, Patterns and Images (SIBGRAPI)*, 2018, pp. 471–478.

[8] K. Hazelwood, S. Bird, D. Brooks, S. Chintala, U. Diril, D. Dzhulgakov, M. Fawzy, B. Jia, Y. Jia, A. Kalro *et al.*, "Applied machine learning at facebook: A datacenter infrastructure perspective," in *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 2018, pp. 620–629.

[9] C. D. Manning, "Introduction to information retrieval-cs 276 lecture slides.","" *Introduction to Information Retrieval" Stanford University*, 2009.

[10] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "Deepface: Closing the gap to human-level performance in face verification," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 1701–1708.

[11] Y. Sun, X. Wang, and X. Tang, "Deep learning face representation from predicting 10,000 classes," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2014, pp. 1891–1898.

[12] F. Schroff, D. Kalenichenko, and J. Philbin, "Facenet: A unified embedding for face recognition and clustering," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 815–823.

[13] J. Yang, P. Ren, D. Zhang, D. Chen, F. Wen, H. Li, and G. Hua, "Neural aggregation network for video face recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4362–4371.

[14] Y. Rao, J. Lu, and J. Zhou, "Attention-aware deep reinforcement learning for video face recognition," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 3931–3940.

[15] K. Sohn, S. Liu, G. Zhong, X. Yu, M.-H. Yang, and M. Chandraker, "Unsupervised domain adaptation for face recognition in unlabeled videos," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 3210–3218.

[16] P. J. Rousseeuw, "Silhouettes: a graphical aid to the interpretation and validation of cluster analysis," *Journal of computational and applied mathematics*, vol. 20, pp. 53–65, 1987.

[17] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," *IEEE Signal Processing Letters*, vol. 23, no. 10, pp. 1499–1503, Oct 2016.

[18] E. Jose, M. Greeshma, M. H. TP, and M. Supriya, "Face recognition based surveillance system using facenet and mtcnn on jetson tx2," in *2019 5th International Conference on Advanced Computing & Communication Systems (ICACCS)*. IEEE, 2019, pp. 608–613.

[19] A. Ghofrani, R. M. Toroghi, and S. Ghanbari, "Realtime face-detection and emotion recognition using mtcnn and minishufflenet v2," in *2019 5th Conference on Knowledge Based Engineering and Innovation (KBEI)*. IEEE, 2019, pp. 817–821.

[20] G. Bezerra and R. Gomes, "Recognition of occluded and lateral faces using mtcnn, dlib and homographies," 2018.

[21] J. Hu, L. Shen, and G. Sun, "Squeeze-and-excitation networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 7132–7141.

[22] Q. Cao, L. Shen, W. Xie, O. M. Parkhi, and A. Zisserman, "Vggface2: A dataset for recognising faces across pose and age," in *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*. IEEE, 2018, pp. 67–74.

[23] J. H. Ward Jr, "Hierarchical grouping to optimize an objective function," *Journal of the American statistical association*, vol. 58, no. 301, pp. 236–244, 1963.

[24] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.