

SUSiNet: See, Understand and Summarize it

Petros Koutras and Petros Maragos

School of E.C.E., National Technical University of Athens, Greece

{pkoutras, maragos}@cs.ntua.gr

Abstract

In this work we propose a multi-task spatio-temporal network, called SUSiNet, that can jointly tackle the spatio-temporal problems of saliency estimation, action recognition and video summarization. Our approach employs a single network that is jointly end-to-end trained for all tasks with multiple and diverse datasets related to the exploring tasks. The proposed network uses a unified architecture that includes global and task specific layer and produces multiple output types, i.e., saliency maps or classification labels, by employing the same video input. Moreover, one additional contribution is that the proposed network can be deeply supervised through an attention module that is related to human attention as it is expressed by eye-tracking data. From the extensive evaluation, on seven different datasets, we have observed that the multi-task network performs as well as the state-of-the-art single-task methods (or in some cases better), while it requires less computational budget than having one independent network per each task.

1. Introduction

During the last decade, the extensive usage of Convolutional Neural Networks (CNNs) has boosted the performance throughout the majority of spatial tasks in computer vision, such as object detection or semantic segmentation [52, 27, 26]. Nowadays, automatic video understanding becomes one of the most essential and demanding challenges and research directions due to the increased amount of video data (i.e. YouTube videos, movies, documentaries, home videos). The new problems that have arisen in this field, such as activity recognition, video saliency, scene analysis or video summarization, require the integration and modelling of the temporal evolution instead of applying image based algorithms to each frame independently, such as in video segmentation or pose estimation. In addition,

This research has been co-financed by the European Regional Development Fund of the European Union and Greek national funds through the Operational Program Competitiveness, Entrepreneurship and Innovation, under the call Research - Create - Innovate (project code: T1EDK-01248, i-Walk)

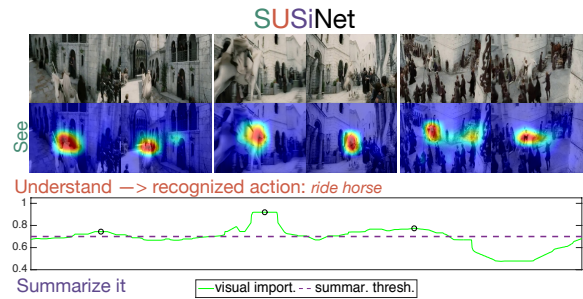


Figure 1. Example of the multi-task spatio-temporal SUSiNet that can jointly perform *saliency estimation*, *action recognition* and *video summarization* in a similar manner as a human who views the video. The second row depicts the saliency maps produced by our net, which highlight the most important part of the video. In parallel, our network recognizes the visual action “ride horse” and computes importance scores (green line) for each frame that indicate which video segments will be included in the summary.

as large-scale video datasets have appeared, an increased performance regarding the video related tasks, i.e., action recognition, can also be achieved [10, 25].

However, till recently, the majority of this great progress in computer vision has been achieved by facing each task independently, without trying to jointly solve multiple tasks or investigating relationships between them. Recent works towards this direction have proceeded by training multi-task networks [34] or finding the structure among visual tasks and apply transfer learning [65]. These approaches open the road for investigating simpler and faster ways to solve multiple task simultaneously, rather than maximize the task-specific performance, and thus providing useful computer vision tools to other scientific communities, i.e., robotics or cognitive sciences.

In this work we propose the See, Understand and Summarize it Network (SUSiNet) a multi-task spatio-temporal network that can jointly tackle the problems of saliency estimation, visual concept understanding and video summarization. Some examples of visual concepts that often appear in videos are human faces, actions, scenes or objects [7, 49]. Let’s consider the example that is depicted in Fig. 1. During watching a video, first we “see” the visual in-

formation and we focus on the most salient spatio-temporal regions, a process which is related with the visual saliency estimation task. Afterwards, we try to “understand” what we have seen by recognizing the underlying visual concept, which in this work is approached through the action recognition task. Finally, if someone ask us to “summarize it” we will think about keeping in the summary the most important parts of the video, which in many cases may include the recognized concept.

Our approach employs a single network that is jointly end-to-end trained for all spatio-temporal tasks and incorporates a unified architecture that includes both global and task specific layers. The proposed multi-task network produces multiple output types, i.e., saliency maps or classification labels, by employing the same video input in the form of 3D video clips. To the best of our knowledge, our approach is the first that deals with these multiple spatio-temporal problems using one single network trained in an end-to-end manner. In [4] the authors have weighted the video clips by the predicted saliency map estimated by their proposed spatio-temporal attention model, in order to compute better feature representations and perform action classification using SVMs classifiers. The work of [49] concatenates deep features, learned independently for the different tasks (scenes, object actions), in order to investigate possible relationships between them. However, none of this approaches employ an end-to-end trained multi-task network for jointly solving multiple spatio-temporal problems simultaneously. Furthermore, one additional contribution is that our new network can be deeply supervised through a proposed attention module that is related to human attention as it is expressed by eye-tracking data.

For the end-to-end training of the SUSiNet we wish to have available a large scale video database that will contain all the required annotations (i.e., eyetracking data, action labels, human created summaries). However, this is not a realistic scenario since the creation of a such dataset requires a lot of human effort for the continuous labelling of a significant amount of videos with all these annotations. So, one important aspect of our proposed network is its ability to be trained with many but diverse datasets, which have been developed for each task independently, following ideas from spatial multi-task networks [34].

We have extensively evaluated SUSiNet performance for all the three tasks on seven different datasets. As we will see the multi-task network performs very close to other state-of-the-art single-task methods that are based on more sophisticated architectures or employ heavyweight inputs [10] (two streams networks and long high-resolution clips) and have been fine-tuned for maximizing a task-specific performance. On the other hand, our approach employs a single network for all tasks and requires less computational budget than having one independent network per each task.

2. Related Work

Since the literature regarding the recognition of the three spatio-temporal tasks we are exploring is quite broad, we refer only some of the most recent deep learning based work about each task. For general reviews about more classical approaches you can see [5, 46, 17].

Visual Saliency: The early CNN-based approaches for saliency were based on the adaptation of pretrained CNN models for visual recognition tasks [39, 58]. Later, in [45] both shallow and deep CNN were trained end-to-end for saliency prediction while [28, 29] trained the networks by optimizing common saliency evaluation metrics. In [44] the authors employed end-to-end Generative Adversarial Networks (GAN), while [62] has utilized multi-level saliency information from different layer through skip connections. Long Short-term Memory (LSTM) networks have also been used for tracking visual saliency both in static images [12] and video stimuli [63]. In order to improve saliency estimation in videos, many approaches employ multi-stream networks, such as RGB/Optical Flow (OF) [3], RGB/OF/Depth [40], or multiple subnets such as objectness/motion [31] or saliency/gaze [22] pathways.

Action Recognition: The work of [32] explored several approaches for fusing information over temporal dimension, while in [30] 3D spatio-temporal convolutions have been proposed, whose performance can be boosted when trained on large datasets [55, 57] or employing ResNet architectures [25]. Recently, many approaches have tried to separate the spatial and temporal parts of 3D convolutions [54, 59, 56], which has achieved improvements over the conventional 3D networks. The work of [51] introduced the two-stream networks that employ RGB and optical flow inputs for modelling the appearance and motion information. This framework has become the basis for many other methods [60, 20, 18, 61, 11, 19]. In [10] the two-stream networks were combined with 3D convolutions, forming the I3D model that holds the state-of-the-art performance on the most action datasets.

Video Summarization: Many of the recent methods approach video summarization as an optimization problem [24, 53, 66, 16, 47], or they employ recurrent neural networks [67, 69], i.e., LSTMs, in order to model long-time dependencies inside the video which are very important in summary creation. In [41] the authors proposed a summarization framework based on GANs, where generator consists of an LSTM-based autoencoder and the discriminator is another LSTM. The work of [68] combines sequential models for the summary creation with a retrospective encoder which maps the summaries to an abstract semantic space. Following a different framework, the authors of [50], inspired by the progress in semantic segmentation, proposed fully convolutional sequence networks as an alternative approach to recurrent networks.

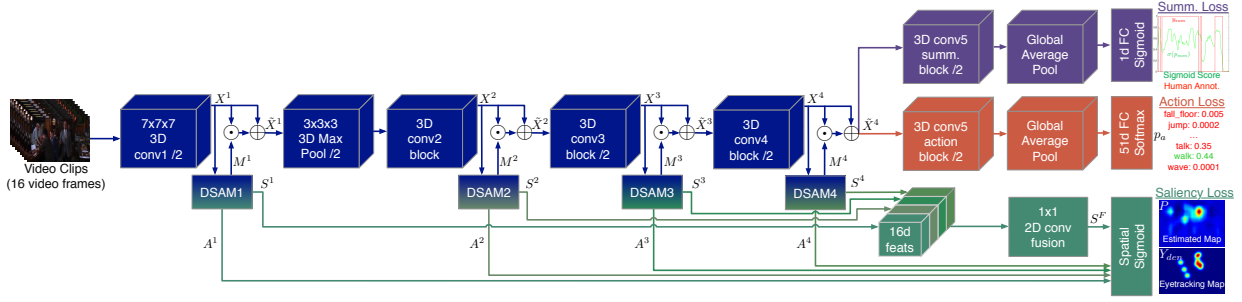


Figure 2. SUSiNet architecture: the multi-task spatio-temporal network is based on the ResNet architecture and has three different branches associated with the different spatio-temporal tasks.

3. Multi-task Spatio-Temporal Network

The proposed spatio-temporal network deals with three different tasks simultaneously and produces different types of outputs by employing the same video input in the form of small video clips. For the saliency estimation, where we face a spatial estimation problem, the network output consists of a saliency map, while for the action classification task we have the classical softmax scores. For the video summarization we need to estimate video’s segments importance scores, which indicate whether a segment will be included in the summary, as the sigmoid scores of a binary classification problem.

3.1. Global Architecture with Deep Supervision

The whole architecture of our multi-task network, which is shown in Fig. 2, is based on the general ResNet architecture [27] and specifically the 3D extension proposed in [25] for the problem of action classification. The global pathway of the network with parameters \mathbf{W}_{GL} (dark blue), which is shared among all the tasks, includes the first four convolutional blocks conv1, conv2, conv3, conv4 from the employed ResNet version that provides outputs $X^m, m = 1, \dots, 4$ in different spatial and temporal scales. In order to enhance the most salient regions of these feature representations, we apply an attention mechanism by taking the element-wise product between each channel of the feature map X^m and the attention map M^m :

$$\tilde{X}_m = (1 + M^m) \odot X_m, \quad m = 1, \dots, 4. \quad (1)$$

The attention map is obtained by our proposed Deeply Supervised Attention Module (DSAM) based on the idea of deep supervision that has been used in edge detection [64], object segmentation [9] and static saliency [62]. In contrary to these previous works the proposed module is used for both enhancing the feature representations of the global network as well as providing the multi-level saliency maps for the task of spatio-temporal saliency. Thus, the DSAM parameters \mathbf{W}_{AM}^m are trained by both the main-path of the

network, which is shared among all the tasks, and the eye-tracking data that are used for the task of saliency estimation through the skip connections of the Fig. 2. In this way, we enrich our network with an attention module that is related to human attention as it is expressed by eye-tracking data.

Figure 3 shows the architecture of the attention module applied at level m . It includes an averaging pooling in the temporal dimension followed by two spatial convolution layers that provide the saliency features S^m and the activation map A^m . Both of these representations are up-sampled (using the appropriate deconvolution layers) to the initial image dimensions and used for the deep supervision of the module as well as for the multi-level saliency estimation. The attention map $M^m(x, y)$ is given through a spatial softmax operation applied at the activation map $A^m(x, y)$:

$$M^m(x, y) = \frac{\exp(A^m(x, y))}{\sum_x \sum_y \exp(A^m(x, y))}. \quad (2)$$

3.2. Task-specific Sub-Networks

3.2.1 Visual Saliency Module

Since in the saliency estimation we face a dense prediction problem, we need to employ a fully convolution sub-network with parameters \mathbf{W}_{sal} (see green parts of Fig. 2) that takes advantage from the concatenated multi-level saliency features S^m of the DSAM components and produces the final fused saliency map S^F which corresponds to the given video clip. For the training of the network parameters $\mathbf{W}'_{sal} = [\mathbf{W}_{sal}, \mathbf{W}_{AM}^{1:4}, \mathbf{W}_{GL}]$, which are associated with visual saliency, the deep attention supervision of the whole multi-task network and the global branch we construct a loss that compares the saliency map S^F and the activations A^m with the ground truth maps Y_{sal} obtained by the eye-tracking data:

$$\mathcal{L}_{sal}(\mathbf{W}'_{sal}) = \mathcal{D}(\mathbf{W}'_{sal} | \sigma(S^F), Y_{sal}) + \sum_{m=1}^4 \mathcal{D}(\mathbf{W}_{AM}^m | \sigma(A^m), Y_{sal}), \quad (3)$$




Figure 3. Deeply Supervised Attention Module (DSAM) enhances the global network’s representations and provides the multi-level saliency maps for the task of spatio-temporal saliency.

where $\sigma(\cdot)$ denotes the sigmoid non-linearity and $\mathcal{D}(\cdot)$ is a loss function between the estimated and the ground truth 2D maps. In the saliency evaluation several different metrics are employed in order to compare the predicted saliency map $P \in [0, 1]^{N_X \times N_Y}$ with the eyetracking data [8]. As ground truth maps we are using either the map of fixation locations $Y_{fix} \in \{0, 1\}^{N_X \times N_Y}$ on the image plane of size $N_X \times N_Y$ or the dense saliency map $Y_{den} \in [0, 1]^{N_X \times N_Y}$, which arises by convolving the binary fixation map with a gaussian kernel. Thus, as $\mathcal{D}(\cdot)$ we employ three loss functions associated with the different aspects of saliency evaluation. The first is the cross-entropy loss between the predicted map P and the thresholded dense map \tilde{Y}_{den} :

$$\mathcal{D}_{CE}(\mathbf{W}|P, \tilde{Y}_{den}) = - \sum_{x,y} \tilde{Y}_{den}(x,y) \odot \log(P(x,y; \mathbf{W})) + (1 - \tilde{Y}_{den}(x,y)) \odot (1 - \log(P(x,y; \mathbf{W}))). \quad (4)$$

In order to handle the strong imbalance between the salient and non-salient pixels we take a variant of the above loss, which has been effectively used in other imbalanced tasks as boundary detection [64, 33, 42]:

$$\tilde{\mathcal{D}}_{CE}(\mathbf{W}|P, \tilde{Y}_{den}) = -\beta \cdot \sum_{x,y \in \mathcal{Y}_+} \log(P(x,y; \mathbf{W})) - (1 - \beta) \cdot \sum_{x,y \in \mathcal{Y}_-} (1 - \log(P(x,y; \mathbf{W}))), \quad (5)$$

where \mathcal{Y}_+ , \mathcal{Y}_- are the set of salient and non-salient pixels respectively and $\beta = |\mathcal{Y}_-|/(|\mathcal{Y}_+| + |\mathcal{Y}_-|)$. The second employed loss function is based on the linear Correlation Coefficient (CC) that is widely used in saliency evaluation and measures the linear relationship between the predicted saliency P and the dense ground truth map Y_{den} :

$$\mathcal{D}_{CC}(\mathbf{W}|P, Y_{den}) = - \frac{\text{cov}(P(x,y; \mathbf{W}), Y_{den}(x,y))}{\rho(P(x,y; \mathbf{W})) \cdot \rho(Y_{den}(x,y))}, \quad (6)$$

where cov , ρ denote the covariance and the standard deviation respectively.

The last loss is derived from the Normalized Scanpath Saliency (NSS) metric, which is computed as the estimated map values $\tilde{P}(x,y; \mathbf{W}) = \frac{P(x,y; \mathbf{W}) - \mu(P(x,y; \mathbf{W}))}{\rho(P(x,y; \mathbf{W}))}$, after zero mean normalization and unit standardization, at human fixation locations ($Y_{fix}(x,y) = 1$):

$$\mathcal{D}_{NSS}(\mathbf{W}|\tilde{P}, Y_{fix}) = - \frac{1}{N_f} \sum_{x,y} \tilde{P}(x,y; \mathbf{W}) \odot Y_{fix}(x,y), \quad (7)$$

where $N_f = \sum_{x,y} Y_{fix}(x,y)$ denotes the total number of fixation points.

The final loss of the j -th input sample for the task of visual saliency estimation is given by a weight combination of the losses $\mathcal{L}_{CE}^j, \mathcal{L}_{CC}^j, \mathcal{L}_{NSS}^j$, which are given by (3) using the corresponding loss functions $\mathcal{D}_{CE}^j, \mathcal{D}_{CC}^j, \mathcal{D}_{NSS}^j$:

$$\mathcal{L}_{sal}^j(\mathbf{W}'_{sal}) = w_1 \mathcal{L}_{CE}^j + w_2 \mathcal{L}_{CC}^j + w_3 \mathcal{L}_{NSS}^j, \quad (8)$$

where w_1, w_2, w_3 are the weights of each loss type.

3.2.2 Action Recognition Module

For the action recognition problem, which constitutes a classical multi-class problem, we build the task specific layers (with parameters \mathbf{W}_{act}) after the output \tilde{X}_4 of the global branch. As we can see from Fig. 2 (orange blocks), we have a 3D convolutional block, which has identical structure as the *conv5* block of the employed ResNet architecture, a global average pooling across the temporal dimension and a C_a -dimension fully connected layer, where C_a is the number of classes. For the training of action-related parameters $\mathbf{W}'_{act} = [\mathbf{W}_{act}, \mathbf{W}_{AM}^{1:4}, \mathbf{W}_{GL}]$ we employ the standard multi-class cross-entropy for the softmax activations $p_a(c; \mathbf{W}'_{act})$, $c = 1, \dots, C_a$ of the final layer:

$$\mathcal{L}_{act}^j(\mathbf{W}'_{act}) = - \log p_a^j(c_j; \mathbf{W}'_{act}), \quad (9)$$

where p_a^j, c_j denote the activation and the ground truth class of the j -th input sample respectively.

3.2.3 Summarization Module

Regarding the summarization task we employ a sub-network with parameters \mathbf{W}_{sum} (Fig. 2 - purple blocks) that has a similar structure with the one we used for action recognition, with the difference that the last fully connected layer has only one dimension since we have a binary classification problem (important vs. non-important video segments). The importance score of each video clip is given by the sigmoid activation of the final full-connected layer $\sigma(p_{sum}) \in [0, 1]$, while for the training of the all task-related parameters $\mathbf{W}'_{sum} = [\mathbf{W}_{sum}, \mathbf{W}_{AM}^{1:4}, \mathbf{W}_{GL}]$ we employ the binary cross-entropy (BCE) loss. Since in most annotated databases only a small portion of the whole video

is annotated as important and selected for final summary, the ground truth data are heavily biased. Thus we use a weighted variant of the BCE based on the ratio $\gamma = \frac{|S_-|}{|S_+|}$ between the number of negative and positive samples in the whole training dataset:

$$\mathcal{L}_{sum}^j(\mathbf{W}'_{sum}) = -\gamma \cdot y_{sum}^j \cdot \log(\sigma(p_{sum}^j(\mathbf{W}'_{sum}))) - (1 - y_{sum}^j) \cdot (1 - \log(\sigma(p_{sum}^j(\mathbf{W}'_{sum})))), \quad (10)$$

where $y_{sum}^j \in [0, 1]$ denotes the ground truth annotation regarding the importance of the j -th video clip.

3.3. Multi-task Training

For the end-to-end training of the whole multi-task spatio-temporal network ($\mathbf{W}_{all} = [\mathbf{W}_{GL}, \mathbf{W}_{AM}^{1:4}, \mathbf{W}_{sal}, \mathbf{W}_{act}, \mathbf{W}_{sum}]$) we can simply minimize the sum of the above task-specific losses over all the samples of the batch \mathcal{B} :

$$\mathcal{L}(\mathbf{W}_{all}) = \alpha_{sal} \sum_{j \in \mathcal{B}} \mathcal{L}_{sal}^j(\mathbf{W}'_{sal}) + \alpha_{act} \sum_{j \in \mathcal{B}} \mathcal{L}_{act}^j(\mathbf{W}'_{act}) + \alpha_{sum} \sum_{j \in \mathcal{B}} \mathcal{L}_{sum}^j(\mathbf{W}'_{sum}), \quad (11)$$

where $\alpha_{sal}, \alpha_{act}, \alpha_{sum}$ are weights that control the contribution of each task. This approach, which has been followed in many static multi-task networks [13, 15, 21], assumes that each sample of the batch has annotations for all tasks. However, as [34] has mentioned, this is not a realistic scenario, especially for our tasks where none of the annotations can be derived from the other tasks' annotations, as in object detection and semantic segmentation. Thus, we use the Asynchronous Stochastic Gradient Descent (SGD) algorithm, which has been proposed in [34], that allows us to have different effective batchsizes $\mathcal{B}_{sal}, \mathcal{B}_{act}, \mathcal{B}_{sum}$ and update the parameters of the task-specific layers once we have seen enough samples. The updates for the shared parameters $\mathbf{W}'_{GL} = [\mathbf{W}_{GL}, \mathbf{W}_{AM}^{1:4}]$ of the multi-task network are based on the sum of the gradients from all losses:

$$d\mathbf{W}'_{GL} = \sum_{j \in \mathcal{B}} \alpha_{sal} \nabla_{\mathbf{W}'_{GL}} \mathcal{L}_{sal}^j(\mathbf{W}'_{sal}) + \alpha_{act} \nabla_{\mathbf{W}'_{GL}} \mathcal{L}_{act}^j(\mathbf{W}'_{act}) + \alpha_{sum} \nabla_{\mathbf{W}'_{GL}} \mathcal{L}_{sum}^j(\mathbf{W}'_{sum}), \quad (12)$$

where $\mathcal{B} = \mathcal{B}_{sal} \cup \mathcal{B}_{act} \cup \mathcal{B}_{sum}$ is the total minibatch that contains all the training samples. The updates for the task-specific parameters depend only on the gradient of each different loss:

$$\begin{aligned} d\mathbf{W}_{sal} &= \sum_{j \in \mathcal{B}_{sal}} \alpha_{sal} \nabla_{\mathbf{W}_{sal}} \mathcal{L}_{sal}^j(\mathbf{W}'_{sal}) \\ d\mathbf{W}_{act} &= \sum_{j \in \mathcal{B}_{act}} \alpha_{act} \nabla_{\mathbf{W}_{act}} \mathcal{L}_{act}^j(\mathbf{W}'_{act}) \\ d\mathbf{W}_{sum} &= \sum_{j \in \mathcal{B}_{sum}} \alpha_{sum} \nabla_{\mathbf{W}_{sum}} \mathcal{L}_{sum}^j(\mathbf{W}'_{sum}) \end{aligned} \quad (13)$$

3.4. Implementation

Our implementation and experimentation with the proposed multi-task network uses as backbone the 3D ResNet-50 architecture [25] that has showed competitive performance against other deeper architectures for the task of action recognition in terms of performance and computational budget. As starting point we used the weights from the pre-trained model in the Kinetics 400 database.

Training: For the training we used the asynchronous version of stochastic gradient descent with momentum 0.9 while we also assign a weight decay of 1e-5 for regularization. We have also employed effective batchsizes of 128 samples for all tasks while the learning rate has started from 0.01 and divided by 10 when the loss saturated. The weights w_1, w_2, w_3 for the saliency loss are selected 0.1, 2, 1 after experimentation, while the ratio γ in the summarization loss was set to 3.06 based on the statistics of the employed training datasets. The weights $\alpha_{sal}, \alpha_{act}, \alpha_{sum}$, which control the importance of each task, have been experimentally tuned to 0.1, 1, 1 (based on the losses' ranges) in order to avoid the overfitting of the network to one task.

Data Augmentation: The input samples in the network consist of 16-frames RGB video clips spatially resized at 112×112 pixels. We have also applied data augmentation for random generation of training samples. For the action recognition task we randomly sampled a 16-frame clip from each training video and afterwards we followed the procedure of random multi-scale spatial cropping and flipping, which is described in [61]. For the summarization task we divided the initial long-duration videos into 90-frames non-overlapping segments and generated the 16-frames clip following the same procedure as in action recognition task. Regarding the human annotations we took its average inside the created clips, that gave training samples with slightly different annotation scores and helped us to avoid the network's overfitting. For data augmentation in the saliency estimation task, we followed a similar approach as in summarization task without the random cropping step. We applied the same spatial transformations to the 16 frames of the video clip and the eye-tracking based saliency maps of the median frame, which has been considered as the ground truth map of the whole clip.

Testing: During the testing phase, for the action recognition task we extracted the network predictions using a 16-frames non-overlapping sliding window where each clip is spatially cropped at the center position with scale 1. Then, we computed the final action label for each video by simply averaging the clips' predictions, while for the summarization task we took frame-wise importance scores by repeating the values of the 16-frame clips' scores. Finally, for saliency estimation we obtained an estimated saliency map per frame using a 16-frame sliding window with step 1 without any spatial cropping.

Dataset Method	DIEM				DFK1K				ETMD			
	CC \uparrow	NSS \uparrow	AUC-J \uparrow	sAUC \uparrow	CC \uparrow	NSS \uparrow	AUC-J \uparrow	sAUC \uparrow	CC \uparrow	NSS \uparrow	AUC-J \uparrow	sAUC \uparrow
SUSiNet (1-task) [ST]	0.6138	2.4267	0.8736	0.6747	0.4676	2.5908	0.8843	0.6991	0.5523	2.8365	0.9173	0.7312
SUSiNet (multi) [ST]	0.5614	2.1398	0.8810	0.6736	0.4116	2.2092	0.8910	0.6980	0.4780	2.3642	0.9162	0.7272
Deep-Net [45] [S]	0.4305	1.6238	0.8401	0.6262	0.2969	1.5804	0.8421	0.6432	0.3438	1.6523	0.8712	0.6588
DVA [62] [S]	0.5179	2.1607	0.8599	0.6400	0.3593	2.0644	0.8609	0.6572	0.4228	2.2507	0.8848	0.6843
SAM [12] [S]	0.5352	2.2482	0.8651	0.6429	0.3684	2.1180	0.8680	0.6562	0.4345	2.3155	0.8890	0.6875
ACLNet [63] [ST]	0.5626	2.2168	0.8717	0.6228	0.4167	2.2962	0.8883	0.6523	0.4508	2.2058	0.9073	0.6482
DeepVS [31] [ST]	0.4885	2.0352	0.8448	0.6248	0.3500	1.9680	0.8561	0.6405	0.4316	2.3030	0.8955	0.6672

Table 1. Evaluation results for the visual saliency estimation task. In most cases, the proposed multi-task SUSiNet outperforms the existing state-of-the-art methods for video saliency over all three different datasets according the four evaluation metrics. [ST] stands for spatio-temporal models while [S] denotes a spatial only model that is applied to each frame independently.

4. Multiple Tasks Evaluation

4.1. Datasets

For the training and the evaluation of the proposed multi-task network we wish to have a large scale video database that will contain eyetracking annotation, labelling of the performed actions as well as continuous human annotation of the frames importance or equivalently human created summaries. However, this is not a realistic scenario since many datasets have been developed for each task but none of them contains all of the three required types of annotation. Very recently, [49] proposed a multi-task and multi-label video dataset aiming to the recognition of different visual concepts (scenes, objects, actions) which are different from our investigated tasks. Note that since our multi-task network is modular, it could be extended to recognize and understand more visual concepts such as objects or scenes.

The most relevant dataset to our tasks is the COGNIMUSE database [70, 1], which constitutes a video database annotated with ground-truth annotations for frame-wise sensory and semantic importance as well as audio and visual events. It is a generic database that has been used for video summarization [36], as well as audio-visual concept recognition [7]. The creators of the database have also developed the Eye-Tracking Movie Database (ETMD) [35], which contains eyetracking annotations for a subset of the COGNIMUSE videos. For our experiments we have used the 30-minutes excerpts from the seven movies ¹ as well as the full movie “Gone With the Wind” (GWW) that they have at least two of the three annotation types (see Table 4). For the training we followed an 8-fold (leaving one movie out) cross-validation approach.

One important aspect of the proposed multi-task network is its ability to be trained with diverse datasets. So, we employ five more state-of-the-art datasets, containing annotations only for a specific task in order to increase the training set as well as compare our results with other state-of-the-art methods. Specifically, for the visual saliency estimation we

¹“A Beautiful Mind” (BMI), “Gladiator” (GLA), “Chicago” (CHI), “Finding Nemo” (FNE), “Lord of the Rings - the Return of the King” (LOR), “Crash” (CRA), “The Departed” (DEP)

employ the DIEM dataset [43], which contains eyetracking data for 84 videos with duration between 27-217 sec from 50 observers, and the DFK1K [63], with eyetracking data from 17 observers over 1000 videos with duration 17-42 sec. During the experiments, for the DIEM we followed the “train-test” split of [6], while for the DFK1K we used the validation set for our testing since the test set is not publicly available. Regarding the action recognition task we employ the HMDB51 dataset [38] that includes 6766 video from 51 human action classes. We decided to use this additional dataset because its classes have also been included in the COGNIMUSE dataset. Finally, for the summarization task we used the SumMe [23] and the TVSum50 [53] datasets that include 25 (1.5 to 6.5 minutes length) and 50 (1 to 5 minutes length) videos respectively, mainly from YouTube resources. For the experiment that involves these datasets we have followed a 5-fold cross validation.

4.2. Experimental Results

For the evaluation of the multi-task network we have constructed two types of experiments. In the first, which we refer as “SUSiNet (1-task)” we trained our network independently for each task using only the task-related datasets. In the “SUSiNet (multi)” we trained the multi-task network jointly for all the three task employing all the available datasets. Next, we evaluate our results for each different task and compare them with several methods that have achieved state-of-the-art performance for each task independently.

Saliency Estimation Evaluation: In Table 1 we present the evaluation results of the proposed SUSiNet on the 3 different datasets and compare its performance against 5 state-of-the-art methods (using their publicly available codes). We employed four widely-used evaluation metrics [8]: CC, NSS, AUC-Judd (AUC-J) and shuffled AUC (sAUC). In the sAUC we have selected the negative samples from the union of all viewers’ fixations across all other frames except the frame for which we compute the AUC. As we see, our method outperforms all the other methods over all the datasets according to all employed metrics. Note there is a small decrease in the CC and NSS scores of the

Method	Pre-train dataset	Aver. Accuracy
SUSiNet (1-task)	Kinetics	60.2
SUSiNet (multi)	Kinetics	62.7
C3D [55]	Sports-1M	51.6
3D ResNet-18 [25]	Kinetics	56.4
3D ResNet-50 [25]	Kinetics	61.0
3D ResNeXt-101 [25]	Kinetics	63.8
RGB I3D (64f) [10]	ImageNet & miniKinetics	66.4
Other Methods		
F _{ST} CN [54]	ImageNet	59.1
ARTNet [59]	Kinetics	67.6
R(2+1)D-RGB [56]	Kinetics	74.5
Two-Stream TSN [61]	ImageNet	68.5
Two-Stream STM-Nets [19]	ImageNet	68.9
Two-Stream I3D (64f) [10]	ImageNet & Kinetics	80.7
Two-Stream R(2+1)D [56]	Kinetics	78.7

Table 2. Evaluation results for the action recognition task over all splits of HMDB51. The proposed multi-task SUSiNet outperforms the single-task as well as several state-of-the-art methods that rely on 3D convolutions.

multi-task network compared to the single-task, while according to AUC-based metrics the multi-task network performs in equally or better than the single one. Moreover, the SUSiNet, which is based on 3D spatio-temporal convolutions, achieves to outperform other spatio-temporal methods (ACLNet, DeepVS) that rely on the LSTM based modelling of the visual saliency. In Figures 1, 4 we see examples of our saliency predictions, which in most cases are focused on humans or actions.

Action Classification Evaluation: In Table 2 we see the evaluation results of our method on the HMDB51. We can observe that the multi-task network achieves better performance than the single-task. Comparing our method against several other approaches, which are based on 3D CNN networks, we see that our network performs better than the most of them. Note that our network is based on ResNet-50 architecture and uses 16-frames inputs, while 3D ResNeXt or I3D are employing more complex networks or longer clips. For completeness we also report several other methods from literature, which are based on techniques for decoupling the spatial and temporal parts of 3D convolutions or employ two-streams (RGB, Optical Flow) networks, that are not directly compared with our method. However, our proposed network is modular and can be modified and extended to include such techniques but we leave this direction for future work as the scope of this paper is to propose a multi-task spatio-temporal network rather than achieve the best performance for each single task.

Finally, in Table 4 we report our method’s results for action recognition in COGNIMUSE database, where we see again that the multi-task network slightly outperforms the single-task. Regarding the lower recognition scores (comparing with HMDB51), we have observed that COGNIMUSE dataset contains many background actions or supplementary actions that may overlap with a main ac-

Method	SumMe (F-score)	TVSum50 (F-score)
SUSiNet (1-task)	41.10	59.20
SUSiNet (multi)	40.80	57.00
vsLSTM [67]	37.6 [41.6]	54.2 [57.9]
HSA-RNN [69]	44.1	59.8
SEQ2SEQ [68]	40.8	56.3
SUM-FCN [50]	47.5 [51.1]	56.8 [59.2]
Other Methods		
Gygli et al. [24]	39.7	-
dppLSTM [67]	38.6 [42.9]	54.7 [59.6]
SUM-GAN [41]	41.7 [43.6]	56.3 [61.2]
re-SEQ2SEQ [68]	[44.9]	[63.9]

Table 3. Evaluation results for the video summarization task over the SumMe and TVSum50 datasets.

Task	Saliency (sAUC)		Action (Acc.)		Summar. (AUC)	
SUSiNet	1-task	multi	1-task	multi	1-task	multi
BMI	-	-	51.54	49.88	0.7831	0.8023
GLA	0.6859	0.6727	48.92	46.77	0.7863	0.7843
CHI	0.7601	0.7565	49.41	50.82	0.7901	0.7826
FNE	0.7224	0.7236	-	-	0.5490	0.5306
LOR	0.7297	0.7325	50.70	54.93	0.7602	0.7557
CRA	0.7056	0.7058	49.83	47.83	0.7424	0.7105
DEP	0.7837	0.7721	58.86	60.76	0.8069	0.8279
GWV	-	-	36.24	37.70	0.6762	0.6806
Aver.	0.7312	0.7272	49.36	49.81	0.7368	0.7343

Table 4. Evaluation results of the proposed multi-task SUSiNet for the three different tasks over the COGNIMUSE database. We report results for each movie independently as well as the average performance for each task.

tion (i.e., in Fig. 4 the action “turn” overlaps with the “run”) and thus it constitutes a very challenging dataset.

Video Summarization Evaluation: For the evaluation over the SumMe and TVSum50 datasets we have employed the evaluation protocol of [67] that is based on the F-score between a generated keyshot-based summary (shots are temporally segmented using KTS [48]), with length 15% of the original video duration, and the user created summaries. In Table 3, we present the evaluation results for the summarization task over the SumMe and TVSum50 datasets compared against various other state-of-the-art approaches. Brackets “[.]” denote results obtained using an augmented dataset that includes videos from auxiliary datasets (YouTube [14], OVP [14, 2]) and it is not directly compared to our method. As we see, the multi-task SUSiNet performs very close to its single-task variant and outperforms many methods that are based on the sequential estimation of the clip based importance score, i.e., using LSTM networks. On the other hand, our network cannot perform better than other methods of literature that operate on the whole video (i.e., using retrospective encoders [68]) or employ a larger number of video frames (i.e., 128 frames in SUM-FCN [50]), especially when they are trained on the augmented dataset. However, these approaches could be added as post-hoc task-specific components and increase the performance of our network.

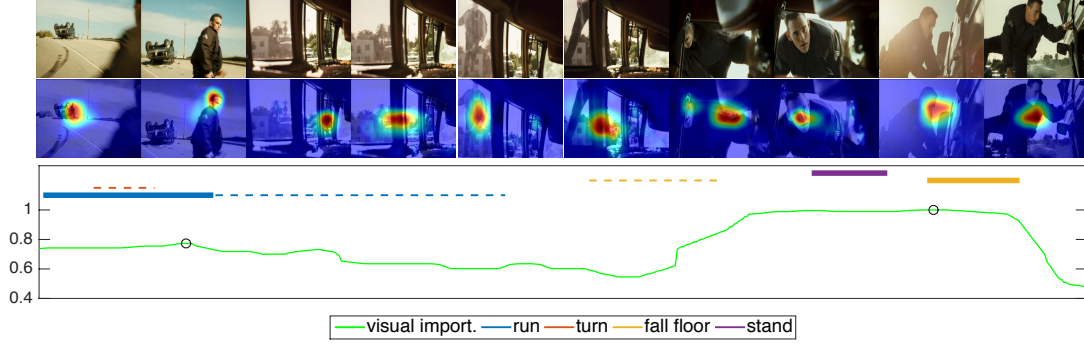


Figure 4. Example results of the proposed multi-task SUSiNet for the three different tasks: saliency estimation, action recognition, video summarization. The second row consists of the estimated saliency maps, while in the figure we also see the action recognition results (with dotted lines the annotated actions that are not correctly recognized) and the green curve that indicates the frame-wise importance score.




Figure 5. ROC curves and corresponding AUC scores for the different methods on COGNIMUSE database (summarization task).

Regarding the evaluation of the summarization task in the COGNIMUSE database we have followed the evaluation procedure described in [70]. Specifically, the summarization task is approached as a two-class classification problem, where multiple thresholds are applied to the estimated frame-wise importance scores in order to obtain results for various compression rates and thus produce summaries of various lengths. Then, the AUC metric is computed from the ROC curve that results from the different thresholds. Table 4 presents the AUC scores for each movie independently as well as the average across all movies. We see that the evaluation scores of the multi- and single-task network are very close to each other, while in some cases, such as the full-length movie GWW, the multi-task network achieves better performance. It is worth to note that the low performance (around chance) for the FNE movie can be attributed to the absence of any other animation movie in the training set. As in the other cases, we have also compared our method with two other approaches that have been appeared in literature for the COGNIMUSE database: ICIP15 [36] which employs hand-crafted features and IVMSP18 [37] that is based on C3D network. As we see in Fig. 5 the SUSiNet outperforms the two other baseline methods

according to the ROC-AUC metric.

Discussion: In Table 4 we present evaluation results for the proposed SUSiNet in all the three different tasks over the COGNIMUSE database. We see that the multi-task SUSiNet, which has been jointly trained end-to-end over all tasks using very diverse datasets, can efficiently face three different problems, achieving almost the same (or in some cases better) performance to a similar single-task network, which has been trained explicitly for the specific task. In addition, as we have observed, the multi-task network performs very close to other state-of-the-art single-task methods that are based on more complex architectures or employ multiple streams networks. Figure 4 depicts a qualitative example of SUSiNet on a movie excerpt from COGNIMUSE database, which contains sequential actions. Our network focuses on the salient regions of the video, performs action recognition and computes frame-wise importance scores simultaneously. We also observe in Fig. 4 that segments with the correctly recognized actions have been selected by multi-task network to be included in the video summary.

5. Conclusions

In this work, we have proposed a multi-task spatio-temporal network that can jointly tackle the problems of saliency estimation, action recognition and video summarization. Our method employs only a single network for all tasks, which is jointly trained for all tasks using diverse datasets. The extensive evaluation indicates that the multi-task network performs equally well or in some cases even better than the single-task methods while it requires less computational budget than having one different network for each task. As future work, we intend to extend our network by employing more complex or multiple-streams architectures and improve the task-specific layers by incorporating recent advances from state-of-the-art single-task methods.

References

- [1] <http://cognimuse.cs.ntua.gr/database>.
- [2] Open Video Project. <http://www.open-video.org/>.
- [3] Cagdas Bak, Aysun Kocak, Erkut Erdem, and Aykut Erdem. Spatio-temporal saliency networks for dynamic saliency prediction. *IEEE Trans. Multimedia*, 20(7):1688–1698, 2018.
- [4] Loris Bazzani, Hugo Larochelle, and Lorenzo Torresani. Recurrent mixture density network for spatiotemporal visual attention. In *Int. Conf. on Learning Representations (ICLR)*, 2017.
- [5] Ali Borji and Laurent Itti. State-of-the-art in visual attention modeling. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 35(1):185–207, 2013.
- [6] Ali Borji, Dicky N. Sihite, and Laurent Itti. Quantitative analysis of human-model agreement in visual saliency modeling: A comparative study. *IEEE Trans. Image Process.*, 22(1):55–69, 2013.
- [7] Giorgos Bouritsas, Petros Koutras, Athanasia Zlatintsi, and Petros Maragos. Multimodal visual concept learning with weakly supervised techniques. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [8] Zoya Bylinskii, Tilke Judd, Aude Oliva, Antonio Torralba, and Frédo Durand. What do different evaluation metrics tell us about saliency models? corr abs/1604.03605 (2016), 2016.
- [9] Sergi Caelles, Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Laura Leal-Taixé, Daniel Cremers, and Luc Van Gool. One-shot video object segmentation. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [10] Joao Carreira and Andrew Zisserman. Quo vadis, action recognition? a new model and the kinetics dataset. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [11] Quan-Qi Chen and Yu-Jin Zhang. Sequential segment networks for action recognition. *IEEE Signal Processing Letters*, 24/5:712–716, May 2017.
- [12] Marcella Cornia, Lorenzo Baraldi, Giuseppe Serra, and Rita Cucchiara. Predicting Human Eye Fixations via an LSTM-based Saliency Attentive Model. *IEEE Trans. Image Process.*, 27(10):5142–5154, 2018.
- [13] Jifeng Dai, Kaiming He, and Jian Sun. Instance-aware semantic segmentation via multi-task network cascades. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [14] Sandra Eliza Fontes De Avila, Ana Paula Brandão Lopes, Antonio da Luz Jr, and Arnaldo de Albuquerque Araújo. Vsumm: A mechanism designed to produce static video summaries and a novel evaluation method. *Pattern Recognition Letters*, 32(1):56–68, 2011.
- [15] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, 2015.
- [16] Ehsan Elhamifar and M Clara De Paolis Kaluza. Online summarization via submodular and convex optimization. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [17] Georgios Evangelopoulos, Athanasia Zlatintsi, Alexandros Potamianos, Petros Maragos, Konstantinos Rapantzikos, Georgios Skoumas, and Yannis Avrithis. Multimodal saliency and fusion for movie summarization based on aural, visual, textual attention. *IEEE Trans. Multimedia*, 15(7):1553–1568, Nov. 2013.
- [18] Christoph Feichtenhofer, Axel Pinz, and Richard P Wildes. Spatiotemporal Residual Networks for Video Action Recognition. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2016.
- [19] Christoph Feichtenhofer, Axel Pinz, and Richard P Wildes. Spatiotemporal multiplier networks for video action recognition. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [20] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [21] Georgia Gkioxari, Ross Girshick, and Jitendra Malik. Contextual action recognition with r* cnn. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, 2015.
- [22] Siavash Gorji and James J. Clark. Going from image to video saliency: Augmenting image salience with dynamic attentional push. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [23] Michael Gygli, Helmut Grabner, Hayko Riemenschneider, and Luc Van Gool. Creating summaries from user videos. In *Proc. European Conf. on Computer Vision (ECCV)*, 2014.
- [24] Michael Gygli, Helmut Grabner, and Luc Van Gool. Video summarization by learning submodular mixtures of objectives. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [25] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [26] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, 2017.
- [27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [28] Xun Huang, Chengyao Shen, Xavier Boix, and Qi Zhao. Salicon: Reducing the semantic gap in saliency prediction by adapting deep neural networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [29] Saumya Jetley, Naila Murray, and Eleonora Vig. End-to-end saliency mapping via probability distribution prediction. *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [30] Shuiwang Ji, Wei Xu, Ming Yang, and Kai Yu. 3D Convolutional Neural Networks for Human Action Recognition. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 35:221–231, Jan 2013.

- [31] Lai Jiang, Mai Xu, Tie Liu, Minglang Qiao, and Zulin Wang. Deepvps: A deep learning based video saliency prediction approach. In *Proc. European Conf. on Computer Vision (ECCV)*, 2018.
- [32] Andrej Karpathy, George Toderici, Sanketh Shetty, Thomas Leung, Rahul Sukthankar, and Li Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [33] Iasonas Kokkinos. Pushing the boundaries of boundary detection using deep learning. In *Int. Conf. on Learning Representations (ICLR)*, 2015.
- [34] Iasonas Kokkinos. Ubertnet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [35] Petros Koutras and Petros Maragos. A perceptually based spatio-temporal computational framework for visual saliency estimation. *Signal Processing: Image Communication*, 38:15–31, 2015.
- [36] Petros Koutras, Athanasia Zlatintsi, Elias Iosif, Athanasios Katsamanis, Petros Maragos, and Alexandros Potamianos. Predicting audio-visual salient events based on visual, audio and text modalities for movie summarization. In *Proc. IEEE Int. Conf. on Image Processing (ICIP)*, 2015.
- [37] Petros Koutras, Athanasia Zlatintsi, and Petros Maragos. Exploring cnn-based architectures for multimodal salient event detection in videos. In *Proc. 13th IEEE Image, Video, and Multidimensional Signal Processing Workshop (IVMSP)*, 2018.
- [38] Hildegard Kuehne, Hueihan Jhuang, Estíbaliz Garrote, Tomaso Poggio, and Thomas Serre. Hmdb: a large video database for human motion recognition. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, 2011.
- [39] Matthias Kümmeler, Lucas Theis, and Matthias Bethge. Deep gaze I: Boosting saliency prediction with feature maps trained on imagenet. In *Int. Conf. on Learning Representations Workshops (ICLRW)*, 2015.
- [40] George Leifman, Dmitry Rudoy, Tristan Swedish, Eduardo Bayro-Corrochano, and Ramesh Raskar. Learning gaze transitions from depth to improve video saliency estimation. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, 2017.
- [41] Behrooz Mahasseni, Michael Lam, and Sinisa Todorovic. Unsupervised video summarization with adversarial lstm networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, volume 1, 2017.
- [42] Kevis-Kokitsi Maninis, Jordi Pont-Tuset, Pablo Arbeláez, and Luc Van Gool. Convolutional oriented boundaries. In *Proc. European Conf. on Computer Vision (ECCV)*, 2016.
- [43] Parag K Mital, Tim J Smith, Robin L Hill, and John M Henderson. Clustering of gaze during dynamic scene viewing is predicted by motion. *Cognitive Computation*, 3(1):5–24, 2011.
- [44] Junting Pan, Cristian Canton, Kevin McGuinness, Noel E. O’Connor, Jordi Torres, Elisa Sayrol, and Xavier and Giro-i Nieto. Salgan: Visual saliency prediction with generative adversarial networks. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition Workshop (CVPRW)*, 2017.
- [45] Junting Pan, Elisa Sayrol, Xavier Giro-i Nieto, Kevin McGuinness, and Noel E. O’Connor. Shallow and deep convolutional networks for saliency prediction. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [46] Xiaojiang Peng, Limin Wang, Xingxing Wang, and Yu Qiao. Bag of visual words and fusion methods for action recognition: Comprehensive study and good practice. *Computer Vision and Image Understanding*, 150:109–125, 2016.
- [47] Bryan A Plummer, Matthew Brown, and Svetlana Lazebnik. Enhancing video summarization via vision-language embedding. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [48] Danila Potapov, Matthijs Douze, Zaid Harchaoui, and Cordelia Schmid. Category-specific video summarization. In *Proc. European Conf. on Computer Vision (ECCV)*, 2014.
- [49] Jamie Ray, Heng Wang, Du Tran, Yufei Wang, Matt Feiszli, Lorenzo Torresani, and Manohar Paluri. Scenes-objects-actions: A multi-task, multi-label video dataset. In *Proc. European Conf. on Computer Vision (ECCV)*, 2018.
- [50] Mrigank Rochan, Linwei Ye, and Yang Wang. Video summarization using fully convolutional sequence networks. 2018.
- [51] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. In *Proc. Advances in Neural Information Processing Systems (NIPS)*, 2014.
- [52] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *Int. Conf. on Learning Representations (ICLR)*, 2015.
- [53] Yale Song, Jordi Vallmitjana, Amanda Stent, and Alejandro Jaimes. Tvsum: Summarizing web videos using titles. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [54] Lin Sun, Kui Jia, Dit-Yan Yeung, and Bertram E Shi. Human action recognition using factorized spatio-temporal convolutional networks. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, 2015.
- [55] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, 2015.
- [56] Du Tran, Heng Wang, Lorenzo Torresani, Jamie Ray, Yann LeCun, and Manohar Paluri. A closer look at spatiotemporal convolutions for action recognition. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [57] Gül Varol, Ivan Laptev, and Cordelia Schmid. Long-term temporal convolutions for action recognition. *IEEE Trans. Pattern Anal. Mach. Intell. (PAMI)*, 40(6):1510–1517, 2018.
- [58] Eleonora Vig, Michael Dorr, and David Cox. Large-scale optimization of hierarchical features for saliency prediction in natural images. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2014.
- [59] Limin Wang, Wei Li, Wen Li, and Luc Van Gool. Appearance-and-relation networks for video classification.

- In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [60] Limin Wang, Yu Qiao, and Xiaoou Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2015.
 - [61] Limin Wang, Yuanjun Xiong, Zhe Wang, Yu Qiao, Dahua Lin, Xiaoou Tang, and Luc Van Gool. Temporal segment networks: Towards good practices for deep action recognition. In *Proc. European Conf. on Computer Vision (ECCV)*, 2016.
 - [62] Wenguan Wang and Jianbing Shen. Deep visual attention prediction. *IEEE Trans. Image Process.*, 27(5):2368–2378, 2018.
 - [63] Wenguan Wang, Jianbing Shen, Fang Guo, Ming-Ming Cheng, and Ali Borji. Revisiting video saliency: A large-scale benchmark and a new model. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
 - [64] Saining Xie and Zhuowen Tu. Holistically-nested edge detection. In *Proc. IEEE Int. Conf. on Computer Vision (ICCV)*, 2015.
 - [65] Amir R Zamir, Alexander Sax, William Shen, Leonidas Guibas, Jitendra Malik, and Silvio Savarese. Taskonomy: Disentangling task transfer learning. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
 - [66] Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman. Summary transfer: Exemplar-based subset selection for video summarization. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2016.
 - [67] Ke Zhang, Wei-Lun Chao, Fei Sha, and Kristen Grauman. Video summarization with long short-term memory. In *Proc. European Conf. on Computer Vision (ECCV)*, 2016.
 - [68] Ke Zhang, Kristen Grauman, and Fei Sha. Retrospective encoders for video summarization. In *Proc. European Conf. on Computer Vision (ECCV)*, 2018.
 - [69] Bin Zhao, Xuelong Li, and Xiaoqiang Lu. Hsa-rnn: Hierarchical structure-adaptive rnn for video summarization. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2018.
 - [70] Athanasia Zlatintsi, Petros Koutras, Georgios Evangelopoulos, Nikolaos Malandrakis, Niki Efthymiou, Katerina Passtra, Alexandros Potamianos, and Petros Maragos. COGNIMUSE: A multimodal video database annotated with saliency, events, semantics and emotion with application to summarization. *EURASIP Journal on Image and Video Processing*, 2017(1):54, Aug 2017.