



Hierarchical Domain-Adapted Feature Learning for Video Saliency Prediction

G. Bellitto¹ · F. Proietto Salanitri¹ · S. Palazzo¹ · F. Rundo² · D. Giordano¹ · C. Spampinato¹

Received: 3 October 2020 / Accepted: 13 August 2021 / Published online: 5 October 2021
© The Author(s) 2021

Abstract

In this work, we propose a 3D fully convolutional architecture for video saliency prediction that employs hierarchical supervision on intermediate maps (referred to as *conspicuity maps*) generated using features extracted at different abstraction levels. We provide the base hierarchical learning mechanism with two techniques for *domain adaptation* and *domain-specific learning*. For the former, we encourage the model to unsupervisedly learn hierarchical general features using gradient reversal at multiple scales, to enhance generalization capabilities on datasets for which no annotations are provided during training. As for domain specialization, we employ domain-specific operations (namely, priors, smoothing and batch normalization) by specializing the learned features on individual datasets in order to maximize performance. The results of our experiments show that the proposed model yields state-of-the-art accuracy on supervised saliency prediction. When the base hierarchical model is empowered with domain-specific modules, performance improves, outperforming state-of-the-art models on three out of five metrics on the DHF1K benchmark and reaching the second-best results on the other two. When, instead, we test it in an unsupervised domain adaptation setting, by enabling hierarchical gradient reversal layers, we obtain performance comparable to supervised state-of-the-art. Source code, trained models and example outputs are publicly available at <https://github.com/perceivelab/hd2s>.

Keywords Video saliency Prediction · Conspicuity networks · Conspicuity maps · Domain adaptation · Gradient reversal layer · Domain specific learning

1 Introduction

Video saliency prediction is the task of predicting human gaze fixation when perceiving dynamic scenes, and it is typically carried out by estimating spatio-temporal saliency maps from an input video sequence. Saliency prediction, in general, can be seen as the upstream processing step of multiple applications that include object detection (Girshick 2015), behavior understanding (Lim et al. 2014; Lu et al. 2017), video surveillance (Li and Lee 2007; Mark et al. 2018; Guraya et al. 2010; Yubing 2011) and video captioning (Nguyen et al. 2013; Wang et al. 2018a; Yangyu et al. 2018).

Existing video saliency prediction methods generally apply single-image saliency estimation on individual frames, and combine the results with recurrent layers to model temporally frame-level features. However, the two separate analysis stages in these models make them unable to fully capture spatio-temporal features simultaneously. Recently, 3D fully-convolutional models have addressed this limitation by progressively aggregating spatio-temporal cues, achieving state-of-the-art performance on standard benchmarks. For example, TASED-Net (Min and Corso 2019) adopts a standard encoder-decoder architecture, as largely used in semantic segmentation tasks (Ronneberger et al. 2015; Badrinarayanan et al. 2017; Noh et al. 2015), that learns a compact spatio-temporal representation, and feeds it to a decoder sub-network to perform saliency prediction. While these methods perform well, saliency prediction is constrained by the aggregated representation learned at the model's bottleneck.

To overcome this limitation, and following the success of 3D convolutional architectures, in this paper we propose *Hierarchical Decoding for Dynamic Saliency* prediction—

Communicated by Karteek Alahari.

G. Bellitto and F. Proietto Salanitri contributed equally to this work.

✉ S. Palazzo
simone.palazzo@unict.it

¹ PeRCeiVe Lab - University of Catania, Catania, Italy

² STMicroelectronics, ADG Central R&D, Catania, Italy

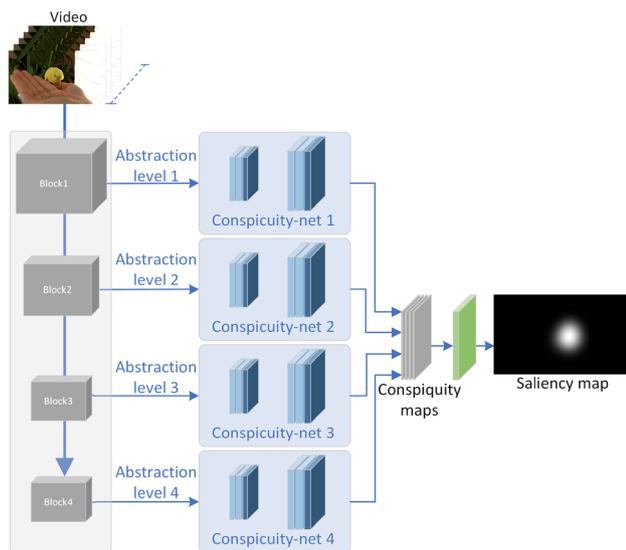


Fig. 1 HD₂S overview. Our proposed model generates multiple intermediate saliency maps by using features extracted at different abstraction levels, and combines them to predict the output map. We refer to the intermediate saliency maps as *conspicuity maps*

HD²S—model that, instead of using a compact spatio-temporal representation as in (Min and Corso 2019), generates multiple saliency maps by using features learned at different abstraction levels and then combines them to compute the final output. We refer to the intermediate saliency maps as *conspicuity maps*, as the employed architecture recalls the multi-scale model proposed in (Itti et al. 1998). Using representations extracted at different abstraction levels (from shallow to deeper) allows the model to learn both generic (and more dataset-independent) and dataset-specific features. The twofold advantage we obtain is to enhance performance on a specific dataset and, at the same time, to improve adaptation capabilities.

Our approach takes inspiration from DVA (Wang et al. 2018e), but extends it to the video domain by learning spatio-temporal cues for predicting visual saliency. More specifically, HD²S, shown in Fig. 1, is a 3D fully-convolutional network that employs an ensemble of multiple prediction models, each producing a *conspicuity-like* map at a specific abstraction level, for better saliency estimation.

As an additional contribution, we tackle the problem of generalization for video saliency prediction. Indeed, state-of-the-art methods lack domain adaptation capabilities and require a mandatory fine-tuning step to perform well on datasets that they were not trained on. As the deep learning community is moving to build more generalizable models, we argue that this holds, even more so, for saliency prediction research, given its fundamental nature in an artificial vision pipeline. To address this issue, our saliency prediction network is provided with a multi-scale domain adaption

mechanism, based on gradient reversal (Ganin 2016), that encourages the model to learn *domain-independent* features. In particular, each abstraction level of HD²S is provided with a gradient reversal layer that prevents the learned representation from becoming dataset-specific.

We also address the opposite problem, i.e., *domain-specific learning*, by adding to the model some dataset-specific modules whose parameters are learned in order to maximize performance on a given dataset.

We carry out extensive experiments testing of HD²S on multiple video saliency benchmarks (DHF1K (Wang et al. 2018c), UCF Sports (Marszalek et al. 2009; Soomro et al. 2014), Hollywood2 (Mathe and Sminchisescu 2014)) obtaining state-of-the-art performance and outperforming existing models. Furthermore, performance is boosted, as expected, when domain-specific learning is enabled. We also test thoroughly the domain adaptation capabilities of HD²S to datasets for which no annotations are available during training. Our model shows remarkable results, achieving performance comparable to state-of-the-art models that, instead, are trained (or fine-tuned) on those datasets in a standard supervised fashion.

2 Related Work

Saliency prediction has been long investigated in AI and computer vision research. In general, saliency models can be categorized in: *saliency prediction* (Wang 2019) approaches that attempt to predict the fixation points of a human observer during free-viewing (e.g., they aim to predict where people look at in a scene), and *salient object detection* (Liu 2010) methods that, instead, focus on assessing the saliency of pixels w.r.t. objects of interest (e.g., they aim to separate the salient objects from the background). Saliency methods can be further categorized according to whether they process still images (static saliency) or videos (dynamic saliency).

Static saliency has been studied for decades. Initial models, biologically-inspired (Itti et al. 1998) and employing hand-crafted features, were followed by recent CNN-based attempts (Huang et al. 2015; Pan et al. 2016, 2017; Kummerer et al. 2017; Wang et al. 2018e; Fan et al. 2018; Cornia 2018; Che 2019; Kroner 2020; Jia et al. 2020) that yield superior performance, rapidly becoming state of the art for static saliency prediction. To overcome the lack of large eye fixation datasets, CNN-based static methods rely mainly on image classification models, as backbone, exploiting their capability to extract features useful for other visual tasks. Different encoder-decoder architectures with various strategies to combine the extracted features have then been proposed. The release of larger dataset for saliency, such as MIT300 (Judd et al. 2012), SALICON (Jiang et al. 2015), and CAT2000 (Borji et al. 2015), led to a performance gain.

DeepGaze II (Kummerer et al. 2017) investigated the benefit of employing low- and high-level features in saliency prediction. Similarly, ML-NET (Cornia et al. 2016) proposed to combine low- and high-level features at the bottleneck, while (Kroner 2020) concatenates the outputs from several layers and processes them with multiple convolutional layers with different dilation rates. Another approach is to use a two-stream encoder architecture as in (Huang et al. 2015), where the image at different spatial scales is fed as input to the model, in order to extract low and high resolution information. (Fan et al. 2018), based on (Huang et al. 2015), used a similar network adding, after feature extraction, a channel weighting subnetwork that encodes contextual information. Differently from the above models, other works exploit adversarial training (Goodfellow et al. 2014) for saliency prediction, such as SalGAN (Pan et al. 2017) and GazeGAN (Che 2019). Compared to saliency models for still images, saliency prediction in videos is an even more complex problem, due to the presence of the temporal dimension and to the additional computational effort it requires. Static saliency models have been adapted to dynamic saliency by using them in frame-by-frame mode, but they are outperformed by the dynamic models that jointly process the temporal dimension.

In recent years, a common strategy has been to extend static saliency models to the video scenario by incorporating motion features (Wang et al. 2017; Shokri et al. 2020; Sun 2018). For example, (Wang et al. 2017) proposes a two-model architecture to exploit spatio-temporal features: the first module performs frame-level saliency prediction; the second module, instead, takes pairs of frames with saliency predicted by the first module, and generates a dynamic saliency map. (Shokri et al. 2020) basically employs the same architecture as (Wang et al. 2017) and self-attention, through non-local operations (Wang et al. 2018d). SaleMA (Linardos et al. 2019), instead, proposes a 2D encoder-decoder architecture with a recurrent module added to the bottleneck for integrating temporal information provided by the previous frames. Motion cues have been also included in saliency prediction through either recurrent neural networks applied to spatial feature encodings or convolutional recurrent networks. OM-CNN (Jiang et al. 2017) is a dual-stream network that extracts spatial and temporal features using YOLO (Redmon et al. 2016) and FlowNet (Dosovitskiy et al. 2015), whose respective objectness and motion features are then combined via a two-layer ConvLSTM. Similarly, ACLNet (Wang et al. 2018c) performs static saliency prediction through attention module that performs a global spatial operation on learned features. These features are then given to a ConvLSTM to model temporal information. The recent SalSAC model (Wu et al. 2020), leveraging the success of self-attention for saliency prediction (Cornia 2018; Wang et al. 2018c), proposes an architecture with a shuffled attention mechanism on multi-level features for better modeling of spatial

saliency. Correlation features between multi-level features and shuffled attention on the same features are provided to a ConvLSTM for learning temporal cues.

With the recent availability of a large-scale saliency benchmark, i.e., DHF1K (Wang et al. 2018c), 3D fully-convolutional models (Bazzani et al. 2016; Min and Corso 2019), jointly extracting spatial and temporal features, have been proposed. RMDN (Bazzani et al. 2016) processes video clips with a 3D convolutional neural network based on C3D (Tran et al. 2015), and then employs LSTMs to enforce temporal consistency among the segments. TASED-Net (Min and Corso 2019) is a 3D fully-convolutional network, based on a standard encoder-decoder architecture, for video saliency prediction without any additional feature processing steps. Similarly to the above approaches, our HD²S model is a 3D fully-convolutional network extending the multi-abstraction level analysis, proposed in (Wang et al. 2018e) for static saliency, to the video domain by learning spatio-temporal cues.

Multi-level feature learning has been already applied in several application domains, most notably in object detection through the use of *feature pyramid networks* (FPN) (He 2020). Most relevant to our approach are the works that carry out salient object detection using multi-level feature hierarchies, such as Amulet (Zhang et al. 2017) and DSS (Hou 2019). However, beside targeting static saliency prediction in images (and not in videos), those approaches apply an early-fusion mechanism of multi-level features, that are combined (through different concatenation schemas) before being further processed. Our method, instead, performs a late fusion of features: we encourage each decoding path to independently extract information from a certain abstraction layer, making sure that no inter-branch “contamination” may happen except at the very last layer, and thus pushing it to learn scale-specific and complementary saliency features.

HD²S can also be used in a domain adaption scenario to generalize across datasets without the need to be fine-tuned. Indeed, in all prediction tasks, shifts in train and test data distributions may lead to a significant degradation of the model’s performance. Trying to train a predictor capable of handling these shifts is commonly referred to as domain adaptation. Among the different domain adaptation settings,¹ we focus on *unsupervised* domain adaptation, which is the task of aligning features extracted from the model across source and target domains, without any labelled samples from the latter. Several techniques have been proposed (though not for saliency prediction), such as regularizing the maximum mean discrepancy (Long et al. 2015), minimizing correlation (Sun et al. 2016) or domain discriminability (Ganin 2016; Tzeng

¹ An extensive review of domain adaptation approaches is out of the scope of this paper and can be found in (Pan et al. 2009; Wang and Deng 2018b).

et al. 2017). An effective approach to transfer the feature distribution from source to target domains is proposed in (Ganin 2016) through the use of the *gradient reversal layer*, treating domain invariance as a binary classification problem. This approach addresses domain adaptation by adversarially forcing a model to solve a given task while learning features that are non-discriminative across datasets. In HD²S we apply this strategy on multi-level features (unlike typical single-branch usage), in order to support the generalization of the saliency prediction task to datasets for which no annotations are available during training. While unsupervised domain adaptation has been applied to image classification (Ganin 2016; Tzeng et al. 2017), face recognition (Kan et al. 2015), object detection (Tang 2016), semantic segmentation (Zhang et al. 2020) and video action recognition (Li et al. 2018) (among others), our work is, to our knowledge, the first to deal with unsupervised domain adaptation on video saliency prediction. It is worthwhile to note that this is technically and fundamentally different from the form of domain adaptation proposed in UNISAL (Droste et al. 2020), that, instead, learns domain-specific parameters. This means that, at inference time, UNISAL requires to know the source dataset of a given input in order to select domain-specific learned parameters. Our approach, instead, is domain-agnostic as it employs the learned parameters on any tested domain. It is also different from unsupervised salient object detection (Zhang et al. 2018), which, instead, attempts to predict saliency by exploiting large unlabelled or weakly-labelled samples. However, we also provide HD²S with domain-specific learning capabilities as in (Droste et al. 2020), showing how this mechanism improves performance but cannot be applied in unsupervised domain adaptation scenarios.

3 Method

3.1 Architecture Overview

The proposed architecture is a fully-convolutional multi-branch encoder-decoder network for saliency prediction, illustrated in Fig. 2. An input sequence of consecutive video frames is first processed by a *feature extraction* path, which computes spatio-temporal features at different scales and abstraction levels. The extracted features serve as input to separate network branches that estimate a set of *conspicuity maps* at the corresponding points in the model, while at the same time providing skip paths to ease gradient flow during training. At the output of the model, conspicuity maps are combined to predict the saliency map for the last frame in the input sequence.

Our model is trained in a supervised way on a *source* dataset, for which saliency annotations are available.

Furthermore, the base model is provided with two additional mechanisms (that can be both disabled or enabled exclusively):

- *Domain Adaptation* modules that aim to make the model learn, in an unsupervised way, generalizable features (see red items in Fig. 2). In particular, each conspicuity sub-network forks to a *domain classification* path, that is trained to classify whether an input video sequence (more precisely, the corresponding features at that abstraction level) is taken from the source domain or from a *target* domain, which cannot be employed for training through direct supervision since annotations are not available. In order to perform this adaptation, we apply the gradient reversal technique: the feature extraction layer, shared by the conspicuity networks and the domain classifiers, is trained in an adversarial way, in order to force the model to learn features that are both discriminative and predictive—saliency-wise—as well as domain-invariant, in order to achieve satisfactory results even on the target domain.
- *Domain-Specific Learning* mechanism that learns specific parameters to enhance the prediction on a given dataset. More specifically, we add modules (shown as light gray items in Fig. 2), used in a multi-source training scenario (i.e., when using in training multiple datasets at the same time), whose parameters are optimized on each individual dataset. These modules aim to modulate features shared across multiple datasets based on the test data domain and include: domain-specific priors, batch-normalization and prediction smoothing.

At inference time, saliency maps are predicted for each frame by applying the model in a sliding window fashion, as in (Min and Corso 2019); the saliency map \mathbf{S}_t at time t is predicted from a sequence $\mathbf{V}_t = \{\mathbf{I}_{t-T+1}, \dots, \mathbf{I}_t\}$, where \mathbf{I}_t is the video frame at time t . To predict the first $T - 1$ frames, we reverse the chronological order of the corresponding input clips: each \mathbf{S}_t for $1 \leq t \leq T - 1$ is predicted from the sequence $\mathbf{V}_t = \{\mathbf{I}_{t+T-1}, \dots, \mathbf{I}_t\}$. As a final post-processing step, we apply a Gaussian filter ($\sigma = 5$) for smoothing the output saliency map.

In the following, we describe each of the components of our architecture.

3.2 Feature Extractor

The employed feature extractor performs spatio-temporal encoding of an input videoclip (16 frames of size 128×192), using S3D (Xie et al. 2018) as a backbone. It then progressively reduces the dimensions of the feature maps through 3D max pooling to $2 \times 4 \times 6$ (time \times height \times width), while increasing the number of channels to 1024. However, in order

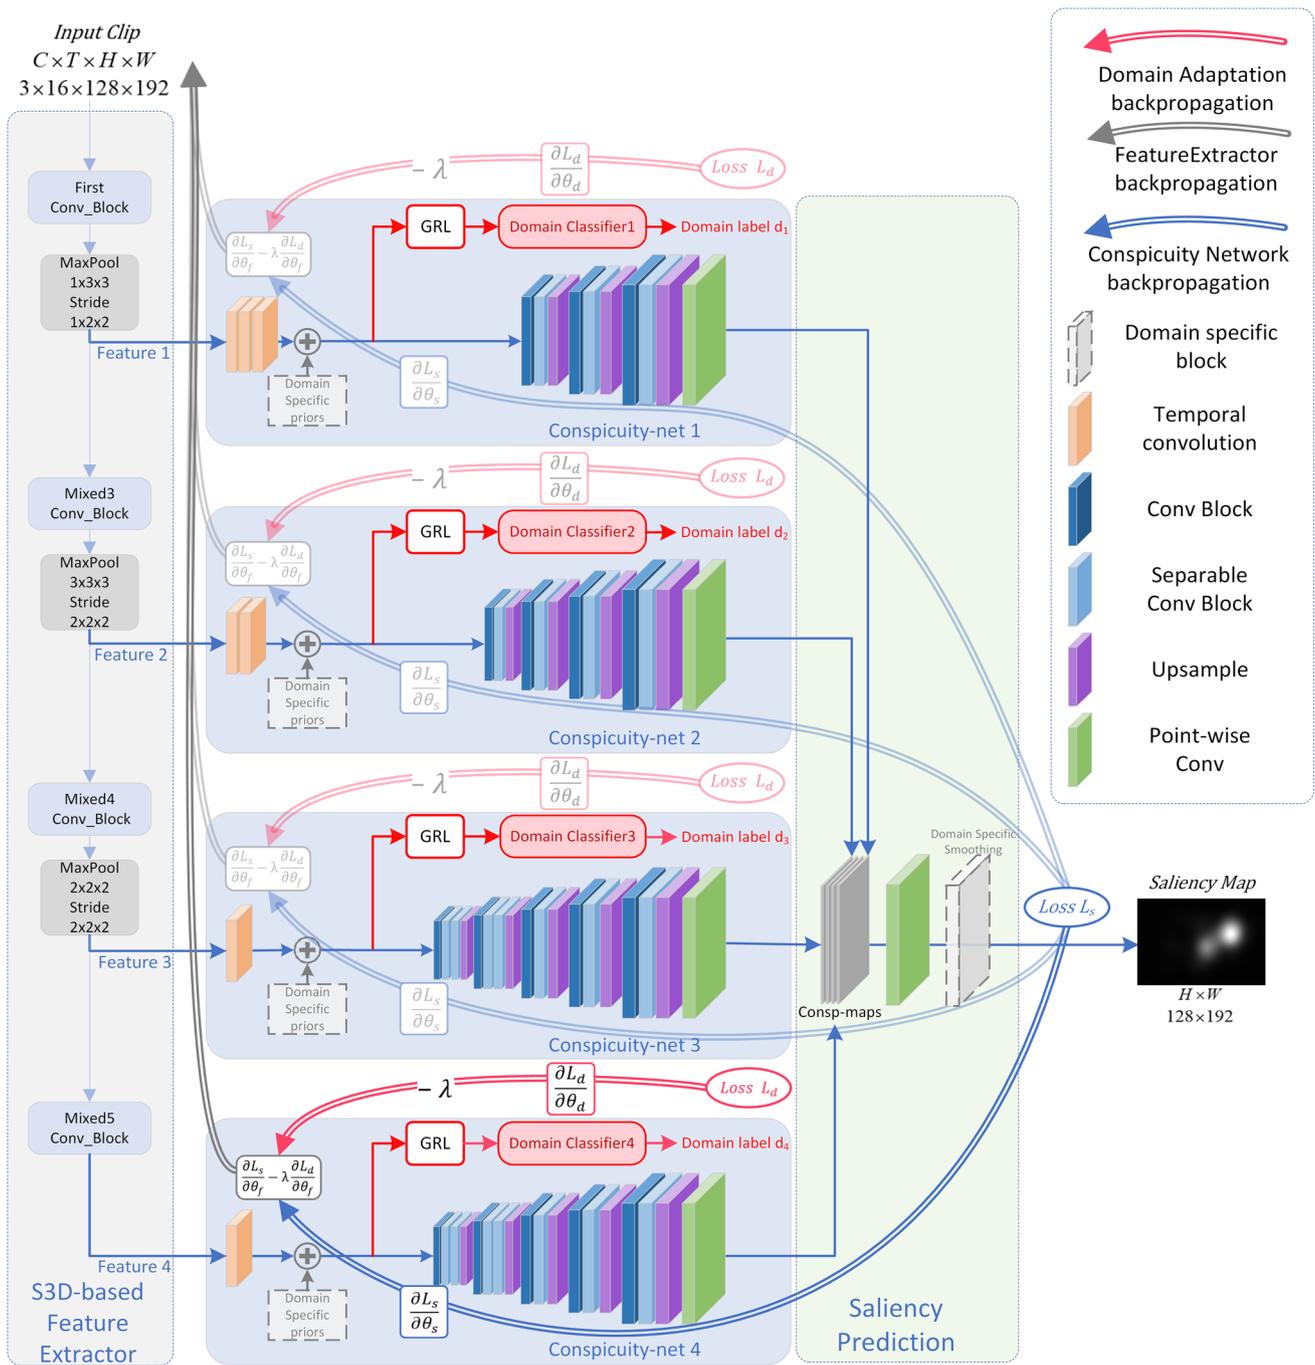


Fig. 2 HD_2S architecture: Our multi-branch decoder predicts four *conspicuity maps* at different feature abstraction levels, which are then integrated into the final saliency prediction, on which KL-divergence loss \mathcal{L}_s is minimized. As for *unsupervised domain adaptation*, each decoder branch is equipped with a gradient reversal layer (see red items) that encourages the model to learn features that generalize to a target

data domain in an unsupervised way, by maximizing the classification error \mathcal{L}_d on the prediction of an input sample’s domain. Finally, HD²S is also provided with *domain-specific* priors added to encoded features, with removed temporal dimension, and domain-specific smoothing as a last final layer

to exploit the full potential of the learned hierarchical representations, we select feature maps at different levels of the extractor, corresponding to different abstraction details, in

order to build a skip architecture able to capture multi-headed saliency responses. In our implementation, we select feature maps from the S3D backbone at the output of the second,

third and fourth pooling layers and at the input of the last average pooling layer.

3.3 Conspicuity Networks

After feature encoding, we learn several *conspicuity maps* from the partial information produced at different levels of the feature extraction stack through multiple decoder networks (referred as conspicuity networks in Fig.2).

Each conspicuity network in the model processes one of the spatio-temporal feature blocks coming from the feature extractor and returns a single-channel saliency map, encoding the conspicuity of spatial locations at that level of abstraction. In detail, the temporal dimension of the input feature maps is gradually removed, by applying a cascade of spatially point-wise convolutions (i.e., with kernel $3 \times 1 \times 1$ and stride $2 \times 1 \times 1$) that halve the temporal dimension at each step. The number of point-wise convolutions varies for each conspicuity network, depending on the size of the input feature maps.

After that, the (now purely spatial) set of feature maps is processed by a stack of 2D convolutional layers, interleaved with bilinear upsampling blocks, each of which doubles the spatial size of the feature maps until the original resolution of each frame is recovered.

3.4 Saliency Prediction

The four conspicuity maps produced by the above sub-networks are finally fused to predict saliency on the last frame of the input video. The global fusion layer consists of concatenating the four maps and performing pixel-wise 1×1 convolution followed by logistic activation.

At training time, the whole model (feature extractor, conspicuity networks and saliency predictor) is trained supervisedly on the source dataset in order to minimize the Kullback-Leibler (KL) divergence (Min and Corso 2019; Huang et al. 2015), between the predicted saliency map and conspicuity maps, and the correct target. More formally, given the predicted output saliency map S_t , the four conspicuity maps $C_{t,i}$ with $i = 1, 2, 3, 4$ and the ground-truth map G_t for a given target frame, all normalized over pixels appropriately, our *multi-level saliency loss* \mathcal{L}_s is computed as follows:

$$\mathcal{L}_s(S_t, C_t, G_t) = \sum_{j=1}^4 \sum_i G_{t,i} \log \frac{G_{t,i}}{C_{t,j,i}} + \sum_i G_{t,i} \log \frac{G_{t,i}}{S_{t,i}} \tag{1}$$

where index i iterates over all pixels, index j iterates over the four conspicuity maps, $G_{t,i}$, $S_{t,i}$ and $C_{t,i,j}$ are correspond-

ing pixels of, respectively, the ground truth map, the output saliency map and the j -th conspicuity map.

3.5 Domain Adaptation

In addition to training the model in a supervised way on the source domain, we also encourage the feature extractor to generalize over a target domain, without any supervision. Our unsupervised domain adaptation strategy relies on the *gradient reversal layer* (GRL) approach.

In particular, we integrate domain adaptation by inserting, in all of the conspicuity subnetworks, a branch with a gradient reversal layer and a domain classifier after the temporal-dimension removal layer (see Fig. 2). More formally and generally, given an input video clip V_t with associate binary domain label $d \in \{0, 1\}$ (source or target, respectively), we compute a set of associated domain classification losses $\{\mathcal{L}_{d,1}, \dots, \mathcal{L}_{d,4}\}$ from 4 domain classifiers attached to the conspicuity networks. If we indicate by \hat{d}_i the probability of the input being from the target domain estimated by the i -th classifier, the corresponding negative log-likelihood loss is defined as:

$$\mathcal{L}_{d,i}(d, \hat{d}_i) = -d \log \hat{d}_i - (1 - d) \log (1 - \hat{d}_i) \tag{2}$$

The overall domain classification loss is simply computed as the sum of the individual contributions, since the interaction between saliency prediction and domain adaptation is controlled by the λ hyperparameter in the gradient reversal layers. As a result, the comprehensive loss for model training with domain adaptation is the following:

$$\mathcal{L} = \mathcal{L}_s + \sum_{i=1}^4 \mathcal{L}_{d,i} \tag{3}$$

During training, we alternately pass a batch of videos from the source domain and a batch of videos from the target domain: on the former, we compute and backpropagate both the saliency prediction loss \mathcal{L}_s and the domain classification loss \mathcal{L}_d (with target $d = 0$); on the latter, we can only compute and backpropagate the domain classification loss \mathcal{L}_d (with target $d = 1$), since no saliency annotation is available on the target domain. Minimizing the domain classification loss has the effect to train the classifiers to better discriminate between the source and the target domains, while at the same time adversarially training the feature extractor (and the initial temporal-removal layers in the conspicuity networks) to produce features that confuse the classifier, and hence that are domain-independent.

Architecturally, each domain classifier consists of a stack of 1×1 spatial convolutions aimed at reducing the number of features, followed by fully-connected layers, the last of

which provides binary classification prediction of the input video's domain.

3.6 Domain-Specific Learning

In certain multi-source training scenarios (e.g., as done in (Droste et al. 2020)), one may assume that annotations are available for all employed datasets, thus enabling supervised training on all of them. When applying our saliency prediction model to this scenario, we provide it with *domain-specific operations* (Droste et al. 2020), which address the domain shift among different datasets. Unlike the unsupervised domain adaption setting, where we attempt to unsupervisedly learn features that generalize over multiple datasets, we here explicitly tailor learned features to the specific characteristics of each dataset.

In practice, we adopt a set of domain-specific techniques which have demonstrated to be effective as in (Droste et al. 2020):

Domain-Specific Priors. (Droste et al. 2020) thoroughly analyzed multiple video saliency benchmarks, identifying the sources of data shift among them and encoding these sources into a set of Gaussian prior maps. We employ the same strategy by initializing domain priors as in (Droste et al. 2020), and then letting the model learn the most suitable filters to weight the encoded spatio-temporal features depending on the input data domain. Domain priors are used to modulate the encoded features, after removing the temporal dimension (see light gray blocks in Fig. 2).

Domain-Specific Smoothing. The optimal way in which the output map should be smoothed varies between different datasets and depends mostly on how ground truth is created. To address this issue, we learn a different Gaussian kernel (i.e., with a different value of σ) for each input data domain. Unlike (Droste et al. 2020), our layer is parameterized by σ only, with convolution coefficients computed accordingly to make the filter Gaussian, while (Droste et al. 2020) initialize domain-specific convolutional filters to be Gaussian, but they may drift to non-Gaussian as the network updates its parameters. This smoothing is applied to the global saliency map (see Fig. 2).

Domain-Specific Batch Normalization aims at mitigating the impact of data distribution shift on the statistics estimated by batch normalization for inference, which may become inaccurate when computed over different benchmarks (Li et al. 2016; Chang et al. 2019; Droste et al. 2020). Thus, we learn batch normalization statistics for each dataset independently and accordingly apply them at inference time, depending on the input domain, as in (Droste et al. 2020).

4 Experimental Results

4.1 Datasets

This section describes briefly the datasets commonly employed for benchmarking video saliency prediction methods:

- *DHF1K* (Wang et al. 2018c) consists of 1,000 high-quality videos with a large diversity of scenes, objects, types of motion, complexity of backgrounds. In total, it includes 582,605 frames annotated with fixation points from 17 observers during a free-viewing experiment. The dataset is split into 600/100/300 videos for training, validation and test sets, respectively. The test set is not released and the results are maintained by the dataset curators.²
- *UCF Sports* (Marszalek et al. 2009) contains 150 videos taken from the UCF Sports Action Dataset (Soomro et al. 2014). Fixations are collected from 16 subjects while attempting to identify the action that occurred in the video. The dataset is split into 103 videos for training, and the remaining 47 for test, for a total of around 6,500 frames for training and 3,000 frames for test. The length of the videos varies between 20 and 140 frames.
- *Hollywood2* (Mathe and Sminchisescu 2014) contains 6,659 video sequences and derives, like UCF Sports, from a dataset for action recognition (Marszalek et al. 2009). The videos are collected from 69 Hollywood movies divided into 33 training movies and 36 test movies. Similarly to UCF Sports, the annotations are collected in a task-driven way. The videos are split into 3,100 clips for training and 3,559 clips for testing.
- *LEDOV* (Jiang et al. 2018) includes 538 videos of daily action, sports, social activity and art performance; we employ this dataset only as a target dataset for unsupervised domain adaptation.

Figure 3 provides statistics on the training splits of the datasets employed for training our model: (1) UCF Sports is the smallest one in terms of available videos and average number of frames per video, thus it seems to be unsuitable for models with high capacity as they likely overfit it; (2) Hollywood2 contains the highest number of videos but the majority has a very short number of frames (see the right histogram in Fig. 3), thus it may disadvantage methods that model temporal cues; (3) DHF1K is the most balanced in terms of videos and number of frames per videos.

² The DHF1K benchmark is available at <https://mmcheng.net/videosal/>.

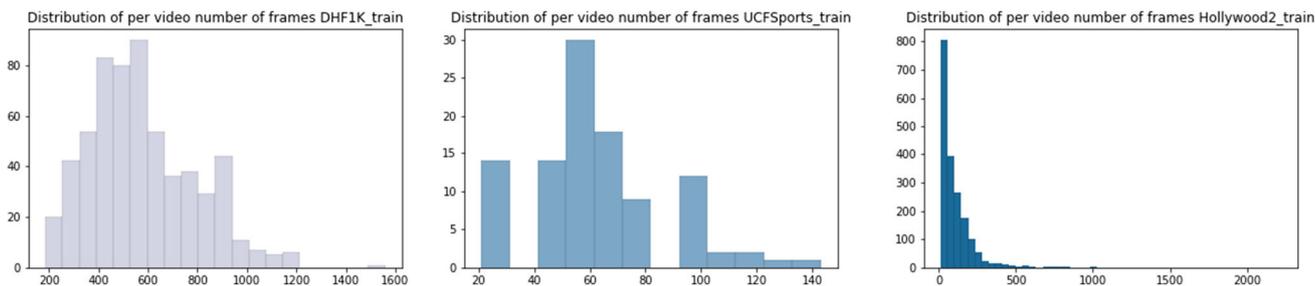


Fig. 3 Statistics of the training sets of DHF1K, Hollywood and UCF Sports

4.2 Training Procedure

In our experiments, we pre-train the S3D backbone on the Kinetics-400 (Kay et al. 2017) dataset; backbone parameters are not frozen, so they are updated during saliency prediction training. After empirically testing different hyperparameter configurations in order to find the best combination, the networks are trained for 2500 iterations, using Adam as optimizer (Kingma and Ba 2014) with learning rate of 10^{-3} . To reduce overfitting, L_2 regularization is applied, with a weight decay factor of 2×10^{-7} . The λ parameter of the gradient reversal layers during training gradually varies from 0 to 1:

$$\lambda = \frac{2}{1 + e^{-10 \cdot p}} - 1 \tag{4}$$

where p linearly goes from 0 to 1 according to the formula:

$$p = \frac{\text{current_iteration}}{\text{total_iterations}} \tag{5}$$

Gradually increasing λ also acts as an additional regularizer, since it prevents the model from focusing too much on the saliency prediction objective as training goes on. During training, sequences of $T = 16$ consecutive frames are randomly sampled from the dataset’s videos, and each frame is spatially resized to 128×192 . We employed a batch size of 200, although for memory limitations we forward batches of 8 samples at each time, which accumulating gradients and updating the model’s parameters every 25 such forward steps. When training with domain adaptation, we also forward a batch of samples from the source domain and one of samples from the target domain, and use them to update the domain classifier only.

To evaluate performance, we use each dataset training/test split when available, with 10% of the training data used as validation split. An exception is represented by DHF1K, since ground-truth annotations for the test set are not provided for blind assessment: in this case, when comparing to state-of-the-art methods (Table 1), we report the test accuracy as computed by the dataset curators; while for ablation study

(Tables 4 and 5) and domain adaptation analysis (Tables 7 and 9), we employ the original validation set as test set.

Validation results are used to perform model selection for inference on the test set. When evaluating test performance in single-dataset experiments, the training, validation and test sets all come from the same domain.

In domain adaptation experiments (with labeled source and unlabeled target datasets), training and validation splits are from the source domain (whose annotations can be used at training time), while the test set is from either an unseen portion of the target domain or from a different dataset altogether.

In multi-dataset experiments, we combine the training splits of DHF1K, UCF Sports and Hollywood2 datasets into a single training set; as validation set, we employ only DHF1K’s validation split (because of its better balance compared to the other datasets, as mentioned in Sect. 4.1), while inference is carried out on each dataset’s test split. In this setting, in order to support domain-specific learning and correctly update domain-specific modules, each training mini-batch contains videos from a dataset at a time, alternating between datasets to deal with different dataset sizes.

To compare the results obtained by the models, we use five commonly used evaluation metrics for video saliency prediction (Bylinskii 2018): Normalized Scanpath Saliency (NSS), Linear Correlation Coefficient (CC), Area under the Curve by Judd (AUC-J), Shuffled-AUC (s-AUC) and Similarity (SIM). Higher scores on each metric mean better performance.

4.3 Video Saliency Prediction Performance

We first test the performance of our base model (without any form of adaptation) in the supervised scenario on the DHF1K test benchmark, to evaluate its capabilities in the video saliency prediction task. We then integrate domain adaptation by means of GRL layers (as shown in Fig. 2), using the LEDOV dataset as a target domain, due to its wider subject variability than Hollywood2 and UCF Sports. Finally, we compute the performance of HD²S when using domain-specific learning, which is the form of adaptation that is most suitable with supervised learning settings and that can lever-

Table 1 Comparison of HD²S, with domain adaptation (HD^2S_{DA}) and with domain-specific learning (HD^2S_{DSL} using LEDOV as target dataset), with other state-of-the-art methods on the DHF1K test set

DHF1K	NSS	CC	SIM	AUC-J	s-AUC
GBVS (Harel et al. 2007)	1.474	0.283	0.186	0.828	0.554
STSCovNet (Bak 2017)	1.632	0.325	0.197	0.834	0.581
Deep Net (Pan et al. 2016)	1.775	0.331	0.201	0.855	0.592
SALICON (Huang et al. 2015)	1.901	0.327	0.232	0.857	0.590
OM-CNN (Jiang et al. 2017)	1.911	0.344	0.256	0.856	0.583
DVA (Wang et al. 2018e)	2.013	0.358	0.262	0.860	0.595
SalGAN (Pan et al. 2017)	2.043	0.370	0.262	0.866	0.709
ACLNet (Wang et al. 2018c)	2.354	0.434	0.315	0.890	0.601
SalEMA (Linardos et al. 2019)	2.574	0.449	0.466	0.890	0.667
STRA-Net (Lai 2019)	2.558	0.458	0.355	0.895	0.663
TASED-Net (Min and Corso 2019)	2.667	0.470	0.361	0.895	0.712
SalSAC (Wu et al. 2020)	2.673	0.479	0.357	0.896	0.697
UNISAL (Droste et al. 2020)	2.776	0.490	0.390	0.901	0.691
HD²S	2.781	0.497	0.406	0.901	0.699
HD²S_DA	2.709	0.491	0.381	0.902	0.709
HD²S_DSL	2.812	0.503	0.406	0.908	0.702

Our variant with domain-specific learning outperforms all state-of-the-art methods on three out of five metrics (NSS, CC, AUC-J), while ranking second-best on SIM and s-AUC

age all available annotated datasets (DHF1K, Hollywood2, UCF Sports).

Table 1 shows the performance of our approach compared to the state of the art. HD²S, without domain adaptation (referred to in Table 1 simply as HD²S), outperforms all state-of-the-art methods on three out of five metrics (NSS, AUC-J, CC) and ranks second-best on SIM and third-best on s-AUC. Note that this variant also outperforms UNISAL (Droste et al. 2020), which already employs domain-specific learning, on four out of five metrics. When we also enable domain-specific learning modules HD²S (HD^2S_{DSL}), performance (especially NSS, CC and AUC-J) increases sensibly, and it outperforms UNISAL on all metrics, demonstrating better representational and specialization capabilities. When using HD²S, with the hierarchical gradient reversal mechanism for domain adaptation (HD^2S_{DA}), performance slightly degrades as the model attempts to adapt the learned features to the target datasets (in this case, UCF Sports, Hollywood2 and LEDOV). However, remarkably, despite this adaption mechanism, the model yields performance comparable with state-of-the-art ones.

Comparing HD²S with TASED-Net (Min and Corso 2019), which also employs S3D (Xie et al. 2018) as backbone, it is possible to notice that our method (with and without adaptation) significantly outperforms TASED-Net in four out of five metrics using only half of the frames employed by TASED-Net (16 versus 32). TASED-Net slightly outperforms HD²S on s-AUC only, a metric that measures performance at the peripheral areas of the image, where a larger temporal context may allow to better capture

the motion of an object. The generally better performance obtained by our method w.r.t TASED-Net demonstrates the importance of hierarchical feature learning, with equal backbone features. While our model yields the highest video saliency performance on DHF1K, and performance comparable to the state of the art on Hollywood2, its performance on UCF Sports is lower than UNISAL (Droste et al. 2020) and SalSAC (Wu et al. 2020), as reported in Table 2. This is explained first with the smaller size of UCF Sports w.r.t. DHF1K and Hollywood2. Indeed, during training, although we use all three datasets, UCF Sports accounts to about 1% of the total number of training video frames (DHF1K: 62%, Hollywood2: 37%, UCF Sports 1%). This imbalance causes the model to overfit UCF Sports.

However, the suitability of Hollywood2 and UCF Sports for saliency prediction deserves a further discussion. Indeed, both datasets' saliency annotations are collected in task-driven experiments (i.e., action recognition) and, as such, human observers tend to mainly observe specific actions rather than focusing on the salient objects themselves, which defeats the very purpose of saliency prediction. An example is given in Fig. 4 where our model fails to match the ground truth: indeed, it focuses on the girl's face at the front (correctly, as it is the most salient area), but the ground truth mostly highlights the action of the person behind the girl. Furthermore, both datasets show a huge center bias (Droste et al. 2020) and have a rather limited variability of spatio-temporal features, especially Hollywood2, where the majority of video clips is very short in time. Analogously, UCF Sports is significantly smaller in terms of video frames, making it hard

Table 2 Comparison of HD²S and its variants (with domain adaptation: HD²S_{DA}; with domain-specific learning: HD²S_{DSL}) with other state-of-the-art methods on Hollywood2 and UCF Sports datasets

Method	NSS	CC	SIM	AUC-J	s-AUC
<i>Hollywood2</i>					
SALICON	2.013	0.452	0.321	0.856	0.711
DVA	2.459	0.482	0.372	0.886	0.727
ACLNet	3.086	0.623	0.542	0.913	0.757
SalEMA	3.186	0.613	0.487	0.919	0.708
STRA-Net	3.478	0.662	0.536	0.923	0.774
TASED-Net	3.302	0.646	0.507	0.918	0.768
SalSAC	3.356	0.670	0.529	0.931	0.712
UNISAL	3.901	0.673	0.542	0.934	0.795
HD²S	3.426	0.668	0.558	0.927	0.797
HD²S_DA	3.139	0.653	0.520	0.927	0.774
HD²S_SL	3.352	0.670	0.551	0.936	0.807
<i>UCF Sports</i>					
Method	NSS	CC	SIM	AUC-J	s-AUC
SALICON	1.838	0.375	0.304	0.848	0.738
DVA	2.311	0.439	0.339	0.872	0.725
ACLNet	2.567	0.510	0.406	0.897	0.744
SalEMA	2.638	0.544	0.431	0.906	0.740
STRA-Net	3.018	0.593	0.479	0.910	0.751
TASED-Net	2.920	0.582	0.469	0.899	0.752
SalSAC	3.523	0.671	0.534	0.926	0.806
UNISAL	3.381	0.644	0.523	0.918	0.775
HD²S	3.001	0.594	0.493	0.913	0.773
HD²S_DA	2.756	0.579	0.478	0.905	0.759
HD²S_DSL	3.114	0.604	0.507	0.904	0.768

In bold the best results, in italic the second-best results. In the domain adaptation scenario, we use, respectively, UCF Sports/Hollywood2 as source domain, DHF1K as target domain, and we test the performance on UCF Sports/Hollywood2

to train 3D convolutional models (or deep learning models in general). For all above reasons, we believe that both Hollywood2 and UCF Sports should not be used for saliency prediction.

Table 3 compares our model with state-of-the-art techniques in terms of processing times and model size. Reference values for compared approaches are from (Droste et al. 2020). UNISAL is the most resource-efficient approach (thanks to its MobileNetV2 (Sandler et al. 2018) backbone); our approach achieves average values on those metrics, while performing better than the others in terms of prediction accuracy, as shown in the previous section.

5 Ablation Studies

To validate the importance and effectiveness of the HD²S architectural design choices, we test some model variants

Table 3 Size (in MB) and processing time (in seconds) for the proposed model and state-of-the-art approaches

Model	Size (MB)	Runtime (s)
Deep Net (Pan et al. 2016)	103	0.080
SALICON (Huang et al. 2015)	117	0.500
DVA (Wang et al. 2018e)	96	0.100
SalGAN (Pan et al. 2017)	130	0.020
ACLNet (Wang et al. 2018c)	250	0.020
SalEMA (Linardos et al. 2019)	364	0.010
STRA-Net (Lai 2019)	641	0.020
TASED-Net (Min and Corso 2019)	82	0.060
UNISAL (Droste et al. 2020)	16	0.009
HD²S	116	0.027

Best values in bold

(without any domain adaptation or domain-specific learning) on the validation set of the DHF1K:

1. We first investigate the performance of our network, adding the different conspicuity nets one at a time;
2. We quantitatively and qualitatively evaluate the individual contribution of each conspicuity net, testing them in simple encoder-decoder architecture.

For the ablation study, we define as *Baseline* our network in a simple encoder-decoder configuration, i.e., without the intermediate conspicuity maps and multi-level loss. More specifically, in the baseline model, the feature extractor remains unchanged, but only the deepest decoder branch (*Conspicuity-net 4*) is used.

The model variants and their performance are reported in Table 4. The results show that: a) each conspicuity net makes its own contribution to improving the final performance; b) multi-level loss on conspicuity maps enhances saliency prediction too. Overall, these results clearly verify the effectiveness of all important design features in HD²S.

In our control experiments, we also evaluate the individual contribution of each conspicuity net by testing the performance of the model when the other decoder streams are ablated. For example, when testing the contribution of the first conspicuity map, we use only Feature 1 (see Fig. 2) from the encoder stream and the related decoder stream (*Conspicuity-net 1* in Fig. 2) and so on for the other conspicuity nets. Results, reported in Table 5, indicate that individually the third conspicuity net performs better than the others.

To further elucidate this behavior, Fig. 5 shows the (normalized) weights learned by the fusion layer of HD²S model when integrating the four conspicuity maps for final prediction on the DHF1K dataset. The obtained values confirm that in the *base* and *domain-adaptation settings*, all the four maps

Table 4 Comparison of various HD²S (without DA and DSL) configurations

	NSS	CC	SIM	AUC-J	s-AUC
Consp-net4(Baseline)	2.602	0.468	0.355	0.902	0.697
+ Consp-net3	2.612	0.474	0.373	0.897	0.706
+ Consp-net2	2.699	0.482	0.374	0.901	0.706
+ Consp-net1	2.743	0.491	0.378	0.904	0.704
+multi loss (HD ² S)	2.806	0.489	0.403	0.904	0.705

The Consp-net4 configuration refers to the network in a simple encoder-decoder configuration, i.e., with Conspicuity-net 4 only. The full model includes all four conspicuity networks with multi-level loss on conspicuity and saliency maps, defined in Eq. 3

Table 5 Individual contribution of each conspicuity net

	NSS	CC	SIM	AUC-J	s-AUC
Only Consp-net1	2.191	0.392	0.301	0.871	0.689
Only Consp-net2	2.605	0.461	0.359	0.893	0.690
Only Consp-net3	2.663	0.480	0.359	0.902	0.697
Only Consp-net4	2.602	0.468	0.355	0.902	0.697
All Consp-nets	2.743	0.491	0.378	0.904	0.704
HD²S	2.806	0.489	0.403	0.904	0.705

Each configuration refers to a simple encoder-decoder architecture with different sets of encoded features. With *All conspicuity-nets*, we refer to the model with all decoding branches but without multi-loss, while *HD²S* is the final model including multi-loss



Fig. 4 An example of failure, taken from Hollywood2. Despite a good prediction, HD²S misses to match the ground truth, as it is collected in a task-driven experiment (action recognition), thus highlighting more actions than salient objects

contribute almost equally, with Map 3 slightly more. In the *domain-specialization* setting, instead, the fourth (deepest) map has a larger weight: this may indicate that domain-specialized blocks focus mostly on higher-level features to achieve good performance.

A qualitative interpretation of this behavior and on the contribution of each conspicuity map in the hierarchy is shown in Fig. 6. When comparing the behaviour of the different decoder branches on the standard, domain adaptation, and domain-specific learning regimes, the following considerations can be drawn: in standard training case (top line in Fig. 6), Map 4 contains similar information to that of Map 3; in the domain adaptation scenario (middle line in Fig. 6), all feature maps appear to contribute equally; in the domain specific learning case (bottom line in Fig. 6),

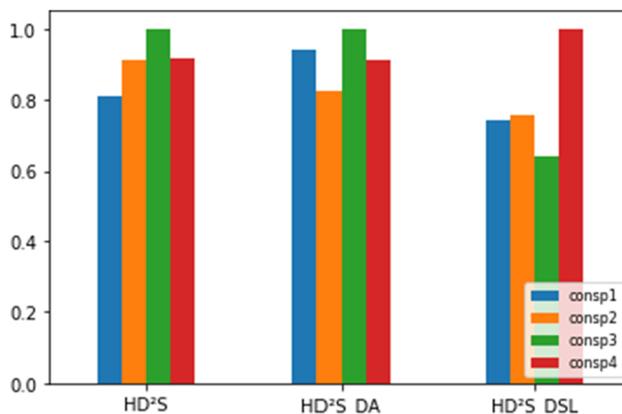


Fig. 5 Normalized weights learned by the fusion layer when integrating the four conspicuity maps on DHF1K dataset: Full HD²S model (*left block*), HD²S model with domain adaptation (*middle block*), HD²S model with domain specific learning (*right block*). For the HD²S model with domain adaptation, we use LEDOV as a target dataset

Map 4 provides additional (motion) information to Map 3. This provides an interpretation to the parameters learned by the fusion layers, reported in Fig. 5. Analyzing the intermediate maps in the domain specific learning (bottom line in Fig. 6), we can observe that the four intermediate maps encode saliency at different levels of detail: Map 1 extracts small background motion, Map 2 focuses mainly on the bull, Map 3 starts highlighting the bullfighter and, finally, Map 4 puts more emphasis on the bullfighter. Furthermore, we can also observe that, in the domain adaptation scenario, predicted salient regions are larger than in the other two settings (no adaptation and domain-specific learning). This is due to the domain adaptation strategy that enforces the model to make less crisp estimations in its attempt to learn common features between datasets; in the domain-specific learning scenario, instead, the model specializes its parameters to the characteristics of each dataset, thus learning more specific features for matching the ground truth more precisely. The same happens, of course, in the no-adaptation scenario, where the model can focus on a single training dataset.

We also quantify the level of similarity among the different conspicuity maps by computing the KL divergence among all pairs of maps over the entire DHF1K test set, and results are reported in Table 6. In particular, Map 1 is the one encoding the most different information w.r.t. to the other maps, while Map 3 and Map 4 encode similar cues (as also shown in Fig. 6), although Map 4 contains additional saliency information (indeed, KL between Map 4 and Map 3 is higher than the one between Map 3 and Map 4), possibly encoding more details about object motion. However, the highest gap among consecutive maps is in the transition from Map 2 to Map 3. Indeed, Map 3 includes most of the information available in Map 2 (KL div between Map 2 and Map 3 is

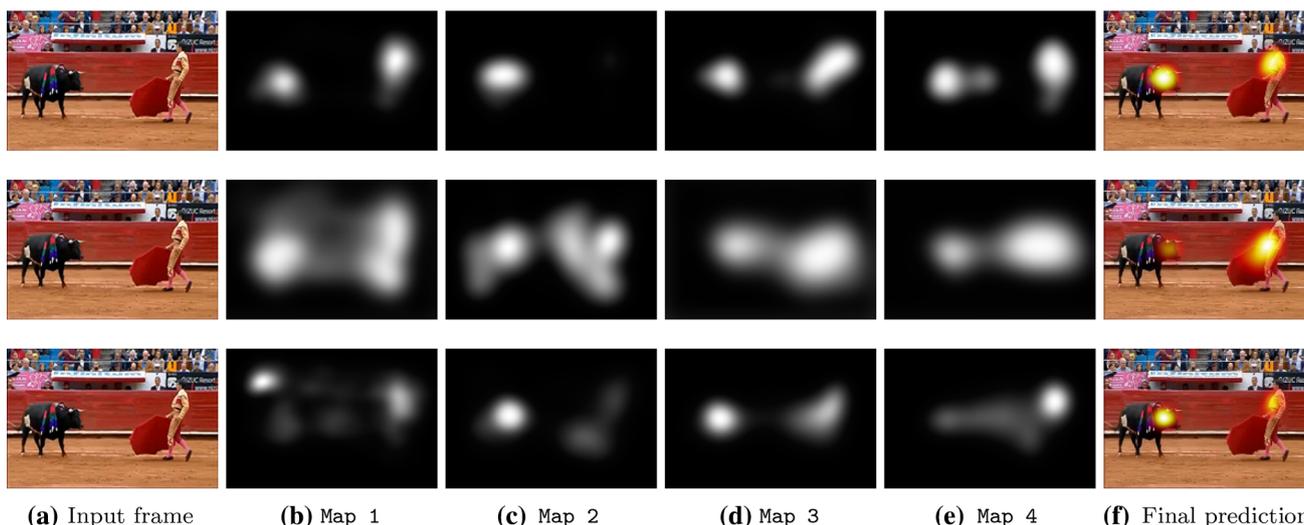


Fig. 6 Qualitative interpretation of the contribution of hierarchical decoding used under different settings. HD²S (*Top line*), HD²S with domain adaptation (*Middle line*) and HD²S with domain-specific learning (*Bottom Line*)

Table 6 KL divergence among all pairs of conspicuity maps

Map \	Map 1	Map 2	Map 3	Map 4
Map 1	0	1.110	1.017	1.015
Map 2	1.090	0	0.772	0.806
Map 3	1.531	1.253	0	0.390
Map 4	1.974	1.545	0.531	0

0.772), but complements it with additional visual details that tend to appear also in Map 4 (see again Fig. 6).

6 Domain Adaptation Performance

When testing domain adaption performance, we distinguish two cases: a) the capabilities of the model to address domain shift issues, i.e., the case of reducing the shift between training and test data; and b) the capabilities of the model to learn generalizable features that can be employed, without any additional tuning.

Domain-Shift. To assess the performance of our hierarchical domain adaptation approach in tackling the problem of domain shift, we run a set of experiments by selecting different combinations of datasets to be employed as source domain (used in a supervised way during training) and target domain, used in an unsupervised way during training; as test set, an unseen portion from the target domain is used. The assumption in these experiments is to perform unsupervised learning on the test domain through our hierarchical gradient reversal approach before testing on it (on the appropriate test split not used for unsupervised learning).

In particular, we compare the performance of our base model in the three scenarios:

- *No Domain Adaptation*, how well the model generalizes to new domains, i.e., the model trained supervisedly on the source domain and directly tested on the target domain, with no additional information on the test dataset used during training;
- *Domain Adaptation*, i.e., the model trained with unsupervised adaptation on the target domain, enabled through the hierarchy of GRL layers as in our full model in Fig. 2;
- *Transfer Learning*, i.e., the model (with gradient reversal disabled) trained on the source dataset and then fine-tuned (in a supervised way) on the target dataset. This scenario represents the upper bound of the evaluation and is, of course, out of the scope of pure domain adaptation, since target domain labels are available at training time.

Table 7 shows the results for different combinations of source and target domains. Two main patterns of results can be identified, depending on whether DHF1K is employed as source domain or not. In the former case (top block of Table 7), it can be noticed that the employment of gradient reversal layers improves performance over all target datasets, compared to simply training on the source dataset. When instead DHF1K is employed as target domain (second and third blocks in Table 7), the use of gradient reversal layers degrades performance. This may be due to the specific characteristics of Hollywood2 and UCF Sports, which were collected in a task-driven experiments while DHF1K in a free-viewing one: as a consequence, training on Hollywood2/UCF Sports encourages the model to focus on visual features that are more indicative of the action being per-

Table 7 Analysis of domain-shift capabilities

Target dataset	Approach	NSS	CC	SIM	AUC-J	s-AUC
<i>Source dataset: DHF1K</i>						
UCF sports	No adaptation	2.483	0.537	0.442	0.890	0.744
	Domain adaptation	2.514	0.539	0.448	0.893	0.750
	Transfer learning	3.001	0.594	0.493	0.913	0.773
Hollywood2	No Adaptation	3.063	0.625	0.505	0.925	0.779
	Domain adaptation	3.101	0.632	0.510	0.925	0.785
	Transfer learning	3.426	0.668	0.558	0.927	0.797
<i>Source dataset: UCF Sports</i>						
DHF1K	No adaptation	2.237	0.405	0.325	0.880	0.658
	Domain adaptation	2.160	0.392	0.309	0.877	0.658
	Transfer learning	2.688	0.477	0.374	0.896	0.700
Hollywood2	No Adaptation	2.469	0.517	0.433	0.899	0.727
	Domain adaptation	2.386	0.503	0.422	0.896	0.727
	Transfer learning	3.298	0.657	0.533	0.925	0.794
<i>Source dataset: Hollywood2</i>						
DHF1K	No adaptation	2.467	0.445	0.338	0.893	0.690
	Domain adaptation	2.461	0.447	0.338	0.894	0.696
	Transfer learning	2.753	0.487	0.384	0.898	0.697
UCF sports	No Adaptation	2.476	0.538	0.442	0.885	0.756
	Domain adaptation	2.389	0.522	0.431	0.882	0.746
	Transfer learning	2.780	0.576	0.486	0.887	0.762

Performance evaluation in the *no adaptation* (supervised training on source; testing on target) and *domain adaptation* (supervised training on source; unsupervised training and test on target) scenarios. Upper-bound performance is measured by the *transfer learning* scenario (supervised training on source and fine-tuning on target)

formed than of visual saliency *per se*, and that may not correspond to visual saliency cues highlighted in DHF1K. Furthermore, the limited variability of spatio-temporal features from videos in Hollywood2, as shown in Fig. 3, makes it harder for the model to move clustered features and to learn more general representations. Similarly, when UCF Sports is used as source domain, the small size of the dataset makes it easier for the model to focus on the supervised saliency prediction task (on which it can easily achieve a low training loss), rather than minimizing the domain adaptation loss. Overall, as expected, the highest performance are obtained in the transfer learning regime.

Learning Generalizable Features. We also test the capabilities of the model to learn general features by using, in the domain adaptation stream, a target dataset different from the one used for test. We specifically compute performance when training on DHF1K, adapting the learned features to LEDOV, and testing on never seen datasets (UCF Sports and Hollywood2). Performance are reported in Table 8, which reports how the performance gain of HD²S, when empowered with hierarchical gradient reversal modules, is higher than in the case of domain-shift experiments (see Table 7). This demonstrates that our hierarchical domain adaptation mechanism

Table 8 Analysis of generalization capabilities

Test	Setting	NSS	CC	SIM	AUC-J
Source: DHF1K - target: LEDOV					
UCF Sports	No Adapt.	2.494	0.536	0.442	0.889
	Adaptation	2.584	0.555	0.452	0.900
Hollywood2	No Adapt.	3.011	0.622	0.502	0.922
	Adaptation	3.066	0.623	0.505	0.926

Domain adaptation performance, respectively, with source: DHF1K, target: LEDOV; test: UCF Sports and Hollywood2. Best results in bold

is better at learning saliency features that generalize well on multiple data domains than at addressing the domain-shift for a given dataset.

6.1 Multi-source Training

A recent trend in video saliency prediction (Droste et al. 2020) proposes *multi-source training* as a means for improving performance by leveraging the larger input variability of multiple data sources. This setup also allows for the integration of *domain-specific learning* capabilities, as mentioned in Sect. 3.6, that attempt to tune general features to spe-

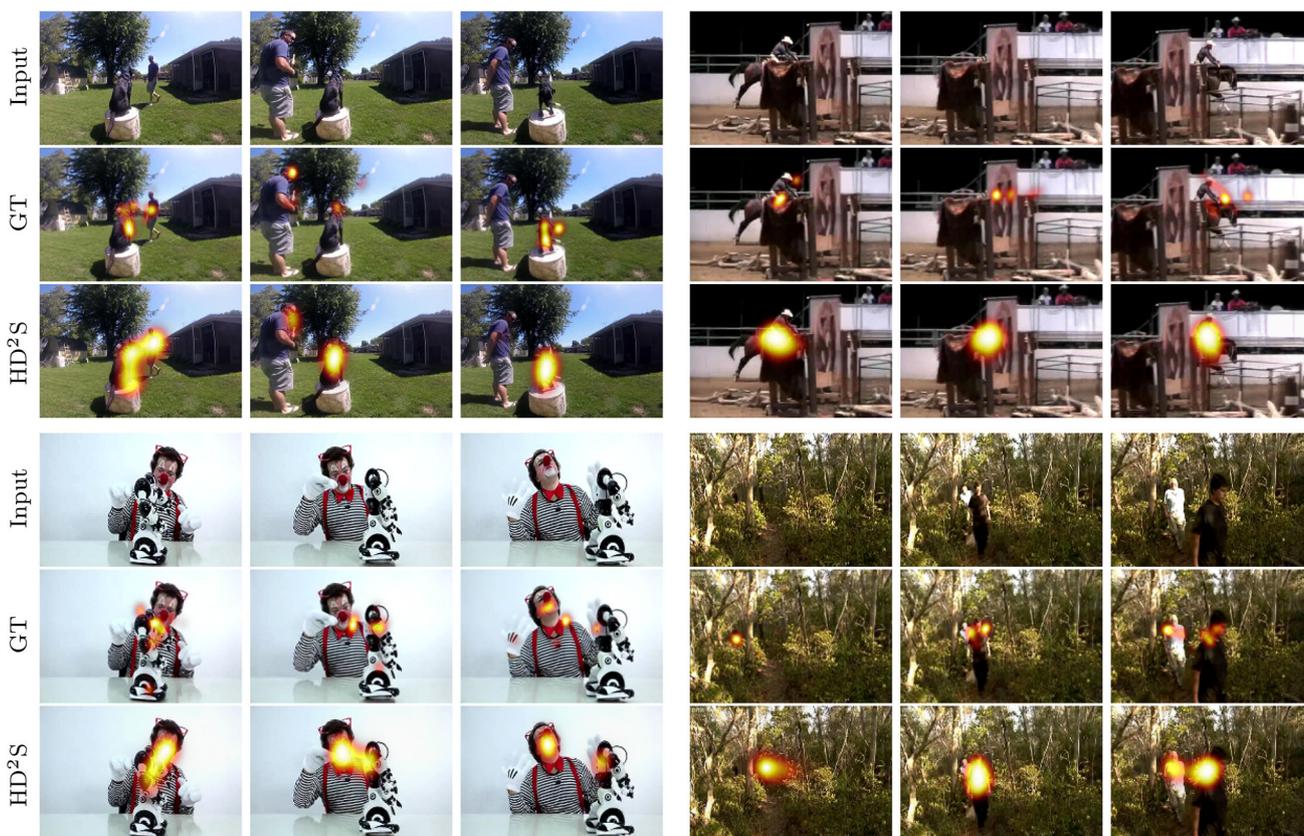


Fig. 7 Qualitative evaluation of the proposed model on the DHF1K validation set. Comparison of the saliency predicted by our model with the ground truth on some frames: saliency with multiple objects (left

block), saliency on an occluded object (upper-right block), saliency on moving objects whose appearance changes rapidly among consecutive frames (lower-right block)

sific datasets. The idea is to have a model that learns shared features across multiple datasets and then to employ domain-specific modules to adapt such features to a particular data domain. Although these domain-specific approaches do not strictly comply with the standard unsupervised domain adaptation formulation, as they go in the exact opposite direction to learning generalizable features (since they assume that target domain labels are available at training time), it is interesting to evaluate the impact of domain-specific learning on our architecture. In Sect. 4.3 and Table 1, we already showed that the integration of domain-specific capabilities into the HD²S model achieves state-of-the-art performance on DHF1K, outperforming (Droste et al. 2020), that introduced those techniques. Here, we complete our analysis by assessing the impact of domain-specific layers compared to multi-source domain learning. More specifically, for multi-source domain learning, we use the integration of DHF1K, Hollywood2 and UCF Sports, as an unified dataset, for training and testing our model. As for domain specific learning, we enable the domain-specific modules (described in Sect. 3.6) and train their parameters using data from each individual dataset and during inference we provide, as an additional

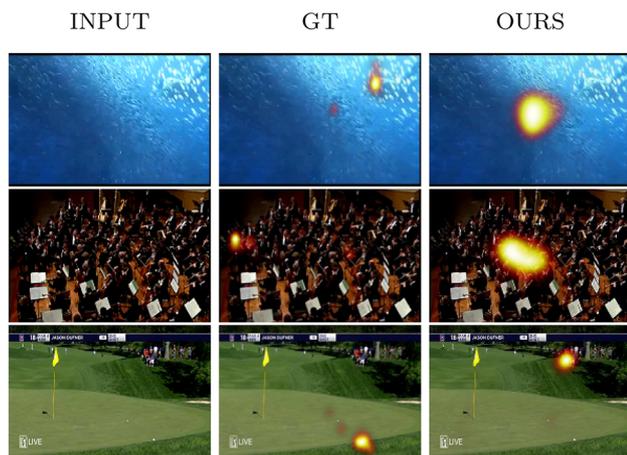
input to the model, the dataset we want to test it. We also compute performance when using single-source domain, i.e., training and test on a single dataset at a time. The results in Table 9 confirm that multi-source training by itself does not provide a much larger boost compared to single-source analysis, while domain-specific learning of dataset characteristics significantly improves performance, confirming that saliency prediction models surely benefit from embedding domain-specific layers from multiple datasets at training time.

7 Qualitative Analysis

We here report quantitative analysis of the results obtained by our model. Figure 7 shows examples of saliency predictions made by our HD²S model with domain-specific learning on the DHF1K benchmark. The model is able to effectively face object occlusion, multiple objects, fast motion, strong camera motion, stationary objects, saliency shift, camera focus change, low-light condition. Sample videos of how our model works are also given in the GitHub page of the project. Figure 8, instead, shows example of failures that typically

Table 9 Performance evaluation on the multi-source and domain specific learning scenarios

Test Dataset	Approach	NSS	CC	SIM	AUC-J	s-AUC
Train datasets: DHF1K, UCF Sports, Hollywood2						
DHF1K	Single-source	2.806	0.489	0.403	0.904	0.705
	Multi-source	2.811	0.491	0.403	0.893	0.708
	Domain-specific	2.875	0.500	0.406	0.910	0.707
UCF Sports	Single-source	2.803	0.589	0.489	0.879	0.759
	Multi-source	2.922	0.594	0.498	0.882	0.767
	Domain-specific	3.114	0.604	0.507	0.904	0.768
Hollywood2	Single-source	3.235	0.660	0.528	0.919	0.778
	Multi-source	3.349	0.665	0.551	0.922	0.797
	Domain-specific	3.352	0.670	0.551	0.936	0.807

**Fig. 8** Examples of failures. The model struggles with small objects and small motion: in the first two cases, the model missed the salient region and highlights a generic prior; in the third example, the model does not manage to identify the golf ball, focusing instead on a man in a red shirt, standing out from the surroundings

happen in case of small global motion or small objects. These failures can be caused by the spatial resolution at which input images are scaled before being processed by the model (128×192). Indeed, in the first two cases of Fig. 8 the models is unable to identify the correct salient region (located in a lateral region of the scene), and instead predicts a generic prior-driven center region. In the last case, the model fails to detect the movement of a golf ball towards the hole (a slow movement of a small object), and erroneously predicts as salient the upper-right region of the scene, where a man with a red shirt significantly stands out from the surroundings.

8 Conclusion

In this work, we propose HD²S, a new fully-convolutional network for video saliency prediction. The key architectural elements of our proposed approach include a multi-branch

decoder which acts at different feature abstraction layers to independently estimate *conspicuity maps*, which are then combined into the final prediction, and an unsupervised domain adaptation mechanism that enables our model to learn features that, at the same time, allow it to reach state-of-the-art performance on supervised saliency prediction, while generalizing to domains for which no annotations are provided at training time. Additionally, when employing domain-specific learning techniques, as introduced in (Droste et al. 2020), our model's performance on the supervised saliency prediction task further improves.

Comparing our approach with state-of-the-art models, we find that our late-fusion mechanism of multi-level saliency features provides a significant boost to performance: our ablation studies show that the gradual integration of multiple abstraction levels positively affects prediction accuracy. This is also confirmed by analyzing the learned weights. Interestingly, the impact of each conspicuity map (and, therefore, of each abstraction level of learned features) seems to vary depending on the employed domain adaptation mechanism: high-level features become predominant when domain-specific learning is applied (possibly due to the larger data distribution variability introduced by multi-source training, which causes shallower features to generalize less), while all conspicuity maps become similarly important when unsupervised domain adaptation is applied, which can be explained through the action of gradient reversal layers, which actively encourage features to become domain-independent and thus to be equally effective at multiple scales.

While the model performs well in several complex cases—e.g., in presence of multiple objects, occlusions and appearance changes—there are certain conditions in which we find room for improvement. Most of them involve the presence of small objects and small motion, where the model fails to correctly locate areas of interest, and the prediction is dominated by the prior.

These situations could benefit from working at higher resolution, given sufficient computing resources, or in a

patch-based fashion, to the detriment of inference times. However, major failures seem to be related to the specific characteristics of datasets: Hollywood2 and UCF Sports, for instance, are annotated with task-driven gaze fixations, rather than free-view the scene. This, of course, negatively affects methods that instead attempt to predict bottom-up saliency. Improved dataset availability curation for video saliency prediction may be an enabling factor for the advancement in the field.

Funding Open access funding provided by Università degli Studi di Catania within the CRUI-CARE Agreement.

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

References

- Badrinarayanan, V., Kendall, A., & Cipolla, R. (2017). Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE TPAMI*, 39(12), 2481–2495.
- Bak, C., et al. (2017). Spatio-temporal saliency networks for dynamic saliency prediction. *IEEE TMM*, 20(7), 1688–1698.
- Bazzani, L., Laroche, H., Torresani L. (2016). Recurrent mixture density network for spatiotemporal visual attention . In: arXiv preprint [arXiv:1603.08199](https://arxiv.org/abs/1603.08199) (2016).
- Borji, A., Itti, L. (2015). Cat2000: A large scale fixation dataset for boosting saliency research . In: arXiv preprint [arXiv:1505.03581](https://arxiv.org/abs/1505.03581)
- Bylinskii, Z., et al. (2018). What do different evaluation metrics tell us about saliency models? *IEEE TPAMI*, 41(3), 740–757.
- Chang W.-G. et al. (2019). Domain-specific batch normalization for unsupervised domain adaptation . In: Proceedings of the IEEE/CVF CVPR. pp. 7354–7362.
- Che, Z., et al. (2019). How is gaze influenced by image transformations? dataset and model. *IEEE TIP*, 29, 2287–2300.
- Yangyu C. et al. (2018). “Saliency-based spatiotemporal attention for video captioning”. In: 2018 IEEE BigMM. IEEE. pp. 1–8.
- Cornia M. et al. (2016). A deep multi-level network for saliency prediction . In: ICPR. IEEE. pp. 3488–3493.
- Cornia, M., et al. (2018). Predicting human eye fixations via an LSTM-based saliency attentive model. *IEEE TIP*, 27(10), 5142–5154.
- Dosovitskiy, A. et al. (2015). Flownet: Learning optical flow with convolutional networks . In: ICCV. pp. 2758–2766.
- Droste, R., Jiao, J., & Alison, N. J. (2020). Unified image and video saliency modeling. *ECCV* (pp. 419–435). Berlin: Springer.
- Fan, S. et al. (2018). Emotional attention: A study of image sentiment and visual attention . In: Proceedings of the IEEE CVPR. pp. 7521–7531.
- Ganin, Y., et al. (2016). Domain-adversarial training of neural networks. *JMLR*, 17(1), 2096–2030.
- Girshick, S. (2015). Fast R-CNN . In: Proceedings of the IEEE ICCV.
- Goodfellow, I. et al. (2014). Generative adversarial networks . In: arXiv preprint [arXiv:1406.2661](https://arxiv.org/abs/1406.2661).
- Guraya, F.F.E., et al. (2010). Predictive saliency maps for surveillance videos. In: DCABES. IEEE. pp. 508–513.
- Harel, J., Koch, C., Perona, P. (2007). Graph-based visual saliency . In: NIPS. pp. 545–552.
- He, K., et al. (2020). Mask R-CNN. In: *IEEE TPAMI*, 42(2), 386–397.
- Hou, Q., et al. (2019). Deeply supervised salient object detection with short connections. *IEEE TPAMI*, 41(4), 815–828.
- Huang, X., et al. (2015). Salicon: Reducing the semantic gap in saliency prediction by adapting deep neural networks . In: ICCV. pp. 262–270.
- Itti, L., Koch, C., & Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE TPAMI*, 20(11), 1254–1259.
- Jia, S., Bruce Neil, D.B. (2020). Eml-net: An expandable multi-layer network for saliency prediction . In: Image and vision computing. vol. 95, p. 103887.
- Jiang, L., Xu, M., Wang, Z. (2017). Predicting video saliency with object-to-motion CNN and two-layer convolutional LSTM . In: arXiv preprint [arXiv:1709.06316](https://arxiv.org/abs/1709.06316).
- Jiang, L., et al. (2018). Deepvs: A deep learning based video saliency prediction approach . In: ECCV. pp. 602–617.
- Jiang, M., et al. (2015). Salicon: Saliency in context . In: Proceedings of the IEEE CVPR. pp. 1072–1080.
- Judd, T., Durand, F., Torralba, A. (2012). A benchmark of computational models of saliency to predict human fixations.
- Kan, M., Shan, S., Chen, X. (2015). Bi- Shifting auto-encoder for unsupervised domain adaptation . In: ICCV.
- Kay, W., et al. (2017). The kinetics human action video dataset . In: arXiv preprint [arXiv:1705.06950](https://arxiv.org/abs/1705.06950).
- Kingma Diederik, P., Ba, J. (2014). Adam: A method for stochastic optimization . In: arXiv preprint [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- Kroner, A., et al. (2020). Contextual encoder-decoder network for visual saliency prediction. *Neural Networks*, 129, 261–270.
- Kummerer, M., et al. (2017). Understanding low and high-level contributions to fixation prediction . In: Proceedings of the IEEE ICCV.
- Lai, Q., et al. (2019). Video saliency prediction using spatiotemporal residual attentive networks. *IEEE TIP*, 29, 1113–1126.
- Li, J., et al. (2018). Unsupervised learning of viewinvariant action representations . In: NIPS. pp. 1254–1264.
- Li, S., Lee M.C. (2007). Fast visual tracking using motion saliency in video . In: ICASSP. IEEE. Vol. 1, pp. I–1073.
- Li, Y., et al. (2016). Revisiting batch normalization for practical domain adaptation . In: arXiv preprint [arXiv:1603.04779](https://arxiv.org/abs/1603.04779).
- Lim, M.K., et al. (2014). Crowd saliency detection via global similarity structure . In: ICPR. IEEE. pp. 3957–3962.
- Linardos, P., et al. (2019). Simple vs complex temporal recurrences for video saliency prediction . In: arXiv preprint [arXiv:1907.01869](https://arxiv.org/abs/1907.01869).
- Liu, T., et al. (2010). Learning to detect a salient object. *IEEE TPAMI*, 33(2), 353–367.
- Long, M., et al. (2015). Learning transferable features with deep adaptation networks . In: ICML. PMLR. pp. 97–105.
- Lu, L., et al. (2017). Crowd behavior understanding through SIOF feature analysis . In: ICAC. IEEE. pp. 1–6.
- Marszalek, M., Laptev, I., Schmid C. (2009). Actions in context . In: CVPR. IEEE. pp. 2929–2936.
- Mathe, S., & Sminchisescu, C. (2014). Actions in the eye: Dynamic gaze datasets and learnt saliency models for visual recognition. *IEEE TPAMI*, 37(7), 1408–1424.
- Min, K., Corso, J.J. (2019). TASED-Net: Temporally- aggregating spatial encoder-decoder network for video saliency detection . In: ICCV. pp. 2394–2403.
- Nguyen, T.V., et al. (2013). Static saliency versus dynamic saliency: A comparative study . In: ACM MM. pp. 987–996.

- Noh, H., Hong, S., Han, B. (2015). Learning deconvolution network for semantic segmentation . In: ICCV. pp. 1520–1528.
- Pan, J., et al. (2017). Salgan: Visual saliency prediction with generative adversarial networks . In: arXiv preprint [arXiv:1701.01081](https://arxiv.org/abs/1701.01081).
- Pan, J., et al. (2016). Shallow and deep convolutional networks for saliency prediction . In: CVPR. pp. 598–606.
- Pan S.J., Yang Q. (2009). A survey on transfer learning. In: IEEE TKDE 22.10, pp. 1345–1359.
- Redmon, J., et al. (2016). You only look once: Unified, real-time object detection . In: CVPR. pp. 779–788.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. *MICCAI* (pp. 234–241). Berlin: Springer.
- Sandler, M., et al. (2018). Mobilenetv2: Inverted residuals and linear bottlenecks . In: CVPR. pp. 4510–4520.
- Mark, S., et al. (2018). Mobilenetv2: Inverted residuals and linear bottlenecks. In: CVPR. pp. 4510–4520.
- Shao, J., Zhou, S.K., Chellappa, R. (2005). Tracking algorithm using background- foreground motion models and multiple cues [surveillance video applications]. In Proceedings (ICASSP'05) IEEE International conference on acoustics, speech, and signal processing, Vol. 2, pp. ii–233.
- Shokri, M., Harati, A., Taba, K. (2020). Saliency object detection in video using deep nonlocal neural networks . In: JVCIR vol. 68, p. 102769.
- Soomro, K., & Zamir, A. R. (2014). Action recognition in realistic sports videos. *Computer vision in sports* (pp. 181–208). Berlin: Springer.
- Sun, B., & Saenko, K. (2016). Deep coral: Correlation alignment for deep domain adaptation. *ECCV* (pp. 443–450). Berlin: Springer.
- Sun, M., et al. (2018). SG-FCN: A motion and memorybased deep learning model for video saliency detection. *IEEE Transactions on Cybernetics*, 49(8), 2900–2911.
- Tang, Y., et al. (2016). Large scale semi-supervised object detection using visual and semantic knowledge transfer. In: CVPR. pp. 2119–2128.
- Tran, D., et al. (2015). Learning spatiotemporal features with 3d convolutional networks . In: ICCV. pp. 4489–4497.
- Tzeng, E., et al. (2017). Adversarial discriminative domain adaptation . In: CVPR. pp. 7167–7176.
- Wang, H., Xu, Y., Han, Y. (2018). Spotting and aggregating salient regions for video captioning. In: ACM MM. pp. 1519–1526.
- Wang, M., & Deng, W. (2018). Deep visual domain adaptation: A survey. *Neurocomputing*, 312, 135–153.
- Wang, W., Shen, J., & Shao, L. (2017). Video salient object detection via fully convolutional networks. *IEEE TIP*, 27(1), 38–49.
- Wang, W., et al. (2019). Revisiting video saliency prediction in the deep learning era. *IEEE TPAMI*, 43(1), 220–237.
- Wang, W., et al. (2018). Revisiting video saliency: A large-scale benchmark and a new model . In: CVPR, pp. 4894–4903.
- Wang, X., et al. (2018). Non-local neural networks . In: CVPR, pp. 7794–7803.
- Wang, J., Shen, W. (2018). Deep visual attention prediction. In: IEEE TIP.
- Wu, X., et al. (2020). SalSAC: A video saliency prediction model with shuffled attentions and correlationbased ConvLSTM . In: AAAI, pp. 12410–12417.
- Xie, S., et al. (2018). Rethinking spatiotemporal feature learning: Speed-accuracy trade-offs in video classification. In: ECCV, pp. 305–321.
- Yubing, T., et al. (2011). A spatiotemporal saliency model for video surveillance. *Cognitive Computation*, 3(1), 241–263.
- Zhang, J., et al. (2018). Deep unsupervised saliency detection: A multiple noisy labeling perspective . In: CVPR, pp. 9029–9038.
- Zhang, P., et al. (2017). Amulet: Aggregating multilevel convolutional features for salient object detection . In: IEEE ICCV.
- Zhang, Y., David, P., Gong, B. (2017). Curriculum domain adaptation for semantic segmentation of urban scenes . In: ICCV, pp. 2020–2030.

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