






Person Search in Videos with One Portrait Through Visual and Temporal Links

Qingqiu Huang¹(✉) , Wentao Liu^{2,3} , and Dahua Lin¹ 

¹ CUHK-SenseTime Joint Lab, The Chinese University of Hong Kong, Shatin, Hong Kong

{hq016,dhlin}@ie.cuhk.edu.hk

² Department of Computer Science and Technology, Tsinghua University, Beijing, China

³ SenseTime Research, Beijing, China
liuwtwinter@gmail.com

Abstract. In real-world applications, *e.g.* law enforcement and video retrieval, one often needs to search a certain person in long videos with *just one portrait*. This is much more challenging than the conventional settings for person re-identification, as the search may need to be carried out in the environments different from where the portrait was taken. In this paper, we aim to tackle this challenge and propose a novel framework, which takes into account the identity invariance along a tracklet, thus allowing person identities to be propagated via both the visual and the temporal links. We also develop a novel scheme called *Progressive Propagation via Competitive Consensus*, which significantly improves the reliability of the propagation process. To promote the study of person search, we construct a large-scale benchmark, which contains 127K *manually annotated* tracklets from 192 movies. Experiments show that our approach remarkably outperforms mainstream person re-id methods, raising the mAP from 42.16% to 62.27% (Code at <https://github.com/hqqasw/person-search-PPCC>).

Keywords: Person search · Portrait · Visual and temporal
Progressive Propagation · Competitive Consensus

1 Introduction

Searching persons in videos is frequently needed in real-world scenarios. To catch a wanted criminal, the police may have to go through thousands of hours of videos collected from multiple surveillance cameras, probably with just a single

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-01261-8_26) contains supplementary material, which is available to authorized users.

portrait. To find the movie shots featured by a popular star, the retrieval system has to examine many hour-long films, with just a few facial photos as the references. In applications like these, the reference photos are often taken in an environment that is very different from the target environments where the search is conducted. As illustrated in Fig. 1, such settings are very challenging. Even state-of-the-art recognition techniques would find it difficult to reliably identify all occurrences of a person, facing the dramatic variations in pose, makeups, clothing, and illumination.




Fig. 1. Person re-id differs significantly from the person search task. The first row shows a typical example in person re-id from the *MARS* dataset [44], where the reference and the targets are captured under similar conditions. The second row shows an example from our person search dataset *CSM*, where the reference portrait is dramatically different from the targets that vary significantly in pose, clothing, and illumination.

It is noteworthy that two related problems, namely *person re-identification* (*re-id*) and *person recognition in albums*, have drawn increasing attention from the research community. However, they are substantially different from the problem of *person search with one portrait*, which we aim to tackle in this work. Specifically, in typical settings of person re-id [8, 13, 16, 22, 38, 44, 45], the queries and the references in the gallery set are usually captured under similar conditions, *e.g.* from different cameras along a street, and within a short duration. Even though some queries can be subject to issues like occlusion and pose changes, they can still be identified via other visual cues, *e.g.* clothing. For person recognition in albums [43], one is typically given a diverse collection of gallery samples, which may cover a wide range of conditions and therefore can be directly matched to various queries. Hence, for both problems, the references in the gallery are often good representatives of the targets, and therefore the methods based on visual cues can perform reasonably well [1, 3, 4, 14, 15, 22, 39, 43, 44]. On the contrary, our task is to bridge a single portrait with a highly diverse set of samples, which is much more challenging and requires new techniques that go beyond visual matching.

To tackle this problem, we propose a new framework that propagates labels through both visual and temporal links. The basic idea is to take advantage of the *identity invariance* along a person trajectory, *i.e.* all person instances along a *continuous* trajectory in a video should belong to the same identity. The connections induced by tracklets, which we refer to as the *temporal links*, are complementary to the *visual links* based on feature similarity. For example, a trajectory can sometimes cover a wide range of facial images that can not be easily associated based on visual similarity. With both *visual* and *temporal* links incorporated, our framework can form a large connected graph, thus allowing the identity information to be propagated over a very diverse collection of instances.

While the combination of visual and temporal links provide a broad foundation for identity propagation, it remains a very challenging problem to carry out the propagation *reliably* over a large real-world dataset. As we begin with only a single portrait, a few wrong labels during propagation can result in catastrophic errors downstream. Actually, our empirical study shows that conventional schemes like linear diffusion [46, 47] even leads to substantially worse results. To address this issue, we develop a novel scheme called *Progressive Propagation via Competitive Consensus*, which performs the propagation *prudently*, spreading a piece of identity information only when there is high certainty.

To facilitate the research on this problem setting, we construct a dataset named *Cast Search in Movies (CSM)*, which contains 127K tracklets of 1218 cast identities from 192 movies. The identities of all the tracklets are *manually annotated*. Each cast identity also comes with a reference portrait. The benchmark is very challenging, where the person instances for each identity varies significantly in makeup, pose, clothing, illumination, and even age. On this benchmark, our approach get 63.49% and 62.27% mAP under two settings, Comparing to the 53.33% and 42.16% mAP of the conventional visual-matching method, it shows that only matching by visual cues can not solve this problem well, and our proposed framework – *Progressive Propagation via Competitive Consensus* can significantly raise the performance.

In summary, the main contributions of this work lie in four aspects: (1) We systematically study the problem of *person search in videos*, which often arises in real-world practice, but remains widely open in research. (2) We propose a framework, which incorporates both the visual similarity and the identity invariance along a tracklet, thus allowing the search to be carried out much further. (3) We develop the *Progressive Propagation via Competitive Consensus* scheme, which significantly improves the reliability of propagation. (4) We construct a dataset *Cast Search in Movies (CSM)* with 120K manually annotated tracklets to promote the study on this problem.

2 Related Work

Person Re-id. Person re-id [6, 7, 41], which aims to match pedestrian images (or tracklets) from different cameras within a short period, has drawn much attention in the research community. Many datasets [8, 13, 16, 22, 38, 44, 45] have been

proposed to promote the research of re-id. However, the videos are captured by just several cameras in nearby locations within a short period. For example, the Airport [16] dataset is captured in an airport from 8 a.m. to 8 p.m. in one day. So the instances of the same identities are usually similar enough to identify by visual appearance although with occlusion and pose changes. Based on such characteristic of the data, most of the re-id methods focus on how to match a query and a gallery instance by visual cues. In early works, the matching process is split into feature designing [9, 11, 26, 27] and metric learning [17, 23, 28]. Recently, many deep learning based methods have been proposed to jointly handle the matching problem. *Li et al.* [22] and *Ahmed et al.* [1] designed siamese-based networks which employ a binary verification loss to train the parameters. *Ding et al.* [4] and *Cheng et al.* [3] exploit triple loss for training more discriminating feature. *Xiao et al.* [39] and *Zheng et al.* [44] proposed to learn features by classifying identities. Although the feature learning methods of re-id can be adopted for the Person Search with One Portrait problem, they are substantially different as the query and the gallery would have huge visual appearances gap in person search, which would make one-to-one matching fail.

Person Recognition in Photo Album. Person recognition [14, 15, 19, 24, 43] is another related problem, which usually focuses on the persons in photo album. It aims to recognize the identities of the queries given a set of labeled persons in gallery. *Zhang et al.* [43] proposed a Pose Invariant Person Recognition method (PIPER), which combines three types of visual recognizers based on ConvNets, respectively on face, full body, and poselet-level cues. The PIPA dataset published in [43] has been widely adopted as a standard benchmark to evaluate person recognition methods. *Oh et al.* [15] evaluated the effectiveness of different body regions, and used a weighted combination of the scores obtained from different regions for recognition. *Li et al.* [19] proposed a multi-level contextual model, which integrates person-level, photo-level and group-level contexts. But the person recognition is also quite different from the person search problem we aim to tackle in this paper, since the samples of the same identities in query and gallery are still similar in visual appearances and the methods mostly focus on recognizing by visual cues and context.

Person Search. There are some works that focus on person search problem. *Xiao et al.* [40] proposed a person search task which aims to search the corresponding instances in the images of the gallery without bounding box annotation. The associated data is similar to that in re-id. The key difference is that the bounding box is unavailable in this task. Actually it can be seen as a task to combine pedestrian detection and person re-id. There are some other works try to search person with different modality of data, such as language-based [21] and attribute-based [5, 35], which focus on the application scenarios that are different from the portrait-based problem we aim to tackle in this paper.

Label Propagation. Label propagation (LP) [46, 47], also known as Graph Transduction [30, 32, 37], is widely used as a semi-supervised learning method. It relies on the idea of building a graph in which nodes are data points (labeled and

Table 1. Comparing CSM with related datasets

Dataset	CSM	MARS[44]	iLIDS[38]	PRID[13]	Market[45]	PSD[40]	PIPA[43]
Task	Search	re-id	re-id	re-id	re-id	det.+re-id	recog.
Type	Video	Video	Video	Video	Image	Image	Image
Identities	1,218	1,261	300	200	1,501	8,432	2,356
Tracklets	127K	20K	600	400	-	-	-
Instances	11M	1M	44K	40K	32K	96K	63K




Fig. 2. Examples of *CSM* Dataset. In each row, the photo on the left is the query portrait and the following tracklets of are ground-truth tracklets of them in the gallery.

unlabeled) and the edges represent similarities between points so that labels can propagate from labeled points to unlabeled points. Different kinds of LP-based approaches have been proposed for face recognition [18,48], semantic segmentation [33], object detection [36], saliency detection [20] in the computer vision community. In this paper, We develop a novel LP-based approach called Progressive Propagation via Competitive Consensus, which differs from the conventional LP in two folds: (1) propagating by competitive consensus rather than linear diffusion, and (2) iterating in a progressive manner.

3 Cast Search in Movies Dataset

Whereas there have been a number of public datasets for person re-id [8,13,16,22,38,44,45] and album-based person recognition [43]. But dataset for our task, namely person search with a single portrait, remains lacking. In this work,

we constructed a large-scale dataset *Cast Search in Movies (CSM)* for this task. *CSM* comprises a *query set* that contains the portraits for 1,218 cast (the actors and actresses) and a *gallery set* that contains 127K tracklets (with 11M person instances) extracted from 192 movies.

We compare *CSM* with other datasets for person re-id and person recognition in Table 1. We can see that *CSM* is significantly larger, 6 times for tracklets and 11 times more instances than MARS [44], which is the largest dataset for person re-id to our knowledge. Moreover, *CSM* has a much wider range of tracklet durations (from 1 to 4686 frames) and instance sizes (from 23 to 557 pixels in height). Figure 2 shows several example tracklets as well as their corresponding portraits, which are very diverse in pose, illumination, and wearings. It can be seen that the task is very challenging (Fig. 3).




Fig. 3. Statistics of *CSM* dataset. (a): the tracklet number distribution over movies. (b): the tracklet number of each movie, both credited cast and “others”. (c): the distribution of tracklet number over cast. (d): the distribution of length (frames) over tracklets. (e): the distribution of height (px) over tracklets.

Query Set. For each movie in *CSM*, we acquired the cast list from IMDB. For those movies with more than 10 cast, we only keep the top 10 according to the IMDB order, which can cover the main characters for most of the movies. In total, we obtained 1,218 cast, which we refer to as the *credited cast*. For each credited cast, we download a portrait from either its IMDB or TMDB homepage, which will serve as the query portraits in *CSM*.

Gallery Set. We obtained the tracklets in the gallery set through five steps:

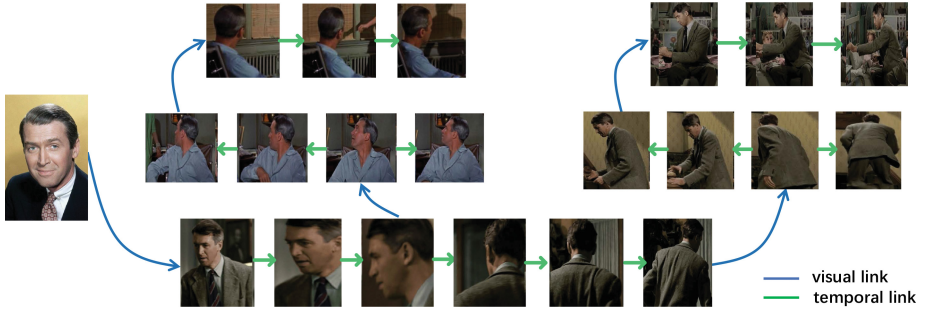


Fig. 4. Visual links and temporal links in our graph. We only keep one strongest link for each pair of tracklets. And we can see that these two kinds of links are complementary. The former allows the identity information to be propagated among those instances that are similar in appearance, while the latter allows the propagation along a continuous tracklet, in which the instances can look significantly different. With both types of links incorporated, we can construct a more connected graph, which allows the identities to be propagated much further.

1. *Detecting shots.* A movie is composed of a sequence of shots. Given a movie, we first detected the shot boundaries of the movies using a fast shot segmentation technique [2, 34], resulting in totally 200K shots for all movies. For each shot, we selected 3 frames as the *keyframes*.
2. *Annotating bounding boxes on keyframes.* We then *manually* annotated the person bounding boxes on keyframes and obtained around 700K bounding boxes.
3. *Training a person detector.* We trained a person detector with the annotated bounding boxes. Specifically, all the keyframes are partitioned into a training set and a testing set by a ratio 7:3. We then finetuned a Faster-RCNN [29] pre-trained on MSCOCO [25] on the training set. On the testing set, the detector gets around 91% mAP, which is good enough for tracklet generation.
4. *Generating tracklets.* With the person detector as described above, we performed per-frame person detection over all the frames. By concatenating the bounding boxes across frames with $\text{IoU} > 0.7$ *within each shot*, we obtained 127K tracklets from the 192 movies.
5. *Annotating identities.* Finally, we manually annotated the identities of all the tracklets. Particularly, each tracklet is annotated as one of the credited cast or as “others”. Note that the identities of the tracklets in each movie are annotated independently to ensure high annotation quality with a reasonable budget. Hence, being labeled as “others” means that the tracklet does not belong to any credited cast of the corresponding movie.

4 Methodology

In this work, we aim to develop a method to find all the occurrences of a person in a long video, *e.g.* a movie, with just a single portrait. The challenge of this

task lies in the vast gap of visual appearance between the portrait (query) and the candidates in the gallery.

Our basic idea to tackle this problem by leveraging the inherent *identity invariance* along a person tracklet and propagate the identities among instances via both visual and temporal links. The visual and temporal links are complementary. The use of both types of links allows identities to be propagated much further than using either type alone. However, how to propagate over a large, diverse, and noisy dataset reliably remains a very challenging problem, considering that we only begin with just a small number of labeled samples (the portraits). The key to overcoming this difficulty is to be *prudent*, only propagating the information which we are certain about. To this end, we propose a new propagation framework called *Progressive Propagation via Competitive Consensus*, which can effectively identify confident labels in a competitive way.

4.1 Graph Formulation

The propagation is carried out over a graph among person instances. Specifically, the propagation graph is constructed as follows. Suppose there are C cast in query set, M tracklets in gallery set, and the length of k -th tracklet (denoted by τ_k) is n_k , *i.e.* it contains n_k instances. The cast portraits and all the instances along the tracklets are treated as graph nodes. Hence, the graph contains $N = C + \sum_{k=1}^M n_k$ nodes. In particular, the identities of the C cast portraits are known, and the corresponding nodes are referred to as *labeled nodes*, while the other nodes are called *unlabeled nodes*.

The propagation framework aims to propagate the identities from the labeled nodes to the unlabeled nodes through both *visual* and *temporal* links between them. The *visual links* are based on feature similarity. For each instance (say the i -th), we can extract a feature vector, denoted as \mathbf{v}_i . Each visual link is associated with an affinity value – the affinity between two instances \mathbf{v}_i and \mathbf{v}_j is defined to be their cosine similarity as $w_{ij} = \mathbf{v}_i^T \mathbf{v}_j / (\|\mathbf{v}_i\| \cdot \|\mathbf{v}_j\|)$. Generally, higher affinity value w_{ij} indicates that \mathbf{v}_i and \mathbf{v}_j are more likely to be from the same identity. The *temporal links* capture the *identity invariance* along a tracklet, *i.e.* all instances along a tracklet should share the same identity. In this framework, we treat the identity invariance as hard constraints, which is enforced via a *competitive consensus* mechanism.

For two tracklets with lengths n_k and n_l , there can be $n_k \cdot n_l$ links between their nodes. Among all these links, the strongest link, *i.e.* the one between the most similar pair, is the best to reflect the visual similarity. Hence, we only keep one strongest link for each pair of tracklets as shown in Fig. 4, which makes the propagation more reliable and efficient. Also, thanks to the temporal links, such reduction would not compromise the connectivity of the whole graph.

As illustrated in Fig. 4, the visual and temporal links are complementary. The former allows the identity information to be propagated among those instances that are similar in appearance, while the latter allows the propagation along a continuous trajectory, in which the instances can look significantly different. With only visual links, we can obtain clusters in the feature space. With only

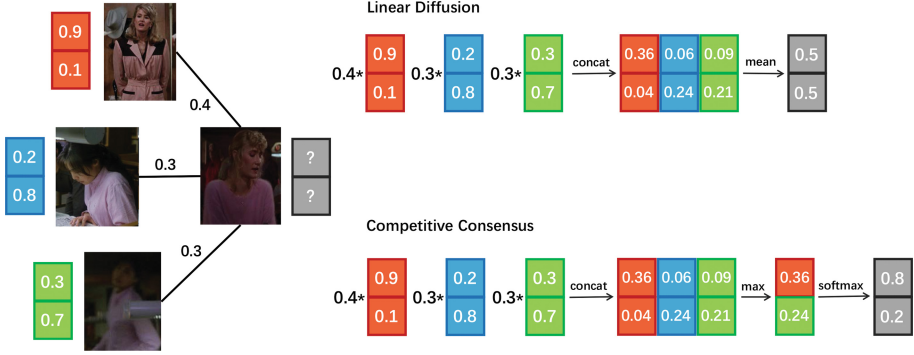


Fig. 5. An example to show the difference between competitive consensus and linear diffusion. There are four nodes here and their probability vectors are shown by their sides. We are going to propagate labels from the left nodes to the right node. However, two of its neighbor nodes are noise. The calculation process of linear diffusion and competitive consensus are shown on the right side. We can see that in a graph with much noise, our competitive consensus, which aims to propagate the most confident information, is more robust.

temporal links, we only have isolated tracklets. However, with both types of links incorporated, we can construct a more connected graph, which allows the identities to be propagated much further.

4.2 Propagating via Competitive Consensus

Each node of the graph is associated with a probability vector $\mathbf{p}_i \in \mathbb{R}^C$, which will be iteratively updated as the propagation proceeds. To begin with, we set the probability vector for each labeled node to be a one-hot vector indicating its label, and initialize all others to be zero vectors. Due to the identity invariance along tracklets, we enforce all nodes along a tracklet τ_k to share the same probability vector, denoted by \mathbf{p}_{τ_k} . At each iteration, we traverse all tracklets and update their associated probability vectors one by one.

Linear Diffusion. Linear diffusion is the most widely used propagation scheme, where a node would update its probability vector by taking a linear combination of those from the neighbors. In our setting with identity invariance, the linear diffusion scheme can be expressed as follows:

$$\mathbf{p}_{\tau_k}^{(t+1)} = \sum_{j \in \mathcal{N}(\tau_k)} \alpha_{kj} \mathbf{p}_j^{(t)}, \quad \text{with } \alpha_{kj} = \frac{\tilde{w}_{kj}}{\sum_{j' \in \mathcal{N}(\tau_k)} \tilde{w}_{kj'}}. \quad (1)$$

Here, $\mathcal{N}(\tau_k) = \cup_{i \in \tau_k} \mathcal{N}_i$ is the set of all visual neighbors of those instances in τ_k . Also, \tilde{w}_{kj} is the *affinity* of a neighbor node j to the tracklet τ_k . Due to the constraint that there is only one visual link between two tracklets (see Sect. 4.1), each neighbor j will be connected to just one of the nodes in τ_k , and \tilde{w}_{kj} is set to the affinity between the neighbor j to that node.

However, we found that the linear diffusion scheme yields poor performance in our experiments, even far worse than the naive visual matching method. An important reason for the poor performance is that errors will be mixed into the updated probability vector and then propagated to other nodes. This can cause catastrophic errors downstream, especially in a real-world dataset that is filled with noise and challenging cases.

Competitive Consensus. To tackle this problem, it is crucial to improve the reliability and propagate the most confident information only. Particularly, we should only trust those neighbors that provide strong evidence instead of simply taking the weighted average of all neighbors. Following this intuition, we develop a novel scheme called *competitive consensus*.

When updating \mathbf{p}_{τ_k} , the probability vector for the tracklet τ_k , we first collect the strongest evidence to support each identity c , from all the neighbors in $\mathcal{N}(\tau_k)$, as

$$\eta_k(c) = \max_{j \in \mathcal{N}(\tau_k)} \alpha_{kj} \cdot p_j^{(t)}(c), \quad (2)$$

where the normalized coefficient α_{kj} is defined in Eq. (1). Intuitively, an identity is *strongly* supported for τ_k if one of its neighbors assigns a high probability to it. Next, we turn the evidences for individual identities into a probability vector via a tempered softmax function as

$$p_{\tau_k}^{(t+1)}(c) = \exp(\eta_k(c)/T) / \sum_{c'=1}^C \exp(\eta_k(c')/T). \quad (3)$$

Here, T is a temperature the controls how much the probabilities concentrate on the strongest identity. In this scheme, all identities compete for getting high probability values in $\mathbf{p}_{\tau_k}^{(t+1)}$ by collecting the strongest supports from the neighbors. This allows the strongest identity to stand out.

Competitive consensus can be considered as a coordinate ascent method to solve Eq. 4, where we introduce a binary variable $z_{kj}^{(c)}$ to indicate whether the j -th neighbor is a trustable source for the class c for the k -th tracklet. Here, \mathcal{H} is the entropy. The constraint means that one trustable source is selected for each class c and tracklet k .

$$\max \sum_{c=1}^C p_{\tau_k}^{(c)} \sum_{j \in \mathcal{N}(\tau_k)} \alpha_{kj} z_{kj}^{(c)} p_j^{(c)} + \sum_{c=1}^C \mathcal{H}(p_{\tau_k}^{(c)}) \quad s.t. \quad \sum_{j \in \mathcal{N}(\tau_k)} z_{kj}^{(c)} = 1. \quad (4)$$

Figure 5 illustrates how linear diffusion and our competitive Consensus work. Experiments on CSM also show that competitive consensus significantly improves the performance of the person search problem.

4.3 Progressive Propagation

In conventional label propagation, labels of all the nodes would be updated until convergence. This way can be prohibitively expensive when the graph contains

Table 2. Train/Val/Test splits of CSM

	Movies	Cast	Tracklets	Credited tracklets
Train	115	739	79K	47K
Val	19	147	15K	8K
Test	58	332	32K	18K
Total	192	1,218	127K	73K

Table 3. Query/Gallery size

Setting	Query	Gallery
IN (per movie)	6.4	560.5
CROSS	332	17,927

a large number of nodes. However, for the person search problem, this is unnecessary – when we are very confident about the identity of a certain instance, we don’t have to keep updating it.

Motivated by the analysis above, we propose a *progressive propagation* scheme to accelerate the propagation process. At each iteration, we will fix the labels for a certain fraction of nodes that have the highest confidence, where the confidence is defined to be the maximum probability value in \mathbf{p}_i . We found empirically that a simple freezing schedule, *e.g.* adding 10% of the instances to the label-frozen set, can already bring notable benefits to the propagation process.

Note that the progressive scheme not only reduces computational cost but also improves propagation accuracy. The reason is that without freezing, the noise and the uncertain nodes will keep affecting all the other nodes, which can sometimes cause additional errors. Experiments in Sect. 5.3 will show more details.

5 Experiments

5.1 Evaluation Protocol and Metrics of CSM

The 192 movies in CSM are partitioned into training (train), validation (val) and testing (test) sets. Statistics of these sets are shown in Table 2. Note that we make sure that there is no overlap between the cast of different sets. *i.e.* the cast in the testing set would not appear in training and validation. This ensures the reliability of the testing results.

Under the Person Search with One Portrait setting, one should rank all the tracklets in the gallery given a query. For this task, we use *mean Average Precision (mAP)* as the evaluation metric. We also report the recall of tracklet identification results in our experiments in terms of R@k. Here, we rank the identities for each tracklet according to their probabilities. R@k means the fraction of tracklets for which the correct identity is listed within the top k results.

We consider two test settings in the CSM benchmark named “search cast in a movie” (IN) and “search cast across all movies” (ACROSS). The setting “IN” means the gallery consists of just the tracklets from one movie, including

the tracklets of the credited cast and those of “others”. While in the “ACROSS” setting, the gallery comprises all the tracklets of credited cast in testing set. Here we exclude the tracklets of “others” in the “ACROSS” setting because “others” just means that it does not belong to any one of the credited cast of a particular movie rather than all the movies in the dataset as we have mentioned in Sect. 3. Table 3 shows the query/gallery sizes of each setting.

5.2 Implementation Details

We use two kinds of visual features in our experiments. The first one is the IDE feature [44] widely used in person re-id. The IDE descriptor is a CNN feature of the whole person instance, extracted by a Resnet-50 [12], which is pre-trained on ImageNet [31] and finetuned on the training set of CSM. The second one is the *face feature*, extracted by a Resnet-101, which is trained on MS-Celeb-1M [10]. For each instance, we extract its IDE feature and the face feature of the face region, which is detected by a face detector [42]. All the visual similarities in experiments are calculated by cosines similarity between the visual features.

Table 4. Results on CSM under two test settings

	IN				ACROSS			
	mAP	R@1	R@3	R@5	mAP	R@1	R@3	R@5
FACE	53.33	76.19	91.11	96.34	42.16	53.15	61.12	64.33
IDE	17.17	35.89	72.05	88.05	1.67	1.68	4.46	6.85
FACE+IDE	53.71	74.99	90.30	96.08	40.43	49.04	58.16	62.10
LP	8.19	39.70	70.11	87.34	0.37	0.41	1.60	5.04
PPCC-v	62.37	84.31	94.89	98.03	59.58	63.26	74.89	78.88
PPCC-vt	63.49	83.44	94.40	97.92	62.27	62.54	73.86	77.44

5.3 Results on CSM

We set up four baselines for comparison: **(1) FACE:** To match the portrait with the tracklet in the gallery by face feature similarity. Here we use the mean feature of all the instances in the tracklet to represent it. **(2) IDE:** Similar to FACE, except that the IDE features are used rather than the face features. **(3) IDE+FACE:** To combine face similarity and IDE similarity for matching, respectively with weights 0.8 and 0.2. **(4) LP:** Conventional label propagation with linear diffusion with both visual and temporal links. Specifically, we use face similarity as the visual links between portraits and candidates and the IDE similarity as the visual links between different candidates. We also consider two settings of the proposed Progressive Propagation via Competitive Consensus method. **(5) PPCC-v:** using only visual links. **(6) PPCC-vt:** the full config with both visual and temporal links.

From the results in Table 4, we can see that: (1) Even with a very powerful CNN trained on a large-scale dataset, matching portrait and candidates by visual cues cannot solve the person search problem well due to the big gap of visual appearances between the portraits and the candidates. Although face features are generally more stable than IDE features, they would fail when the faces are invisible, which is very common in real-world videos like movies. (2) Label propagation with linear diffusion gets very poor results, even worse than the matching-based methods. (3) Our approach raises the performance by a considerable margin. Particularly, the performance gain is especially remarkable on the more challenging “ACROSS” setting (62.27 with ours vs. 42.16 with the visual matching method).

Analysis on Competitive Consensus. To show the effectiveness of *Competitive Consensus*, we study different settings of the Competitive Consensus scheme in two aspects: (1) The max in Eq. (3) can be relaxed to top- k average. Here k indicates the number of neighbors to receive information from. When $k = 1$, it reduces to only taking the maximum, which is what we use in PPCC. Performances obtained with different k are shown in Fig. 6. (2) We also study on the “softmax” in Eq. (3) and compare results between different temperatures of it. The results are also shown in Fig. 6. Clearly, using smaller temperature of softmax significantly boosts the performance. This study supports what we have claimed when designing *Competitive Consensus*: we should only propagate the most confident information in this task.

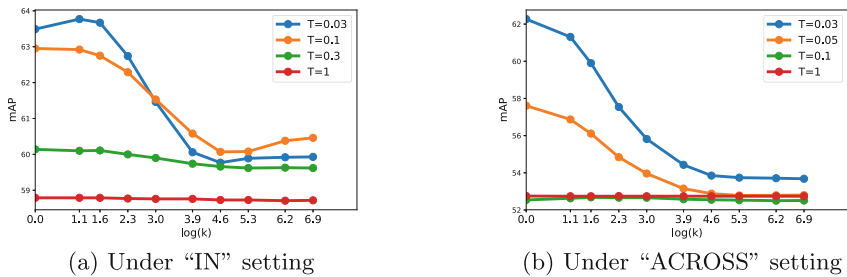


Fig. 6. mAP of different settings of competitive consensus. Comparison between different temperatures(T) of softmax and different settings of k (in top- k average).

Table 5. Results of different updating schemes

	IN				ACROSS			
	mAP	R@1	R@3	R@5	mAP	R@1	R@3	R@5
Conventional	60.54	76.64	91.63	96.70	57.42	54.60	63.31	66.41
Threshold	62.51	81.04	93.61	97.48	61.20	61.54	72.31	76.01
Step	63.49	83.44	94.40	97.92	62.27	62.54	73.86	77.44



Fig. 7. Some samples that are correctly searched in different iterations.

Analysis on Progressive Propagation. Here we show the comparison between our progressive updating scheme and the conventional scheme that updates all the nodes at each iteration. For progressive propagation, we try two kinds of freezing mechanisms: (1) *Step* scheme means that we set the freezing ratio of each iteration and the ratio are raised step by step. More specifically, the freezing ratio r is set to $r = 0.5 + 0.1 \times \text{iter}$ in our experiment. (2) *Threshold* scheme means that we set a threshold, and each time we freeze the nodes whose max probability to a particular identity is greater than the threshold. In our experiments, the threshold is set to 0.5. The results are shown in Table 5, from which we can see the effectiveness of the progressives scheme.

Case Study. We show some samples that are correctly searched in different iterations in Fig. 7. We can see that the easy cases, which are usually with clear frontal faces, can be identified at the beginning. And after iterative propagation, the information can be propagated to the harder samples. At the end of the propagation, even some very hard samples, which are non-frontal, blurred, occluded and under extreme illumination, can be propagated a right identity.

6 Conclusion

In this paper, we studied a new problem named *Person Search in Videos with One Portrait*, which is challenging but practical in the real world. To promote the research on this problem, we construct a large-scale dataset *CSM*, which contains 127K tracklets of 1,218 cast from 192 movies. To tackle this problem, we proposed a new framework that incorporates both visual and temporal links

for identity propagation, with a novel *Progressive Propagation vis Competitive Consensus* scheme. Both quantitative and qualitative studies show the challenges of the problem and the effectiveness of our approach.

Acknowledgement. This work is partially supported by the Big Data Collaboration Research grant from SenseTime Group (CUHK Agreement No. TS1610626), the General Research Fund (GRF) of Hong Kong (No. 14236516).

References

1. Ahmed, E., Jones, M., Marks, T.K.: An improved deep learning architecture for person re-identification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 3908–3916 (2015)
2. Apostolidis, E., Mezaris, V.: Fast shot segmentation combining global and local visual descriptors. In: 2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 6583–6587. IEEE (2014)
3. Cheng, D., Gong, Y., Zhou, S., Wang, J., Zheng, N.: Person re-identification by multi-channel parts-based cnn with improved triplet loss function. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1335–1344 (2016)
4. Ding, S., Lin, L., Wang, G., Chao, H.: Deep feature learning with relative distance comparison for person re-identification. *Pattern Recogn.* **48**(10), 2993–3003 (2015)
5. Feris, R., Bobbitt, R., Brown, L., Pankanti, S.: Attribute-based people search: lessons learnt from a practical surveillance system. In: Proceedings of International Conference on Multimedia Retrieval, p. 153. ACM (2014)
6. Gheissari, N., Sebastian, T.B., Hartley, R.: Person reidentification using spatiotemporal appearance. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, vol. 2, pp. 1528–1535. IEEE (2006)
7. Gong, S., Cristani, M., Yan, S., Loy, C.C. (eds.): *Person Re-Identification*. ACVPR. Springer, London (2014). <https://doi.org/10.1007/978-1-4471-6296-4>
8. Gou, M., Karanam, S., Liu, W., Camps, O., Radke, R.J.: DukeMTMC4ReID: a large-scale multi-camera person re-identification dataset. In: IEEE Conference on Computer Vision and Pattern Recognition Workshops (2017)
9. Gray, D., Tao, H.: Viewpoint invariant pedestrian recognition with an ensemble of localized features. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008. LNCS, vol. 5302, pp. 262–275. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-88682-2_21
10. Guo, Y., Zhang, L., Hu, Y., He, X., Gao, J.: Ms-celeb-1m: challenge of recognizing one million celebrities in the real world. *Electron. Imaging* **2016**(11), 1–6 (2016)
11. Hamdoun, O., Moutarde, F., Stanciulescu, B., Steux, B.: Person re-identification in multi-camera system by signature based on interest point descriptors collected on short video sequences. In: Second ACM/IEEE International Conference on Distributed Smart Cameras (ICDSC 2008), pp. 1–6. IEEE (2008)
12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)

13. Hirzer, M., Beleznai, C., Roth, P.M., Bischof, H.: Person re-identification by descriptive and discriminative classification. In: Heyden, A., Kahl, F. (eds.) SCIA 2011. LNCS, vol. 6688, pp. 91–102. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-21227-7_9
14. Huang, Q., Xiong, Y., Lin, D.: Unifying identification and context learning for person recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2217–2225 (2018)
15. Joon Oh, S., Benenson, R., Fritz, M., Schiele, B.: Person recognition in personal photo collections. In: Proceedings of the IEEE International Conference on Computer Vision, pp. 3862–3870 (2015)
16. Karanam, S., Gou, M., Wu, Z., Rates-Borras, A., Camps, O., Radke, R.J.: A systematic evaluation and benchmark for person re-identification: features, metrics, and datasets. arXiv preprint [arXiv:1605.09653](https://arxiv.org/abs/1605.09653) (2016)
17. Koestinger, M., Hirzer, M., Wohlhart, P., Roth, P.M., Bischof, H.: Large scale metric learning from equivalence constraints. In: 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 2288–2295. IEEE (2012)
18. Kumar, V., Nambodiri, A.M., Jawahar, C.: Face recognition in videos by label propagation. In: 2014 22nd International Conference on Pattern Recognition (ICPR), pp. 303–308. IEEE (2014)
19. Li, H., Brandt, J., Lin, Z., Shen, X., Hua, G.: A multi-level contextual model for person recognition in photo albums. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1297–1305 (2016)
20. Li, H., Lu, H., Lin, Z., Shen, X., Price, B.: Inner and inter label propagation: salient object detection in the wild. *IEEE Trans. Image Process.* **24**(10), 3176–3186 (2015)
21. Li, S., Xiao, T., Li, H., Zhou, B., Yue, D., Wang, X.: Person search with natural language description. In: Proceedings of the CVPR (2017)
22. Li, W., Zhao, R., Xiao, T., Wang, X.: Deepreid: deep filter pairing neural network for person re-identification. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 152–159 (2014)
23. Liao, S., Hu, Y., Zhu, X., Li, S.Z.: Person re-identification by local maximal occurrence representation and metric learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2197–2206 (2015)
24. Lin, D., Kapoor, A., Hua, G., Baker, S.: Joint people, event, and location recognition in personal photo collections using cross-domain context. In: Daniilidis, K., Maragos, P., Paragios, N. (eds.) ECCV 2010. LNCS, vol. 6311, pp. 243–256. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15549-9_18
25. Lin, T.-Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: Fleet, D., Pajdla, T., Schiele, B., Tuytelaars, T. (eds.) ECCV 2014. LNCS, vol. 8693, pp. 740–755. Springer, Cham (2014). https://doi.org/10.1007/978-3-319-10602-1_48
26. Ma, B., Su, Y., Jurie, F.: Local descriptors encoded by fisher vectors for person re-identification. In: Fusiello, A., Murino, V., Cucchiara, R. (eds.) ECCV 2012. LNCS, vol. 7583, pp. 413–422. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-33863-2_41
27. Ma, B., Su, Y., Jurie, F.: Covariance descriptor based on bio-inspired features for person re-identification and face verification. *Image Vis. Comput.* **32**(6–7), 379–390 (2014)
28. Prosser, B.J., Zheng, W.S., Gong, S., Xiang, T., Mary, Q.: Person re-identification by support vector ranking. In: BMVC, vol. 2, p. 6 (2010)

29. Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: towards real-time object detection with region proposal networks. In: *Advances in Neural Information Processing Systems*, pp. 91–99 (2015)
30. Rohrbach, M., Ebert, S., Schiele, B.: Transfer learning in a transductive setting. In: *Advances in nEural Information Processing Systems*, pp. 46–54 (2013)
31. Russakovsky, O., et al.: Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**(3), 211–252 (2015)
32. Sener, O., Song, H.O., Saxena, A., Savarese, S.: Learning transferrable representations for unsupervised domain adaptation. In: *Advances in Neural Information Processing Systems*, pp. 2110–2118 (2016)
33. Sheikh, R., Garbade, M., Gall, J.: Real-time semantic segmentation with label propagation. In: Hua, G., Jégou, H. (eds.) *ECCV 2016. LNCS*, vol. 9914, pp. 3–14. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-48881-3_1
34. Sidiropoulos, P., Mezaris, V., Kompatsiaris, I., Meinedo, H., Bugalho, M., Trancoso, I.: Temporal video segmentation to scenes using high-level audiovisual features. *IEEE Trans. Circ. Syst. Video Technol.* **21**(8), 1163–1177 (2011)
35. Su, C., Zhang, S., Xing, J., Gao, W., Tian, Q.: Deep attributes driven multi-camera person re-identification. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016. LNCS*, vol. 9906, pp. 475–491. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46475-6_30
36. Tripathi, S., Belongie, S., Hwang, Y., Nguyen, T.: Detecting temporally consistent objects in videos through object class label propagation. In: *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1–9. IEEE (2016)
37. Wang, J., Jebara, T., Chang, S.F.: Graph transduction via alternating minimization. In: *Proceedings of the 25th International Conference on Machine Learning*, pp. 1144–1151. ACM (2008)
38. Wang, T., Gong, S., Zhu, X., Wang, S.: Person re-identification by discriminative selection in video ranking. *IEEE Trans. Pattern Anal. Mach. Intell.* **38**(12), 2501–2514 (2016)
39. Xiao, T., Li, H., Ouyang, W., Wang, X.: Learning deep feature representations with domain guided dropout for person re-identification. In: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1249–1258. IEEE (2016)
40. Xiao, T., Li, S., Wang, B., Lin, L., Wang, X.: Joint detection and identification feature learning for person search. In: *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3376–3385. IEEE (2017)
41. Zajdel, W., Zivkovic, Z., Krose, B.: Keeping track of humans: have i seen this person before? In: *Proceedings of the 2005 IEEE International Conference on Robotics and Automation, ICRA 2005*. pp. 2081–2086. IEEE (2005)
42. Zhang, K., Zhang, Z., Li, Z., Qiao, Y.: Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Sig. Process. Lett.* **23**(10), 1499–1503 (2016)
43. Zhang, N., Paluri, M., Taigman, Y., Fergus, R., Bourdev, L.: Beyond frontal faces: improving person recognition using multiple cues. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4804–4813 (2015)
44. Zheng, L., et al.: MARS: a video benchmark for large-scale person re-identification. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) *ECCV 2016. LNCS*, vol. 9910, pp. 868–884. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-46466-4_52
45. Zheng, L., Shen, L., Tian, L., Wang, S., Wang, J., Tian, Q.: Scalable person re-identification: a benchmark. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1116–1124 (2015)

46. Zhou, D., Bousquet, O., Lal, T.N., Weston, J., Schölkopf, B.: Learning with local and global consistency. In: *Advances in Neural Information Processing Systems*, pp. 321–328 (2004)
47. Zhu, X., Ghahramani, Z.: Learning from labeled and unlabeled data with label propagation (2002)
48. Zoidi, O., Tefas, A., Nikolaidis, N., Pitas, I.: Person identity label propagation in stereo videos. *IEEE Trans. Multimedia* **16**(5), 1358–1368 (2014)