# Video Imprint

Zhanning Gao<sup>®</sup>, Le Wang<sup>®</sup>, *Member, IEEE*, Nebojsa Jojic, Zhenxing Niu, *Member, IEEE*, Nanning Zheng, *Fellow, IEEE*, and Gang Hua<sup>®</sup>, *Senior Member, IEEE* 

Abstract—A new unified video analytics framework (ER3) is proposed for complex event retrieval, recognition and recounting, based on the proposed video imprint representation, which exploits temporal correlations among image features across video frames. With the video imprint representation, it is convenient to reverse map back to both temporal and spatial locations in video frames, allowing for both key frame identification and key areas localization within each frame. In the proposed framework, a dedicated feature alignment module is incorporated for redundancy removal across frames to produce the tensor representation, i.e., the video imprint. Subsequently, the video imprint is individually fed into both a reasoning network and a feature aggregation module, for event recognition/recounting and event retrieval tasks, respectively. Thanks to its attention mechanism inspired by the memory networks used in language modeling, the proposed reasoning network is capable of simultaneous event category recognition and localization of the key pieces of evidence for event recounting. In addition, the latent structure in our reasoning network highlights the areas of the video imprint, which can be directly used for event recounting. With the event retrieval task, the compact video representation aggregated from the video imprint contributes to better retrieval results than existing state-of-the-art methods.

Index Terms—Event videos, feature alignment, feature aggregation, reasoning network

## **1** INTRODUCTION

A NALYSIS of event videos is generally considered more challenging than the related video-based action recognition task [6], [41], thanks to the richer contents of such event videos. Typical event videos are much longer (several minutes or even hours) than trimmed action recognition videos, and multiple human actions and a variety of different objects often appear across various scenes. For example, a "birthday party" event may take place at home or in a restaurant, with multiple objects coming into focus, e.g., a birthday cake, and may include a variety of activities that span multiple frames, e.g., singing the birthday song, or blowing out candles.

In the last decade, analysis of complex events in videos has attracted significant attention in the computer vision community [5], [10], [13], [14], [22], [31], [38], [44], [49]. Previous research could be categorized as the unsupervised and the supervised methods. Unsupervised methods were typically used for *event retrieval* [12], [38] where the goal is the retrieval of all related videos semantically relevant to the query video sample. On the other hand, supervised learning has been widely used in *event recognition* [4], [7] and detection tasks [31], [57] in similar ways to its applications in

Recommended for acceptance by H. Ling.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below. Digital Object Identifier no. 10.1109/TPAMI.2018.2866114 action recognition [6], [41], [61] and generic video classification tasks [23], [56], [59]. In the latter case, a classifier is trained on the annotated training set to recognize the event categories of test videos, e.g., the multimedia event detection task of the TRECVID [33]. In practical applications, it is often desirable to provide explainable results by qualifying the category prediction with the localization of the key pieces of evidence that lead to the recognition decision, which is sometimes referred to as the *event recounting*.

A major challenge in *event retrieval* and *event recognition* is the construction of appropriate video representations, which should ideally be both discriminative for efficient disambiguation and compact for computational efficiency. Conventionally, a video representation is a fixed-length per-video global feature vector extracted from many frame-level appearance features [12], [13], [38], [56], [57]. In a typical *event recognition* task, this video representation is fed into a linear classifier [57] or a neural network [23], [56] for classification. However, this procedure is generally incompatible with event recounting, as tracing the decision back to individual frame locations is impossible due to the irreversible video representation (global feature) extraction. Therefore, most existing methods perform event recounting as an extra post-processing step [13], [29], [50].

To address these challenges, the ER3 framework is proposed to simultaneously achieve *event retrieval, recognition* and *recounting.* Fig. 1 illustrates the components and the inputs/outputs of such ER3 system. In ER3, (i) we introduce a feature alignment step which can significantly suppress the redundant information and generate a more comprehensive and compact video representation called *video imprint.* In addition, the video imprint also preserves the local spatial layout among video frames. (ii) Based on the video imprint, in unsupervised setting, we propose an efficient aggregation method for large-scale event retrieval. In

0162-8828 © 2018 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications\_standards/publications/rights/index.html for more information.

Z. Gao, L. Wang, and N. Zheng are with the Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University, Xi'an, Shaanxi 710049, China. E-mail: zhanninggao@gmail.com, {lewang, nnzheng}@mail.xjtu.edu.cn.

N. Jojic and G. Hua are with the Microsoft Research, Redmond, WA 98052. E-mail: jojic@microsoft.com, ganghua@gmail.com.

Z. Niu is with the Alibaba Group, Hangzhou, Zhejiang 311121, China. E-mail: zhenxing.nzx@alibaba-inc.com.

Manuscript received 16 Dec. 2017; revised 12 Aug. 2018; accepted 12 Aug. 2018. Date of publication 20 Aug. 2018; date of current version 31 Oct. 2019. (Corresponding author: Le Wang.)



Fig. 1. Illustration of the ER3 framework for event retrieval, recognition and recounting. The compact video representation from feature aggregation can be used for large-scale event retrieval. With supervised training, ER3 can also recognize the event category of the input video. Event recounting falls directly out of the latent structure of the model in form of statistics displayed as heat maps for each frame indicating key areas related to the event.

supervised setting, we further employ a reasoning network, a modified version of the neural memory networks [43], which can simultaneously recognize the event category and locate the key pieces of evidence for the event category. In fact, the recounting is so naturally integrated in the framework that the experiments show that the recounting step can assist the recognition task and improve the recognition accuracy. (iii) In the recounting task, both temporal key frame identification (attribution of the key frames with respect to the event category, as in [29], [50]) and spatial key areas localization (attribution of the key areas within each frame) are implemented, thanks to the video imprint preserving local semantic and spatial layouts.

This manuscript is an extension of our conference paper [15] with modifications as follows. In the feature alignment step, an alternative generative model (i.e., epitome [25]) is included besides the the tessellated counting grid (TCG) model [34], [35]. Both generative models share the core idea of building condensed representation by exploiting the spatial interdependence among the input features. However, unlike the TCG model which is limited to counts/histogram-style input features, the epitome model is represented as structured Gaussian mixtures which can accommodate general vector or tensor input features. In addition, to accelerate the feature alignment step with the epitome model, we propose an accelerated two-step scheme to update the epitome. More details are provided in Section 3.2. We also provide a comprehensive comparison between the two generative models with various datasets and tasks. The experimental results show that the alternative epitome model achieves higher computational efficiency with comparable results with the TCG model.

The paper is organized as follows. Section 2 discusses related work about event videos analysis. Then, we present the technical details of the ER3 in Section 3. Experimental results are provided in Section 4. Finally, we conclude the paper in Section 5.

#### 3087

## 2 RELATED WORK

With a typical unsupervised event retrieval task, the goal is the retrieval of all related videos semantically relevant to the query video sample. A major challenge is the construction of both compact and discriminative video representations. Conventional methods [12], [37], [38] rely on frame-level local features (e.g., SIFT [30]) and aggregation strategies (e.g., Fisher Vector [36], [39], VLAD [12], [20], [21] or explicit feature maps [37]) for frame-level feature description. Recent works [23], [57] predominantly employ deep Convolutional Neural Network (CNNs) [18], [42] to extract a feature descriptor from each video frame. Latest work [3] revisits temporal match kernels [37] and presents a learnable temporal layer which can further enhance the CNNs based feature descriptor. Subsequently, a video-level representation is typically obtained by directly averaging all frame-level descriptors in the video. Such sum-aggregation strategy disregards the strong temporal correlations among consecutive frames, which may over-emphasize certain long or recurrent shots in the video. We discuss this problem in Section 3 and show that the redundant information among frames can be effectively suppressed by the feature alignment step.

Event recognition and detection have attracted wide attention in the last decade. In general, event video recognition system usually consists of three stages, i.e., feature extraction, feature aggregation/pooling and training/recognition. As in event retrieval, the first two stages aim at building discriminative video representations. Previous methods focused mostly on designing better video features or representations for the classifier, such as hand-crafted visual features [11], [30], motion features [51], [52], audio features [2], and mid-level concept/attribute features [9], [48]. Recently, the advancements of CNN [27], [42] lead to promising results in event recognition task [23], [57], [60]. The video representations are usually constructed by direct aggregation of the frame-level CNN features. Due to limited amount of training data, video representations are typically fed to classifiers such as the Support Vector Machine (SVM) [8] or multi-feature fusion framework [23], [56], [60], [62].

Event recounting refers to the attribution of key pieces of evidence supporting the recognition decision. As with most video analytics datasets, only video-level annotations are provided, making such attribution a challenging task. Event recounting is usually implemented as a post-processing step after the recognition [29], [50]. Chang et al. [9] proposed a joint optimization framework with mid-level semantic concept representations for event recognition and recounting. Sun et al. [45] introduced an evidence localization model learned via a max-margin framework, and Lai et al. [28] applied Multiple-Instance Learning (MIL) which infers temporal instance labels and the video-level labels. In these works, event videos are treated as sets of shots or instances, these recounting procedures perform only temporal localization and usually at a coarse scale. Recently, Gan et al. [13] proposed a deep neural network for event recognition. Specifically, event recounting (both temporal and spatial) is also achieved via passing the classification scores backward. However, this is still an explanatory post-processing step which is never designed to assist the event recognition task.

In contrast with these methods, at the core of our integrated event recognition and recounting framework is a



Fig. 2. Illustration of the frames related to "Dominique Strauss-Kahn arrested". The frames in green box denote the positive frames related to the event. Red boxes show irrelevant frames.

latent structure that contains reverse attribution pointers back to the video frames. In addition, the proposed framework is trained in an unsupervised manner by simultaneously aligning areas across frames and estimating a probability distribution over frame features in the corresponding areas. The obtained representation consists of a grid of distributions with the corresponding mappings from the feature to the frames, similar to the way video frames are mapped to panoramas from pixel space (see the toy examples in Figs. 1 and 3). The obtained grid and such reverse mapping pointers (back to the spatial locations of the specific video frames) form the proposed *video imprint*.



Fig. 3. Illustration of tessellated counting grid (TCG) and epitome models, and their Bayesian networks. The left tensor block represents the TCG with  $\mathbf{E} = 24 \times 24$ ,  $\mathbf{W} = 8 \times 8$ ,  $\mathbf{S} = 4 \times 4$ . The right tensor block represents the epitome with  $\mathbf{E} = 24 \times 24$ ,  $\mathbf{W} = 8 \times 8$ . For TCG, the input CNN feature maps are down-sampled to  $\mathbf{S} = 4 \times 4$ . In fact, the epitome can be regarded as a special case of TCG when  $\mathbf{W} = \mathbf{S}$ . For both TCG and epitome, similar frames are usually represented in the same or nearby windows, e.g., the anchor who we frequently see in the video.

With this video imprint representation, it is possible to design an aggregation to emphasize mere presence instead of frequency of repetition.

Our experiments show that the proposed video imprint aggregation yields better performances in both the supervised and unsupervised tasks than existing algorithms. The video imprint also allows for the reasoning over the spatial layouts of features across frames. Inspired by the attention mechanism in memory networks [43] reasoning over sentences in priming text and video face recognition [58], the proposed reasoning network analyzes evidences at different spatial locations of the compact video imprint while carrying out the recognition task, which also highlights the key areas of the imprint. These key areas of the imprint could readily be mapped back to specific video frames and corresponding spatial locations. In this way, event recounting is implemented as an integral part of recognition instead of a post-processing step.

## 3 THE DETAILS OF ER3

In this section, we present all modules and the operating mechanism of the proposed ER3 framework, as illustrated in Fig. 1.

### 3.1 Feature Extraction

Recently, image descriptors based on the activations of convolutional layers [17], [40], [60] have outperformed previous methods [23], [60] based on features extracted from fully connected layers. Inspired by the spatial information preserving characteristics of convolutional layers outputs, we also choose the activations of the last convolutional layer as the frame-level feature.

### 3.2 Feature Alignment

Existing event recognition algorithms rely on a compact and discriminative video representation, which is typically direct average of the frame-level descriptors [12], [23], [57], [60]. Such sum-aggregation strategy disregards the strong temporal correlations among consecutive frames, which may over-emphasize certain long or recurrent shots in the video. Therefore, irrelevant and repetitive shots in the video might dominate the obtained video representation. For instance, in the event "Dominique Strauss-Kahn arrested" as shown in Fig. 2, many video frames showing the news anchor are irrelevant, but they are visually similar, therefore the simple averaging aggregation strategy for video representation may over-emphasize such irrelevant frames. To

TABLE 1	
Principal Notatio	ns

$\pi_{i}$	Parameters of the counting grid model. $l_1$ -normalized
	counts feature at location i on the counting grid
	$\mathbf{E} = [1 \dots E_x] \times [1 \dots E_y]$
$\pi_{\mathbf{i},z}$	<i>z</i> th dimension of $\pi_i$ and $\sum_{z \in \mathbb{Z}} \pi_{i,z} = 1$ , where
	$z \in \mathbf{Z} = [1, \dots, Z]$
$(\boldsymbol{\mu}_{\mathbf{i}}, \boldsymbol{\sigma}_{\mathbf{i}}^2)$	Parameters of the epitome model. The mean $\mu_{ m i}$ and
	variance $\sigma_i^2$ of the Gaussian distribution aligned at the
	location i of the grid E
$\{\mathbf{c}^s\}_{s\in\mathbf{S}}$	Counts features plugged in a tessellation S, where
0,000	$\mathbf{s} \in \mathbf{S} = [1 \dots S_x] \times [1 \dots S_y]$
F	Tensor features, e.g., the activations from convolutional
	layer of the CNN model
$f_{i}$	The feature vector extracted from <b>F</b> along the spatial
• )	dimensions, i.e., $\mathbf{j} \in \mathbf{W} = [1, \dots, W_x] \times [1, \dots, W_y]$
Wk	The window at the location <b>k</b> of the grid <b>E</b> which assumed
	to generate counts features $\{c^s\}_{s\in S}$ or tensor features F
W <sup>s</sup> <sub>k</sub>	The sub-window in the $W_k$ generating each $c^s$
l	The latent variable that represents the mapping location in
	the grid E
	0

mitigate this problem, we propose the feature alignment procedure to regularize the influence of frame features.

The idea of feature alignment is inspired by panoramic stitching [46], [47], which can remove the redundant or overlapping parts between multiple images. The redundancy across video frames could likewise be significantly reduced to improve the robustness of the obtained video representation against long and/or repetitive irrelevant video frames.

Due to the dynamic and complex nature of event videos, it is impractical to directly stitch video frames by pixel-level alignment. Instead, image features are extracted first using the activations of the last convolutional layer with each frame as input. Afterwards, the tessellated counting grid [34], [35] or epitome model [25] is employed to generate a tensor of frame-level feature distributions, which implies a panorama-style reversible mapping, accommodating the geometric variations in objects and scenes in event videos. This model exploits the spatial interdependence of the frame-level features in relevant frames, which makes possible the subsequent clustering of visually similar shots.

The inquisition for leveraging the epitome model [25] is to accelerate the feature alignment step. Both the TCG and the epitome can capture the spatial interdependence among input features. However, the input features of the epitome model can be more flexible, considering their Gaussian location distribution instead of the discrete categorical distribution in the TCG. In addition, the incorporation of the epitome model also allows us to propose a highly efficient two-step implementation, which leads to approximately an order of magnitude faster feature alignment procedure than the counterpart in the TCG. In the remainder of this section, mathematical notations are first summarized in Table 1, followed by introductions of both generative models.<sup>1</sup>

#### 3.2.1 Tessellated Counting Grid

Tessellated counting grid [34], [35] is designed to capture the spatial interdependence among image features. Given a set of images or a video sequence, it assumes that each image/frame is represented by a set of  $l_1$ -normalized non-negative feature vectors (e.g., bags of visual words vectors)  $\{\mathbf{c}^s\}_{s\in\mathbf{S}}$  plugged in a tessellation  $\mathbf{S} = [1, \ldots, S_x] \times [1, \ldots, S_y]$ .<sup>2</sup>

Formally, the counting grid  $\pi_i$  is a set of normalized counts of features indexed by  $z \in \mathbf{Z} = [1 \dots Z]$  (dimension of image feature) on the 2D discrete grid  $\mathbf{i} = (i_x, i_y) \in \mathbf{E} =$  $[1, \dots, E_x] \times [1, \dots, E_y]$ , where  $\mathbf{i}$  denotes the location on the grid and  $\sum_{z \in \mathbf{Z}} \pi_{\mathbf{i},z} = 1$  [26].

As a generative model, the image features  $\{\mathbf{c}^s\}_{s\in\mathbf{S}}$  are assumed to follow a distribution found in a window into the counting grid. The probability of generating the image features  $\{\mathbf{c}^s\}_{s\in\mathbf{S}}$  from the window  $\mathbf{W}_{\mathbf{k}} = [k_x, \ldots, k_x + W_x - 1] \times [k_y, \ldots, k_y + W_y - 1]$  placed at the location  $\mathbf{k} = (k_x, k_y) \in \mathbf{E}$  of the grid is

$$p(\{\mathbf{c}^s\}_{s\in\mathbf{S}}|l=\mathbf{k}) = \gamma \prod_{z\in\mathbf{Z}} \prod_{s\in\mathbf{S}} \left(\sum_{\mathbf{i}\in\mathbf{W}^s_{\mathbf{k}}} \pi_{\mathbf{i},z}\right)^{c_z^s},$$
(1)

where  $\gamma$  is the normalization constant. l denotes the latent variable, while **i** and **k** represent generic positions in the grid **E**. Then, for a given counting grid  $\pi$ , the joint distribution over the set of image features  $\{\mathbf{c}^{s,t}\}_{s\in\mathbf{S},t\in\mathbf{T}}$ , indexed by  $t \in \mathbf{T} = [1...T]$ , and their corresponding latent window locations  $\{l^t\}$  in the grid can be derived as

$$P(\{\mathbf{c}^{s,t}\}_{s\in\mathbf{S},t\in\mathbf{T}},\{l^t\}_{t\in\mathbf{T}}) \propto \prod_{t\in\mathbf{T}} \sum_{\mathbf{k}\in\mathbf{E}} \prod_{z\in\mathbf{Z}} \prod_{s\in\mathbf{S}} \left(\sum_{\mathbf{i}\in\mathbf{W}^s_{\mathbf{k}}} \pi_{\mathbf{i},z}\right)^{c_z^{s,t}}.$$
 (2)

The counting grid  $\pi$  can be estimated by maximizing the log likelihood of the joint distribution with an EM algorithm

E step : 
$$q(l^t = \mathbf{k}) \propto \exp\left(\sum_{s \in \mathbf{S}} \sum_{z \in \mathbf{Z}} c_z^{s,t} \log \sum_{\mathbf{i} \in \mathbf{W}_{\mathbf{k}}^s} \pi_{\mathbf{i},z}\right),$$
  
M step :  $\pi_{\mathbf{i},z} \propto \pi_{\mathbf{i},z}^{old} \sum_{t \in \mathbf{T}} \sum_{s \in \mathbf{S}} c_z^{s,t} \sum_{\mathbf{k} \mid \mathbf{i} \in \mathbf{W}_{\mathbf{k}}^s} \frac{q(l^t = \mathbf{k})}{\sum_{\mathbf{i} \in \mathbf{W}_{\mathbf{k}}^s} \pi_{\mathbf{i},z}^{old}},$ 
(3)

where  $q(l^t = \mathbf{k})$  denotes the posterior probability  $p(l^t = \mathbf{k} | \{\mathbf{c}^{s,t}\}_{s \in \mathbf{S}})$  and  $\pi_{\mathbf{i},z}^{old}$  denotes the counting grid at the previous iteration.

The iterative process of TCG jointly estimates the counting grid (i.e., *video imprint*)  $\pi$  and aligns all video frame features to it with such correspondences captured in *q*.

## 3.2.2 Epitome

The original epitome model [25] takes raw pixels as input and aims to mine the essence of the textural and shape properties of the image. Formally, the epitome **e** is a set of dependent Gaussian distributions  $\{\mathcal{N}(f_j; \boldsymbol{\mu}_i, \sigma_i^2)\}$  aligned on a grid  $\mathbf{i} = (i_x, i_y) \in \mathbf{E} = [1, \dots, E_x] \times [1, \dots, E_y]$  just like TCG. In the original epitome formulation,  $f_j$  denotes the intensity or the color of the pixel on the image patch **F** indexed by  $\mathbf{j} = (j_x, j_y) \in \mathbf{W} = [1, \dots, W_x] \times [1, \dots, W_y]$ . Here we extend **F** to a general tensor (feature map), specifically,

<sup>1.</sup> Comprehensive introductions to both the TCG and the epitome models can be find in  $\cite{25}\cite{$ 

<sup>2.</sup> With  $l_1$ -normalization and appropriate down-sampling, the feature maps (after ReLU) from the convolutional layer of CNN model naturally satisfy this assumption.

the activations from the last convolutional layer of a CNN model. In other words, **F** is the CNN feature map and  $f_j$  is the feature vector extracted from **F** along the spatial dimensions, i.e.,  $\mathbf{j} \in \mathbf{W}$ .

Given the epitome e, the probability of generating the feature map F from the window  $W_k$  at location k of the epitome e is

$$p(\mathbf{F}|l = \mathbf{k}) = \gamma \prod_{\mathbf{i} \in \mathbf{W}_{\mathbf{k}}} \mathcal{N}(f_{\mathbf{i}-\mathbf{k}}; \boldsymbol{\mu}_{\mathbf{i}}, \boldsymbol{\sigma}_{\mathbf{i}}^{2}),$$
(4)

where  $\gamma$  is the normalization constant.  $\mathcal{N}(f_{i-k}; \mu_i, \sigma_i^2)$  is a Gaussian distribution over  $f_{i-k}$  with mean  $\mu_i$  and variance  $\sigma_i^2$  and  $\mathbf{i} - \mathbf{k} = (i_x - k_x + 1, i_y - k_y + 1)$ . Similar to TCG, the joint distribution over the set of feature maps  $\{\mathbf{F}^t\}_{t \in \mathbf{T}}$ , indexed by t, and their corresponding latent window locations  $\{l^t\}_{t \in \mathbf{T}}$  on the epitome can be derived as

$$P({\mathbf{F}^t}, {l^t}) \propto \prod_{t \in \mathbf{T}} \sum_{\mathbf{k} \in \mathbf{E}} \prod_{\mathbf{i} \in \mathbf{W}_{\mathbf{k}}} \mathcal{N}(f_{\mathbf{i}-\mathbf{k}}^t; \boldsymbol{\mu}_{\mathbf{i}}, \boldsymbol{\sigma}_{\mathbf{i}}^2).$$
(5)

The parameters  $(\mu, \sigma^2)$  are estimated by marginalizing the joint distribution, i.e., optimizing the log likelihood of the data with an iterative EM algorithm

E step : 
$$q(l^{t} = \mathbf{k}) \propto \prod_{\mathbf{i}\in\mathbf{W}_{\mathbf{k}}} \mathcal{N}(f_{\mathbf{i}-\mathbf{k}}^{t}; \boldsymbol{\mu}_{\mathbf{i}}, \sigma_{\mathbf{i}}^{2}),$$
  
M step :  $\boldsymbol{\mu}_{\mathbf{i}} = \frac{\sum_{t\in\mathbf{T}}\sum_{\mathbf{k}\in\mathbf{W}_{\mathbf{i}-\mathbf{W}}} q(l^{t} = \mathbf{k})f_{\mathbf{i}-\mathbf{k}}^{t}}{\sum_{t\in\mathbf{T}}\sum_{\mathbf{k}\in\mathbf{W}_{\mathbf{i}-\mathbf{W}}} q(l^{t} = \mathbf{k})},$  (6)  
 $\sigma_{\mathbf{i}}^{2} = \frac{\sum_{t\in\mathbf{T}}\sum_{\mathbf{k}\in\mathbf{W}_{\mathbf{i}-\mathbf{W}}} q(l^{t} = \mathbf{k})(f_{\mathbf{i}-\mathbf{k}}^{t} - \boldsymbol{\mu}_{\mathbf{i}})^{2}}{\sum_{t\in\mathbf{T}}\sum_{\mathbf{k}\in\mathbf{W}_{\mathbf{i}-\mathbf{W}}} q(l^{t} = \mathbf{k})},$ 

where  $q(l^t = \mathbf{k})$  denotes the posterior probability  $p(l^t = \mathbf{k} | \mathbf{F}^t)$ and  $\mathbf{i} - \mathbf{W} = (i_x - W_x + 1, i_y - W_y + 1)$ .

As presented above, both generative models (TCG and epitome) build condensed representations by capturing the spatial interdependence among input features. The epitome model differs from the TCG in the location distributions (Gaussian versus discrete categorical). Therefore, the input features of the epitome can be more flexible. See Fig. 3 for the illustration of the TCG and epitome.<sup>3</sup> In addition, according to Equation (6), the updating of the epitome parameters, i.e., the M step, only depends on the current *q* distribution, i.e., the current E step, and the input features. This motivates us to propose an efficient two-step scheme to generate the epitome based video imprint, i.e., we can first efficiently estimate the final *q* distribution, and then calculate the epitome based video imprint directly with Equation (6). The following are the details of the efficient two-step implementation.

## 3.2.3 An Efficient Two-Step Scheme

Given a fixed epitome size **E** and window size **W**, the only feasible way to accelerate the learning procedure of epitome model is to reduce the dimension of the input feature maps  $\{\mathbf{F}^t\}_{t\in\mathbf{T}}$ . Although the low dimensional input features produced from  $\{\mathbf{F}^t\}_{t\in\mathbf{T}}$  by dimension reduction are enough to

capture the correspondences among frames within a video, it may also reduce the discriminative capacity of the epitome based video imprint  $\mu$  for large-scale video event analysis. Therefore, we propose an efficient two-step scheme which divides the epitome based feature alignment operation into two steps: the correspondence analysis step and the video imprint generation step.

At the first step, the low dimensional feature maps  $\{\mathbf{G}^t\}_{t\in\mathbf{T}}$  are generated by PCA [19] projection,  $\mathbf{G} = \phi_{PCA}(\mathbf{F})$ , where  $\mathbf{G} \in \mathbb{R}^{W_x \times W_y \times d}$ ,  $\mathbf{F} \in \mathbb{R}^{W_x \times W_y \times D}$  and  $d \ll D$ . Then, the epitome of  $\{\mathbf{G}^t\}_{t\in\mathbf{T}}$  can be learned with

$$\begin{split} \mathbf{E} \ \text{step} : \ \hat{q}(l^{t} = \mathbf{k}) &\propto \prod_{\mathbf{i} \in \mathbf{W}_{\mathbf{k}}} \mathcal{N}\big(\boldsymbol{g}_{\mathbf{i}-\mathbf{k}}^{t}; \hat{\boldsymbol{\mu}}_{\mathbf{i}}, \hat{\boldsymbol{\sigma}}_{\mathbf{i}}^{2}\big), \\ \mathbf{M} \ \text{step} : \ \hat{\boldsymbol{\mu}}_{\mathbf{i}} &= \frac{\sum_{t \in \mathbf{T}} \sum_{\mathbf{k} \in \mathbf{W}_{\mathbf{i}-\mathbf{W}}} \hat{q}(l^{t} = \mathbf{k}) \boldsymbol{g}_{\mathbf{i}-\mathbf{k}}^{t}}{\sum_{t \in \mathbf{T}} \sum_{\mathbf{k} \in \mathbf{W}_{\mathbf{i}-\mathbf{W}}} \hat{q}(l^{t} = \mathbf{k})}, \\ \hat{\boldsymbol{\sigma}}_{\mathbf{i}}^{2} &= \frac{\sum_{t \in \mathbf{T}} \sum_{\mathbf{k} \in \mathbf{W}_{\mathbf{i}-\mathbf{W}}} \hat{q}(l^{t} = \mathbf{k}) (\boldsymbol{g}_{\mathbf{i}-\mathbf{k}}^{t} - \hat{\boldsymbol{\mu}}_{\mathbf{i}})^{2}}{\sum_{t \in \mathbf{T}} \sum_{\mathbf{k} \in \mathbf{W}_{\mathbf{i}-\mathbf{W}}} \hat{q}(l^{t} = \mathbf{k})}. \end{split}$$
(7)

As an approximation of the *q* distribution, we found that  $\hat{q}$  can fully capture the correspondence of video frames (see Fig. 5). In addition, since  $d \ll D$ ,  $\hat{q}$  estimation with Equation (7) achieves much higher efficiency.

At the second step, according to the M step of the Equation (6), the updated parameters  $\mu$  and  $\sigma^2$  only depend on the set of input feature maps  $\{\mathbf{F}^t\}_{t\in T}$  and the *q* distribution from the last iteration of E step. Therefore, if we replace *q* with its approximation  $\hat{q}$  from the first step,  $\mu$  and  $\sigma^2$  can be directly computed with

$$\mu_{\mathbf{i}} = \frac{\sum_{t \in \mathbf{T}} \sum_{\mathbf{k} \in \mathbf{W}_{\mathbf{i}-\mathbf{W}}} \hat{q}(l^{t} = \mathbf{k}) f_{\mathbf{i}-\mathbf{k}}^{t}}{\sum_{t \in \mathbf{T}} \sum_{\mathbf{k} \in \mathbf{W}_{\mathbf{i}-\mathbf{W}}} \hat{q}(l^{t} = \mathbf{k})},$$

$$\sigma_{\mathbf{i}}^{2} = \frac{\sum_{t \in \mathbf{T}} \sum_{\mathbf{k} \in \mathbf{W}_{\mathbf{i}-\mathbf{W}}} \hat{q}(l^{t} = \mathbf{k}) (f_{\mathbf{i}-\mathbf{k}}^{t} - \mu_{\mathbf{i}})^{2}}{\sum_{t \in \mathbf{T}} \sum_{\mathbf{k} \in \mathbf{W}_{\mathbf{i}-\mathbf{W}}} \hat{q}(l^{t} = \mathbf{k})}.$$
(8)

The computational complexity and performance of the efficient two-step scheme will be discussed in Section 4. We shall note that this fast implementation scheme is inapplicable to TCG. It is difficult to transfer a high dimensional histogram vector to a low dimensional space. Most importantly, as shown in Equation (3), the current  $\pi_i$  relies on both  $\pi_i^{old}$  and the *q* distribution at the previous iteration which makes the efficient two-step scheme invalid for TCG. In practice,  $\hat{\sigma}_i^2$  is fixed as 0.1 to avoid overfitting during the EM iterations and only  $\mu$  is taken as video imprint which has the same size with  $\pi$ . The efficient two-step scheme is summarized in Algorithm 1.

## Algorithm 1. The Efficient Two-Step Scheme

**Input:** the feature maps  $\{\mathbf{f}^t\}_{t \in \mathbf{T}}$ , epitome size **E**, window size **W Output:**  $\hat{q}$ ,  $\mu$ 

- 1: for each  $t = 1 \dots T$  do
- 2:  $\mathbf{G}^t = \phi_{PCA}(\mathbf{F}^t)$
- 3: end for
- 4: repeat
- 5: Update  $\hat{q}$  with the EM steps of Equation (7)
- 6: **until** Convergence
- 7: Compute  $\mu$  with Equation (8)
- 8: Return  $\hat{q}$ ,  $\mu$

<sup>3.</sup> For ease of illustration of the video imprint, we accumulate the frames on the location with the maximum posterior probability  $q(l^t = \mathbf{k})$  and draw the mean image.



Fig. 4. Illustration of the reasoning network for event recognition and event recounting.

### 3.3 Feature Aggregation

In this section, we demonstrate how to aggregate the video imprint into a compact video representation for unsupervised event retrieval. We refer each  $\pi_i$  or  $\mu_i$  on the video imprint as an imprint descriptor. As shown in Fig. 3, some imprint descriptors are meaningless since no frames are aligned to their locations. The first step is to generate an active map for the video imprint to eliminate these meaningless imprint descriptors. Formally, the binary active map,  $\mathbf{A} = \{a_i | i \in \mathbf{E}\}, a_i \in \{0, 1\}$ , is computed as

$$a_{\mathbf{i}} = \begin{cases} 1 & \left\{ \mathbf{i} \in \mathbf{W}_{\mathbf{k}} | \mathbf{k} : \sum_{t=1}^{T} q(l_t = \mathbf{k}) > \tau \right\}, \quad (9) \\ 0 & \text{otherwise} \end{cases}$$

where  $\tau$  is the threshold of the active map.

After the active map is obtained, a direct sum-aggregation over the entire activated imprint descriptors is carried out to produce the final video representation. Formally, the aggregation step can be written as

$$TCG: \ \phi_{FA}(\boldsymbol{\pi}, \mathbf{A}) = \sum_{\mathbf{i} \in \mathbf{E}} a_{\mathbf{i}} \boldsymbol{\pi}_{\mathbf{i}},$$
  
Epitome: 
$$\phi_{FA}(\boldsymbol{\mu}, \mathbf{A}) = \sum_{\mathbf{i} \in \mathbf{E}} a_{\mathbf{i}} \boldsymbol{\mu}_{\mathbf{i}}.$$
 (10)

The obtained  $\phi_{FA}$  is subsequently  $l_2$ -normalized and the cosine similarity is computed for event retrieval.



Fig. 5. The average running time of learning epitomes for EVVE dataset with different d and the corresponding retrieval performance (mAP).

### 3.4 Reasoning Over the Imprint

Once the imprint is obtained for each video, it makes reasoning based on this compact imprint representation possible, where each location corresponds to a recurring scene/object part, with spatial layouts of imprint locations reflecting the scene/object spatial layouts in video frames where they were observed. We treat locations in the imprint in a similar way to the sentences in memory networks [43], [53]. Our reasoning network determines the event category in stages, which traverses attention from one set of imprint locations to the next. In this process as illustrated in Fig. 4, the imprint locations of large importance are highlighted, and we also trace these highlights back to the locations in video frames using the  $q/\hat{q}$  distributions as discussed above.

Our reasoning network differs from the memory networks in two ways. First, since there is no query question for event recognition, we initialize the input vector  $u^1$  with Equation (10), i.e., sum-aggregation of video imprint. Second, because the spatial organization in the imprint is meaningful, we incorporate an average spatial pooling layer after the softmax layer to improve the smoothness of the recounting results. The model details are as follows.

*Memory Layers in the Reasoning Network.* As shown in Fig. 4, the video imprint (non-activated locations are ignored) is processed via multiple memory layers (hops). In each layer, the imprint descriptors  $\pi_i$  or  $\mu_i$  from video imprint are first embedded to the output vector space and memory vector space with embedding matrices **B** and **M**, respectively

$$TCG: b_{i} = B\pi_{i}, m_{i} = M\pi_{i},$$
  
Epitome:  $b_{i} = B\mu_{i}, m_{i} = M\mu_{i},$  (11)

where  $b_i$  denotes the output vector and  $m_i$  denotes memory vector. The memory vector  $m_i$  is introduced to compute the weights map  $\mathbf{P} = \{p_i | i \in \mathbf{E}\}$  with the internal state u

$$p_{\mathbf{i}} = \operatorname{avgpooling}(\operatorname{softmax}(\boldsymbol{u}^{\top}\boldsymbol{m}_{\mathbf{i}})).$$
 (12)

The average pooling is performed with  $3 \times 3$  windows, stride 1. The output vector *o* is then computed by a weighted sum over the output vectors  $b_{ir}$  i.e.,

$$\boldsymbol{o} = \sum_{\mathbf{i}} p_{\mathbf{i}} \boldsymbol{b}_{\mathbf{i}}.\tag{13}$$

For the internal state vector u, the initial  $u^1$  is obtained with Equation (10), and the  $u^{k+1}$  in k + 1 layer can be computed by

$$\boldsymbol{u}^{k+1} = \boldsymbol{u}^k + \boldsymbol{o}^k. \tag{14}$$

The obtained output vector is fed into a decision network for event categorization. The decision network can consist of only a single softmax layer or multiple fully connected layers. The recounting heat map<sup>4</sup> (posterior probabilities  $q(l^t = \mathbf{i})$ ) of each frame shown in Fig. 4 is generated via the sum of all weights maps,  $\mathbf{P}^{\text{sum}} = \sum_k \mathbf{P}^k$ . We use  $\mathbf{P}_{\mathbf{W}_i}^{\text{sum}}$  to denote the weights map cropped from  $\mathbf{P}^{\text{sum}}$  in the window  $\mathbf{W}_i$ . Then the recounting map  $\mathbf{R}^t$  of frame *t* is

$$\mathbf{R}^{t} = \sum_{\mathbf{i} \in \mathbf{E}} q(l^{t} = \mathbf{i}) \mathbf{P}_{\mathbf{W}_{\mathbf{i}}}^{\text{sum}}.$$
 (15)

The importance score of each frame is obtained with the sum of the recounting map.

## 4 EXPERIMENTS

#### 4.1 Datasets and Evaluation Protocol

In terms of event retrieval, we validated our method on the large-scale benchmark EVVE dataset [38]. It contains 2,995 videos (620 videos are set as queries) related to 13 specific event classes. Given a single video of an event, the task is to retrieve videos related to the same event from the dataset. The methods are evaluated based on the mean AP (mAP) computed per event. The overall performance is evaluated by averaging the mAPs over the 13 events. In addition, a large distractor dataset (100,000 vedios) is also provided to evaluate the retrieval performance on large-scale data.

To evaluate the event recognition and recounting, we used three datasets: EVVE, Columbia Consumer Videos (CCV) [24] and TRECVID MEDTest 14 (MED14) [33].

In addition, we also configured the EVVE as a small recognition dataset with 13 events. For each event, we set the query video as the test data (620 videos), and treat the ground truth in the dataset as the training data. We report the top-1 classification accuracy for performance evaluation.

The Columbia Consumer Videos dataset [24] contains 9,317 YouTube videos of 20 classes. We follow the protocol defined in [24], with a training set of 4,659 videos and a test set of 4,658 videos. The TRECVID MEDTest 14 (MED14) [33] is one of the most challenging datasets for event recognition containing 20 complex events. In the training section, there are 100 positive exemplars per event, and all events share negative exemplars with about 5,000 videos. The test section contains approximately 23,000 videos.

For these two datasets, mAP is used as the evaluation metric of event recognition according to the NIST standard [33]. Since no ground truth is available for the recounting task, we only provide a user study and qualitative analysis for such results.

#### 4.2 Implementation Details

*Frame-Level Descriptor*. Given an input video, we sample 5 frames per second (5 fps) to extract the CNN features. We

explore various pre-trained CNN models, including Alex-Net [27], VGG [42] and ResNet-50 [18]. We adopt the output from the last convolutional layer (after ReLU) of these models as the frame descriptors. The CNN feature maps are down-sampled to  $4 \times 4$  with linear interpolation to fit the TCG (we set  $\mathbf{S} = 4 \times 4$  in TCG for computational efficiency). For the epitome model, the feature maps are sampled to  $8 \times 8$  to fit the window size  $\mathbf{W}$ . In addition, we also average all the frame descriptors over the video (sum-aggregation), as the baseline to evaluate our framework.

*Post-Processing.* For the baseline video representation, we apply the same post-processing strategy as in [1], [16], i.e., the representation vector of a video is first  $l_2$ -normalized, and then whitened using PCA [19] and  $l_2$ -normalized again. For the imprint descriptors on the video imprint, power normalization ( $\alpha = 0.2$ ) shows better results than  $l_2$ -normalization in our experiments. Therefore, after feature alignment, the imprint descriptors are first power normalized, then PCA-whitened and  $l_2$ -normalized.

*Re-Ranking Methods for Event Retrieval.* For the event retrieval task on the EVVE dataset, we also employ two variants of query expansion method presented by Douze et al. [12]: Average Query Expansion (AQE) and Difference of Neighborhood (DoN). In our experiments, we set  $N_1 = 10$  for AQE and  $N_1 = 10$ ,  $N_2 = 2000$  for DoN.

Training Details for the Reasoning Network. The reasoning network (RNet) is trained with stochastic gradient descent (SGD). The initial learning rate is  $\beta = 0.025$ , which is then annealed every 5 epochs by  $\beta/2$  until 20 epochs are finished. All weights are initialized randomly from a Gaussian distribution with zero mean and  $\sigma = 0.05$ . The weights are shared among different memory layers. The batch size is 128 and the gradients with an  $l_2$  norm larger than 20 are rescaled to norm 20 during the training step.

### 4.3 Complexity Analysis

The most time consuming step is constructing video imprint for input videos. As shown in Fig. 5, the average running time of TCG (with ResNet features) for EVVE (about 1,200 frames per video) implemented on the GPU platform (K40 with MATLAB parallel computing toolbox) is about 18 seconds. As discussed in [34], [35], with efficient use of cumulative sums, the computational complexity of learning TCG with the EM algorithm grows at most linearly with the product of counting grid size and the tessellation sections  $E_x \cdot E_y \cdot S_x \cdot S_y$ .

For the epitome, the computational complexity without the efficient two-step scheme for *D*-dimensional input feature maps and EM iteration times n ( $n \approx 30$ ) is O(nD). With the proposed efficient two-step scheme, the computational complexity becomes O(nd + D). When  $d \ll D$ , the efficient two-step scheme can significantly accelerate the learning stage. Fig. 5 shows the average running time of learning epitomes for the EVVE dataset with different *d* and the corresponding retrieval performance (mAP). When d = 64, without any performance losses, the efficient two-step scheme achieves almost an order of magnitude speed-up.

## 4.4 Evaluation Results on Event Retrieval

#### 4.4.1 Parameter Analysis

*Threshold*  $\tau$  *of the Active Map.* Fig. 6a shows the retrieval performance with different threshold  $\tau$  used for the active map

<sup>4.</sup> More sophisticated recounting inferences can be implemented by computing conditional heat maps based on individual memory layers as we trace the reasoning engine through the layers.



Fig. 6. The influence of different (a) threshold  $\tau$  of active map and (b) grid sizes **E** for Counting grid aggregation (CGA) and Epitome aggregation (EPA). Alex and res denote two CNN models, i.e., AlexNet and ResNet.

construction. We can observe that increasing  $\tau$  helps filter out some very short shots (with small number of frames) which are usually irrelevant. We set  $\tau = 4$  in the subsequent experiments.

*Grid Size* **E**. To evaluate the influence of the grid size **E** of the TCG and epitome models, we first fix the window size (**W** = 8 × 8) of TCG, epitome, and tessellation size (**S** = 4 × 4) of TCG. Subsequently, we choose 7 different counting grid sizes to perform the feature alignment with the epitome and TCG model, respectively. The performance with regard to size choices is presented in Fig. 6b. No further improvement can be obtained when  $E_x = E_y > 24$ . Therefore, the size of the grid is fixed at 24 for all following experiments.

#### 4.4.2 Comparison with Sum-Aggregation

We refer to our unsupervised flow (combining the feature alignment and aggregation steps) on ER3 as counting grid aggregation (CGA) for TCG based video imprint and epitome aggregation (EPA) for epitome based video imprint. Table 2 shows the retrieval performance compared with baselines. We evaluate the CGA and EPA on two different CNN models, AlexNet [27] and ResNet-50 [18]. Fig. 7 shows the retrieval performance for each event class of EVVE dataset. The IDs of events are the same with [38]. We can see that the video representation based on ResNet-50 demonstrates superior performance in most event categories compared with AlexNet. And the proposed aggregation methods can further boost the retrieval performance.

TABLE 2 Comparison with Sum-Aggregation on EVVE Dataset

Representation	Dim.	mAP
Sum-alex	256	38.3
Sum-res	1,024	46.6
Sum-(alex+res)	1,280	47.3
CGA-alex	256	42.6
CGA-res	1,024	51.2
CGA-(alex+res)	1,280	52.3
EPA-alex	256	41.5
EPA-res	1,024	51.6
EPA-(alex+res)	1,280	52.1

Sum-, CGA- and EPA- denote sum-aggregation, counting grid aggregation, epitome aggregation, respectively. alex and res denote two CNN models, AlexNet and ResNet-50. For ResNet based representation, the vectors dimension are reduced to 1024 with PCA-whitening. (alex+res) denotes the concatenated vector.

As shown in Table 2, compared with Sum-aggregation, CGA and EPA both obtain better retrieval performance with the benefits from feature alignment step that can suppress the redundancy among frames. In addition, consistent improvement can be observed for different CNN models (AlexNet and ResNet-50). After concatenating the video representation vectors from these two CNN models, the retrieval performance can be further improved, mAP = 52.3 for CGA and mAP = 52.1 for EPA.

#### 4.4.3 Comparison with State-of-the-Arts

In Table 3, we can see that the sum-aggregation with CNN features already outperforms previous work [12], [37], [38]. After merging with 100K distractors, the mAPs of CGA-(alex +res) and EPA-(alex+res) achieve 42.9 and 43.9 respectively, which are still better than the baseline (mAP = 38.7) and previous work [12], [37], [38]. For fair comparison, we also reimplement previous aggregation methods (CTE, TMK and SHP) based on the same CNN features with CGA and EPA. Note that MMV is in fact the sum-aggregation based on the multi-VLAD frame feature [38], which employs the same aggregation method with our baseline method.

As shown in Table 3, frame-level descriptors have a huge impact on performance. CNN feature-based aggregation methods achieve better retrieval performance than their original handcrafted feature-based implementations. However, based on the same frame-level descriptors, our aggregation methods still obtain better results compared with



Fig. 7. Retrieval performance (mAP) per event.

TABLE 3 Retrieval Performance Compared with Other Methods

		EVVE		EVVE+100K			
Method	Dim.	no QE	AQE	DoN	no QE	AQE	DoN
MMV [38]	512	33.4	_	_	22.0	_	_
CTE [38]	-	35.2	_	_	20.2	_	-
MMV+CTE [38]	-	37.6	-	-	25.4	-	-
TMK [37]	66,560	33.5	36.1	41.3	25.4	-	34.7
SHP [12]	16,384	36.3	38.9	44.0	26.5	30.1	33.1
CTE-(alex+res)	_	44.7	_	_	40.1	_	-
TMK-(alex+res)	42,240	46.7	51.8	53.0	40.9	48.4	48.0
SHP-(alex+res)	40,960	48.7	54.0	55.4	41.5	48.7	48.9
Sum-(alex+res)	1,280	47.3	53.1	55.2	38.7	45.8	47.1
CGA-(alex+res)	1,280	52.3	58.5	60.1	42.9	50.4	52.7
EPA-(alex+res)	1,280	52.1	59.3	59.9	43.9	53.2	53.8

AQE and DoN denote the two Re-ranking methods.

existing methods. In addition, previous aggregation methods usually lead to much higher dimension in the video representation (which may reduce the efficiency of the retrieval task). Based on the initial retrieval results, the query expansion can further boost the performance. We achieve 8.5 and 8.9 percent improvement compared with previous result (mAP = 55.4 for SHP [12] combined with CNN feature) and the baseline (mAP = 55.2) on the EVVE dataset, respectively. Consistent improvement is also observed with query expansion on the large dataset (EVVE+100K).

## 4.5 Evaluation Results on Event Recognition

## 4.5.1 Parameter Analysis

*Structure of the Reasoning Network.* For the EVVE dataset, we set a softmax layer as the decision network. The video imprint is generated based on the ResNet-50 [18] model and

TABLE 4 Comparison with Sum-Aggregation and CGA/EPA

	Method	vgg	res	(vgg+res)
CCV	Sum-	74.3	75.3	78.1
	CGA-	75.7	76.6	79.1
	EPA-	75.2	76.7	78.3
	CG-RNet-	76.7	78.5	<b>79.9</b>
	EP-RNet-	75.8	78.2	79.0
MED14	Sum-	26.0	30.4	32.8
	CGA-	30.5	32.2	33.7
	EPA-	31.2	32.4	34.4
	CG-RNet-	32.8	34.2	<b>36.9</b>
	EP-RNet-	32.6	34.3	36.3

Sum-, CGA- and EPA- denote sum-aggregation, counting grid aggregation and epitome aggregation, respectively. RNet- denote the reasoning network. vgg and res denote two CNN model, VGG and ResNet-50. (vgg+res) denotes the later fusion result.

its imprint descriptors are first reduced to 256 dimension with PCA-whitening before being fed to the reasoning network (RNet). For the CCV and MED14 datasets, we add a fully connected layer in front of the softmax layer as the decision network for better performance. Besides the ResNet-50 model, we also evaluate the framework with VGG (16 layers) [42] model on these two datasets. The dimension of the imprint descriptors is set to 1024 and 512 for ResNet-50 and VGG, respectively. For all the datasets, the internal vectors  $b_i$  and  $m_i$  have the same dimensions with the input imprint descriptors.

*Number of Memory Layers.* Fig. 8 illustrates the influence of RNet with increased hops on the EVVE dataset. To make a fair comparison, we employ the same decision network as classifier for the baselines and the output from RNet. We



Fig. 8. The influence of RNet with increased hops on EVVE dataset. Best viewed in color.

	Method	mAP	Recounting
	Lai et al. [28]	43.6	$\checkmark$
	Jiang et al. [24]	59.5	×
	Wu et al. [54]	70.6	×
CCV	Nagel et al. [32]	71.7	×
	Wu et al. [55]	84.9	×
	CG-RNet-(vgg+res)	79.9	$\checkmark$
	EP-RNet-(vgg+res)	79.0	$\checkmark$
	MA+CG-RNet-(vgg+res)	87.1	$\checkmark$
	MA+EP-RNet-(vgg+res)	86.8	$\checkmark$
MED14	IDT [33], [52]	27.6	×
	Gan et al. [13]	33.3	$\checkmark$
	Xu et al. [57]	36.8	×
	Zha et al. [60]	38.7	×
	CG-RNet-(vgg+res)	36.9	$\checkmark$
	EP-RNet-(vgg+res)	36.3	$\checkmark$
	IDT+CG-RNet-(vgg+res)	40.2	$\checkmark$
	IDT+EP-RNet-(vgg+res)	39.8	$\checkmark$

TABLE 5
Comparison with Other Methods

MA+CG/EP-RNet-(vgg+res) denote the results fused with audio and motion
information using adaptive fusion method [55]. IDT+CG/EP-RNet-(vgg+res)
denote the results fused with improved dense trajectories [52].



Fig. 9. Influence of the average pooling layer in RNet. The middle column shows the recounting map of RNet. The right column shows the recounting map with avg-pooling layer removed. Best viewed in color.

compare with two representations, one is the video representation with sum-aggregation, and the other is either the counting grid aggregation (CGA) or the epitome aggregation (EPA) which depends on the feature alignment step. In fact, if we fix the value of the weights map equal to the active map, the RNet will reduce to the CGA or EPA, i.e., the unsupervised flow in Fig. 1. We can see that CGA/EPA



Fig. 10. The statistical results of user study. The average percentage of key frames belong to *Bad, OK* and *Good* are shown with different color bar. Best viewed in color.



Fig. 11. Examples of the key frames for user study. Each blue box contains the key frames proposed from a test video with three strategies, i.e., random selection, recounting results based on TCG or epitome. Each strategy proposes four key frames to represent the video. the scores ranked by the users are shown with different color boxes, red for *Bad*, yellow for *OK* and green for *Good*. Best viewed in color.

provides better performance than sum-aggregation, and the RNet can further refine the video representation and lead to better recognition accuracy than the two baselines. In addition, the gain is also increased with more hops. Consistent gains are observed on both CCV and MED14 datasets. We set the hops = 3 in the following experiments on the CCV and MED14 datasets.



Fig. 12. Examples of event recounting results. We use heat map to indicate the recounting map. The key areas related to the event in each frame is painted with red color. The importance score which is computed by the sum of recounting map is shown with color bar (red for important frames) upon the video frame flow. Best viewed in color.

#### 4.5.2 Performance on CCV and MED14

Table 4 shows the recognition performance (mAP) of RNet and baseline methods. With the benefit from re-weighting the video imprint, the RNet achieves better results on both CCV (mAP = 79.9/79.0) and MED14 (mAP = 36.9/36.3) datasets compared with sum-aggregation and CGA/EPA. In addition, on the CCV dataset, we also employ the same strategy as [55] to combine motion and audio features with our appearancebased representation. As shown in Table 5, the fusion strategy further boosts the recognition performance (mAP = 87.1/86.8). On MED14 dataset, the proposed method sets a new performance record (mAP = 40.2) after fusing motion features (IDT) in a similar way to [60], with the additional convenience of simultaneously providing recounting results.

### 4.6 Evaluation Results on Event Recounting

*Influence of Average Pooling*. In contrast to the original memory networks [43], we add an average pooling layer inside the memory layer, which takes advantage of the spatial organization of the information in the video imprint. Fig. 9 demonstrates the influence of adding the average pooling layer. We can see that the recounting maps are smoother and more reasonable, especially for TCG based video imprint. In addition, with benefits from finer resolution of input feature maps, the epitome based video imprint can hold more spatial information, and the recounting results are more reasonable than TCG based video imprint.

*User Study.* The evaluation of event recounting is not easy since there is no ground truth information. To assess the quality of our recounting results, we randomly sample 200 videos from 620 test videos of the EVVE dataset for user study. First, for each test video, we sample 4 key frames either randomly or based on the importance score from the recounting map R (3 groups in total: randomly selected, based on important scores from TCG or Epitome). Then, we invited 50 users to score the key frames with {Bad, OK, *Good* } based on the relevance with the labeled event. Here, Bad means that the selected frame is irrelevant with its event label, i.e., users cannot recognize the event category by the selected frame. OK means that users can identify the event category by the selected frame with reasonable amount of guesswork. Good means that the event category can be identified beyond doubt by the selected frame and users are highly confident with the judgment.

The statistical results are shown in Fig. 10. It presents the average percentage of key frames belong to *Bad*, *OK* or *Good*, respectively. Compared with random selection, the key frames proposed with recounting map are more likely to be relevant with the corresponding event, the percentage of bad proposals is reduced from 39 to 26 percent. Note that epitome based recounting results are slightly better than TCG due to richer spatial information. Fig. 11 shows the examples of key frames for user study. Each blue box contains a test video represented by 4 key frames. The key frames are proposed with three strategies, i.e., random

selection, recounting based on TCG or epitome. The scores ranked by the users are shown with different color boxes.

Recounting Map. Fig. 12 illustrates some examples of the recounting results. The heat map is used to visualize the recounting map (the map is rescaled to the same size with the frame). Since no ground truth recounting is available, we can only provide some examples as shown in Fig. 12. We can see that our recounting process can not only provide the importance score of each frame, but also indicate the most relevant areas inside each frame. However, due to the coarse resolution of the input feature maps, the spatial-level recounting results are also very coarse. Nevertheless, the recounting heat map may be used as a good prior for other post-processing methods, e.g., object segmentation.

#### 5 CONCLUSION

In this paper, we propose an efficient and effective event retrieval, recognition and recounting framework (ER3), based on our proposed video imprint representation. In the proposed framework, a dedicated feature alignment module is incorporated for redundancy removal across frames to produce the compact intermediate tensor representation, i.e., the video imprint. Subsequently, the video imprint is processed by a reasoning network for event recognition/ recounting, and by a feature aggregation module for event retrieval. Thanks to the attention mechanism inspired by the memory networks in language modeling, the proposed reasoning network is capable of simultaneous event recognition and event recounting. With the event retrieval task, the compact video representation aggregated from the video imprint contributes to better retrieval results.

#### ACKNOWLEDGMENTS

This work was supported partly by National Key R&D Program of China Grant 2017YFA0700800, National Natural Science Foundation of China Grants 61629301, 61773312, 91748208, and 61503296, China Postdoctoral Science Foundation Grants 2017T100752 and 2015M572563.

## REFERENCES

- A. Babenko and V. Lempitsky, "Aggregating local deep features [1] for image retrieval," in Proc. IEEE Int. Conf. Comput. Vis., 2015, pp. 1269–1277
- M. Baillie and J. M. Jose, "Audio-based event detection for sports [2] video," in Proc. Int. Conf. Image Video Retrieval, 2003, pp. 300–309. L. Baraldi, M. Douze, R. Cucchiara, and H. Jégou, "LAMV: Learn-
- [3] ing to align and match videos with kernelized temporal layers," in
- Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2018, pp. 7804–7813. S. Bhattacharya, M. M. Kalayeh, R. Sukthankar, and M. Shah, [4] "Recognition of complex events: Exploiting temporal dynamics between underlying concepts," in Proc. IEEE Conf. Comput. Vis.
- Pattern Recognit., 2014, pp. 2235–2242.
  [5] H. Bilen, B. Fernando, E. Gavves, and A. Vedaldi, "Action recognition with dynamic image networks," *IEEE Trans. Pattern Anal.* Mach. Intell., 2017. doi: 10.1109/TPAMI.2017.2769085
- [6] F. Caba Heilbron, V. Escorcia, B. Ghanem, and J. Carlos Niebles, "ActivityNet: A large-scale video benchmark for human activity understanding," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2015, pp. 961–970.
- L. Cao, Y. Mu, A. Natsev, S.-F. Chang, G. Hua, and J. R. Smith, [7] "Scene aligned pooling for complex video recognition," in Proc. Eur. Conf. Comput. Vis., 2012, pp. 688-701.
- C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector [8] machines," ACM Trans. Intell. Syst. Technol., vol. 2, no. 3, 2011, Art. no. 27.

- [9] X. Chang, Y.-L. Yu, Y. Yang, and A. G. Hauptmann, "Searching persuasively: Joint event detection and evidence recounting with limited supervision," in Proc. ACM Int. Conf. Multimedia, 2015, pp. 581–590. [10] X. Chang, Y.-L. Yu, Y. Yang, and E. P. Xing, "Semantic pooling for
- complex event analysis in untrimmed videos," IEEE Trans. Pattern Anal. Mach. Intell., vol. 39, no. 8, pp. 1617–1632, Aug. 2017. [11] N. Dalal and B. Triggs, "Histograms of oriented gradients for
- human detection," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2005, pp. 886-893.
- [12] M. Douze, J. Revaud, C. Schmid, and H. Jégou, "Stable hyper-pooling and query expansion for event detection," in *Proc. IEEE* Int. Conf. Comput. Vis., 2013, pp. 1825–1832. C. Gan, N. Wang, Y. Yang, D.-Y. Yeung, and A. G. Hauptmann,
- [13] "DevNet: A deep event network for multimedia event detection and evidence recounting," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2015, pp. 2568-2577.
- [14] C. Gan, T. Yao, K. Yang, Y. Yang, and T. Mei, "You lead, we exceed: Labor-free video concept learning by jointly exploiting web videos and images," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2016, pp. 923-932.
- [15] Z. Gao, G. Hua, D. Zhang, N. Jojic, L. Wang, J. Xue, and N. Zheng, "ER3: A unified framework for event retrieval, recognition and recounting," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2017, pp. 2253–2262.
- [16] Z. Gao, J. Xue, W. Zhou, S. Pang, and Q. Tian, "Democratic diffu-sion aggregation for image retrieval," *IEEE Trans. Multimedia*, vol. 18, no. 8, pp. 1661-1674, Aug. 2016.
- [17] Y. Gong, L. Wang, R. Guo, and S. Lazebnik, "Multi-scale orderless pooling of deep convolutional activation features," in Proc. Eur. Conf. Comput. Vis., 2014, pp. 392-407.
- [18] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2016, pp. 770-778.
- [19] H. Jégou and O. Chum, "Negative evidences and co-occurences in image retrieval: The benefit of PCA and whitening," in Proc. Eur. Conf. Comput. Vis., 2012, pp. 774-787.
- H. Jégou, M. Douze, C. Schmid, and P. Pérez, "Aggregating local descriptors into a compact image representation," in *Proc. IEEE* [20] Conf. Comput. Vis. Pattern Recognit., 2010, pp. 3304-3311.
- [21] H. Jegou, F. Perronnin, M. Douze, J. Sánchez, P. Perez, and C. Schmid, "Aggregating local image descriptors into compact codes," IEEE Trans. Pattern Anal. Mach. Intell., vol. 34, no. 9, pp. 1704–1716, Sep. 2012.
- [22] Y.-G. Jiang, S. Bhattacharya, S.-F. Chang, and M. Shah, "Highlevel event recognition in unconstrained videos," Int. J. Multimedia *Inf. Retrieval*, vol. 2, no. 2, pp. 73–101, 2013. [23] Y.-G. Jiang, Z. Wu, J. Wang, X. Xue, and S.-F. Chang, "Exploiting
- feature and class relationships in video categorization with regu-larized deep neural networks," IEEE Trans. Pattern Anal. Mach. Intell., vol. 40, no. 2, pp. 352-364, Feb. 2018.
- [24] Y.-G. Jiang, G. Ye, S.-F. Chang, D. Ellis, and A. C. Loui, "Consumer video understanding: A benchmark database and an evaluation of human and machine performance," in Proc. Int. Conf. Multimedia Retrieval, 2011, pp. 29–37.
- [25] N. Jojic, B. J. Frey, and A. Kannan, "Epitomic analysis of appear-ance and shape," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2003, pp. 34–41. [26] N. Jojic and A. Perina, "Multidimensional counting grids: Infer-
- ring word order from disordered bags of words," in Proc. Conf. Uncertainty Artif. Intell., 2011, pp. 547–556.
- [27] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in Proc. Conf. Neural Inf. Process. Syst., 2012, pp. 1097-1105.
- [28] K.-T. Lai, X. Y. Felix, M.-S. Chen, and S.-F. Chang, "Video event detection by inferring temporal instance labels," in *Proc. IEEE* Conf. Comput. Vis. Pattern Recognit., 2014, pp. 2251-2258.
- [29] K.-T. Lai, D. Liu, M.-S. Chen, and S.-F. Chang, "Recognizing complex events in videos by learning key static-dynamic evidences," in Proc. Eur. Conf. Comput. Vis., 2014, pp. 675-688.
- [30] D. G. Lowe, "Distinctive image features from scale-invariant key-
- points," Int. J. Comput. Vis., vol. 60, no. 2, pp. 91–110, 2004. Z. Ma, Y. Yang, Z. Xu, S. Yan, N. Sebe, and A. G. Hauptmann, [31] "Complex event detection via multi-source video attributes," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2013, pp. 2627–2633.
- [32] M. Nagel, T. Mensink, C. G. Snoek et al. "Event fisher vectors: Robust encoding visual diversity of visual streams," in Proc. Brit. Mach. Vis. Conf., 2015, pp. 6-18.

- [33] P. Over, J. Fiscus, G. Sanders, D. Joy, M. Michel, G. Awad, A. Smeaton, W. Kraaij, and G. Quénot, "TRECVID 2014–an overview of the goals, tasks, data, evaluation mechanisms and metrics," in *Proc. Text Retrieval Conf. Video Retrieval Eval.*, 2014, pp. 52–74.
- [34] A. Perina and N. Jojic, "Image analysis by counting on a grid," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit., 2011, pp. 1985– 1992.
- [35] A. Perina and N. Jojic, "Capturing spatial interdependence in image features: The counting grid, an epitomic representation for bags of features," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 37, no. 12, pp. 2374–2387, Dec. 2015.
- [36] F. Perronnin and D. Larlus, "Fisher vectors meet neural networks: A hybrid classification architecture," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 3743–3752.
- [37] S. Poullot, S. Tsukatani, A. Phuong Nguyen, H. Jégou, and S. Satoh, "Temporal matching kernel with explicit feature maps," in *Proc. ACM Int. Conf. Multimedia*, 2015, pp. 381–390.
- [38] J. Revaud, M. Douze, C. Schmid, and H. Jégou, "Event retrieval in large video collections with circulant temporal encoding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2013, pp. 2459–2466.
- [39] J. Sánchez, F. Perronnin, T. Mensink, and J. Verbeek, "Image classification with the fisher vector: Theory and practice," Int. J. Comput. Vis., vol. 105, no. 3, pp. 222–245, 2013.
- [40] A. Sharif Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "CNN features off-the-shelf: An astounding baseline for recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. Workshops, 2014, pp. 806–813.
- [41] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Proc. Conf. Neural Inf. Process. Syst.*, 2014, pp. 568–576.
- [42] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," arXiv:1409.1556, 2014.
- [43] S. Sukhbaatar, J. Weston, R. Fergus et al. "End-to-end memory networks," in *Proc. Conf. Neural Inf. Process. Syst.*, 2015, pp. 2440– 2448.
- [44] C. Sun and R. Nevatia, "ACTIVE: Activity concept transitions in video event classification," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2013, pp. 913–920.
- [45] C. Sun and R. Nevatia, "DISCOVER: Discovering important segments for classification of video events and recounting," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2014, pp. 2569–2576.
- [46] R. Szeliski, "Image alignment and stitching: A tutorial," Found. Trends<sup>®</sup> Comput. Graph. Vis., vol. 2, no. 1, pp. 1–104, 2006.
- [47] R. Szeliski and H.-Y. Shum, "Creating full view panoramic image mosaics and environment maps," in *Proc. Conf. Comput. Graph. Interactive Techn.*, 1997, pp. 251–258.
- [48] L. Torresani, M. Szummer, and A. Fitzgibbon, "Efficient object category recognition using classemes," in *Proc. Eur. Conf. Comput. Vis.*, 2010, pp. 776–789.
- [49] D. Tran, J. Yuan, and D. Forsyth, "Video event detection: From subvolume localization to spatiotemporal path search," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 36, no. 2, pp. 404–416, Feb. 2014.
- [50] C.-Y. Tsai, M. L. Alexander, N. Okwara, and J. R. Kender, "Highly efficient multimedia event recounting from user semantic preferences," in *Proc. Int. Conf. Multimedia Retrieval*, 2014, pp. 419–427.
- [51] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu, "Action recognition by dense trajectories," in *Proc. IEEE Conf. Comput. Vis. Pattern Rec*ognit., 2011, pp. 3169–3176.
- [52] H. Wang and C. Schmid, "Action recognition with improved trajectories," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2013, pp. 3551– 3558.
- [53] J. Weston, A. Bordes, S. Chopra, A. M. Rush, B. van Merriënboer, A. Joulin, and T. Mikolov, "Towards AI-complete question answering: A set of prerequisite toy tasks," in *Proc. Int. Conf. Learn. Representations*, 2016.
- [54] Z. Wu, Y.-G. Jiang, J. Wang, J. Pu, and X. Xue, "Exploring interfeature and inter-class relationships with deep neural networks for video classification," in *Proc. ACM Int. Conf. Multimedia*, 2014, pp. 167–176.
- [55] Z. Wu, Y.-G. Jiang, X. Wang, H. Ye, and X. Xue, "Multi-stream multi-class fusion of deep networks for video classification," in *Proc. ACM Int. Conf. Multimedia*, 2016, pp. 791–800.
- [56] Z. Wu, X. Wang, Y.-G. Jiang, H. Ye, and X. Xue, "Modeling spatialtemporal clues in a hybrid deep learning framework for video classification," in *Proc. ACM Int. Conf. Multimedia*, 2015, pp. 461–470.

- [57] Z. Xu, Y. Yang, and A. G. Hauptmann, "A discriminative CNN video representation for event detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1798–1807.
  [58] J. Yang, P. Ren, D. Zhang, D. Chen, F. Wen, H. Li, and G. Hua,
- [58] J. Yang, P. Ren, D. Zhang, D. Chen, F. Wen, H. Li, and G. Hua, "Neural aggregation networks for video face recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 2492–2495.
- [59] J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici, "Beyond short snippets: Deep networks for video classification," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 4694–4702.
- [60] S. Zha, F. Luisier, W. Andrews, N. Srivastava, and R. Salakhutdinov, "Exploiting image-trained CNN architectures for unconstrained video classification," in *Proc. Brit. Mach. Vis. Conf.*, 2015, pp. 60–73.
- [61] P. Zhang, C. Lan, J. Xing, W. Zeng, J. Xue, and N. Zheng, "View adaptive recurrent neural networks for high performance human action recognition from skeleton data," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2117–2126.
- [62] Q. Zhang and G. Hua, "Multi-view visual recognition of imperfect testing data," in Proc. ACM Int. Conf. Multimedia, 2015, pp. 561– 570.



Zhanning Gao received the BS degree in automatic control engineering from Xi'an Jiaotong University, Xi'an, China, in 2012. He is currently working toward the PhD degree in the Institute of Artificial Intelligence and Robtics, Xi'an Jiaotong University. He was a research intern in Visual Computing Group, Microsoft Research Asia from 2015 to 2017. His research interests include compact image/video representation, large scale content based multimedia retrieval, and complex event video analysis.



Le Wang (M'14) received the BS and PhD degrees in control science and engineering from Xi'an Jiaotong University, in 2008 and 2014, respectively. From 2013 to 2014, he was a visiting PhD student with Stevens Institute of Technology. From 2016 to 2017, he is a visiting scholar with Northwestern University. He is currently an associate professor with the Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University. His research interests include computer vision, machine learning, and their application for web

images and videos. He is the author of more than 20 peer reviewed publications in prestigious international journals and conferences. He is a member of the IEEE.



**Nebojsa Jojic** received the PhD degree from the University of Illinois at Urbana-Champaign, in 2001, where he received a Microsoft Fellowship, in 1999 and a Robert T. Chien Excellence in Research award, in 2000. He has been a researcher with Microsoft Research in Redmond, Washington, since 2000. He has published more than 100 papers in the areas of computer vision, machine learning, signal processing, computer graphics, and computational biology.



Zhenxing Niu received the PhD degree in control science and engineering from Xidian University, Xi'an, China, in 2012. From 2013 to 2014, he was a visiting scholar with the University of Texas at San Antonio, Texas. He is a researcher with Alibaba Group, Hangzhou, China. Before joining Alibaba Group, he is an associate professor of the School of Electronic Engineering, Xidian University, Xi'an, China. His research interests include computer vision, machine learning, and their application in object discovery and localization.

He served as PC member of CVPR, ICCV, and ACM Multimedia. He is a member of the IEEE.

#### GAO ET AL.: VIDEO IMPRINT



Nanning Zheng (SM'94-F'06) received the graduate degree from the Department of Electrical Engineering, Xi'an Jiaotong University (XJTU), in 1975, the ME degree in information and control engineering from Xi'an Jiaotong University, in 1981, and the PhD degree in electrical engineering from Keio University, in 1985. He is currently a professor and the director with the Institute of Artificial Intelligence and Robotics of Xi'an Jiaotong University. His research interests include computer vision, pattern recognition, computa-

tional intelligence, and hardware implementation of intelligent systems. Since 2000, he has been the Chinese representative on the Governing Board of the International Association for Pattern Recognition. He became a member of the Chinese Academy Engineering, in 1999. He is a fellow of the IEEE.



Gang Hua (M'03-SM'11) was enrolled in the Special Class for the Gifted Young of Xi'an Jiaotong University (XJTU), in 1994, and received the BS degree in automatic control engineering from XJTU, in 1999, the MS degree in control science and engineering from XJTU, in 2002, and the PhD degree in electrical engineering and computer science from Northwestern University, in 2006. He is currently a principle researcher/research manager with Microsoft Research. Before that, he was an associate professor of computer science with the Stevens Institute of Tech-

nology. He also held an academic advisor position with IBM T. J. Watson Research Center between 2011 and 2014. He was a research staff member with IBM Research T. J. Watson Center from 2010 to 2011, a senior researcher with Nokia Research Center, Hollywood from 2009 to 2010, and a scientist with Microsoft Live Labs Research from 2006 to 2009. He is currently an associate editor in chief of the Computer Vision and Image Understanding, and associate editor of the International Journal of Computer Vision, the IEEE Transactions on Image Processing, the IEEE Transactions on Circuits and Systems for Video Technology, the IEEE Multimedia, and the Journal of Machine Vision and Applications. He also served as the lead quest editor on two special issues in the IEEE Transactions on Pattern Analysis and Machine Intelligence and the International Journal of Computer Vision, respectively. He is a program chair of CVPR'2019&2022. He is an area chair of CVPR'2015&2017, ICCV'2011&2017, ICIP'2012&2013&2016, ICASSP'2012&2013, and ACM MM 2011&2012&2015&2017. He is the author of more than 150 peer reviewed publications in prestigious international journals and conferences. He holds 19 issued US patents and has 20 more US patents pending. He is the recipient of the 2015 IAPR Young Biometrics Investigator Award for his contribution on Unconstrained Face Recognition from Images and Videos, and a recipient of the 2013 Google Research Faculty Award. He is an IAPR fellow, an ACM distinguished scientist, and a senior member of the IEEE.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.