

Disentangling Spatio-temporal Knowledge for Weakly Supervised Object Detection and Segmentation in Surgical Video

Guiqiu Liao¹

Matjaz Jogan¹

Sai Koushik^{1,3}

Eric Eaton²

Daniel A. Hashimoto^{1,2}

¹ PCASO Laboratory, Department of Surgery, University of Pennsylvania

² Department of Computer and Information Science, University of Pennsylvania

³ Department of Electrical and Systems Engineering, University of Pennsylvania

Abstract

Weakly supervised video object segmentation (WSVOS) enables the identification of segmentation maps without requiring extensive annotations of object masks, relying instead on coarse video labels indicating object presence. WSVOS in surgical videos is, however, more challenging due to the complex interaction of multiple transient objects, such as surgical tools moving in and out of the surgical field. In this scenario, state-of-the-art WSVOS methods struggle to learn accurate segmentation maps. We address this problem by introducing ViDeo Spatio-Temporal disentanglement Networks (VDST-Net), a framework to disentangle complex spatio-temporal object interactions using semi-decoupled knowledge distillation to predict high-quality class activation maps (CAMs). A teacher network is designed to help a temporal-reasoning student network resolve activation conflicts, as the student leverages temporal dependencies when specifics about object location and timing in the video are not provided. We demonstrate the efficacy of our framework on a challenging surgical video dataset where objects are, on average, present in less than 60% of annotated frames, and compare our method to state-of-the-art methods on surgical data and on a public dataset commonly used to benchmark WSVOS. Our method outperforms state-of-the-art techniques and generates accurate segmentation masks under video-level weak supervision. Our code is available at: <https://github.com/PCASOlab/VDST-net>.

1. Introduction

Incorporating expert knowledge into machine learning for surgical video analysis demands extensive and costly annotations. Annotated surgical video databases are rare but crucial for applications ranging from video review and analysis to AI supported robotic assistance [46].

The scarcity of annotated video is not unique to surgery. To address this, Weakly Supervised Semantic Segmentation (WSSS) methods were developed [7, 25, 38, 58], using image-level class labels to identify areas in an image that contain the corresponding object. Such methods have also been applied to video with frame-level labels, exploiting temporal information to improve accuracy [31, 33, 53].

In the more difficult Weakly Supervised Video Object Segmentation (WSVOS), object presence labels are only available at the video level, while no image level presence or segmentation supervision signal is provided. This introduces higher uncertainty about object presence in each frame. Some approaches to WSVOS leverage motion cues [42, 56]. However, these methods often lack robustness against background motion and stationary objects.

A recent end-to-end TCAM model [1] has achieved state-of-the-art performance in object detection on two challenging datasets. In these datasets, the primary objects or concepts present in the video are labeled and are constantly present in most frames. We refer to this type of problem as Consistent Object Presence (COP), in contrast to Transient Object Presence (TOP) where no assumptions are made about object constancy (Figure 1).

The latter case covers the significant variation in temporal presence of objects in surgical videos, where annotations are often based on the start and end times of surgical steps with annotated tools intermittently moving in and out of view. In Transient Object Presence (WSVOS-TOP), the algorithm must thus learn only from object classes without any implicit or explicit information on location, size, or temporal presence, and without assumptions about a specific count of object classes per frame.

In order to address these challenges, we propose a novel method for WSVOS that excels at the TOP problem on surgical data while also achieving state-of-the-art performance in COP scenarios. Our contributions include:

1. We propose an efficient network architecture to gener-

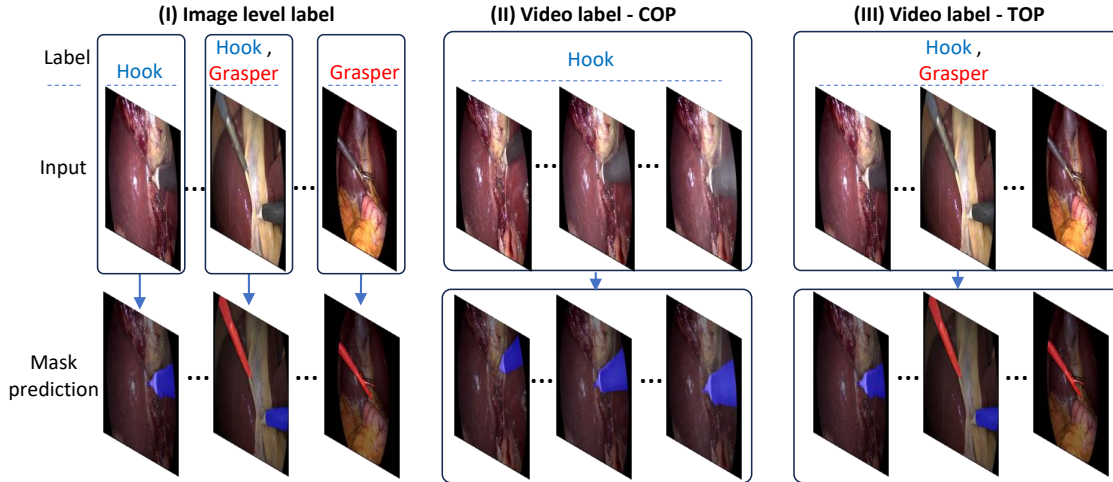


Figure 1. Three types of weakly supervised detection/segmentation in video. Type I has image-level presence labels available, and research for this task usually builds on WSSS adding temporal constraints to improve prediction. Type II WSVOS-COP only has video-level labels, and the object is assumed to be in the video for most of the frames (this is the scenario for YouTube-object dataset). Type III WSVOS-TOP is the most challenging, where the label only indicates the presence of objects for the whole video, yet each object may be present in the video temporarily (i.e., in only a subset of frames).

ate high-quality spatio-temporal class activation maps trained with video-level classification labels.

2. We introduce ViDeo Spatio-Temporal disentanglement Networks (VDST-Net), which captures information associations across time and effectively addresses activation conflicts due to the complex interaction of multiple target objects, including objects appearing and disappearing from the video.
3. Experiments on our TOP surgical video dataset based on Cholec80 [45], which has an average object presence rate of under 60%, demonstrate a clear advantage of our method in WSVOS-TOP.
4. Experiments on YouTube-Objects (v2.2) demonstrate that our method also outperforms state-of-the-art approaches in classic WSVOS-COP.

2. Related Work

2.1. Weakly supervised segmentation for image data

Image-level weakly supervised semantic segmentation has been well-researched, and most existing methods are based on Class Activation Map (CAM) [58]. Raw CAM may be sufficient for detection tasks, but it fails to delineate the target object at the semantic level as it tends to learn the most salient feature. Training strategies such as randomly masking out images [38] or feature maps [13,27] can refine the CAM by enhancing activation of novel features. Loss functions designed to improve CAM have been

proposed, including SEC loss [22], CRF loss [39], regularization losses [52], and contrastive loss [50]. Recently, work on network modeling that explored the long-range dependence and attention mechanism of Vision Transformer (ViT) in CAM [7,47] demonstrated ViT based CAM could outperform Convolutional Neural Networks (CNN) based models in image semantic segmentation tasks.

Emerging work integrates foundation models with weakly supervised segmentation. [30] uses the Contrastive Language-Image Pretraining (CLIP) [35] for WSSS. Segment Anything Model (SAM) was also applied to WSSS [25]. While these methods improve segmentation performance for tasks with image-level labels, they are not directly applicable to weakly supervised segmentation tasks in video.

2.2. Weakly supervised detection and segmentation for video

WSVOS can be categorized into three types of tasks, as shown in Figure 1. In the first type (Type I), frame-level object presence labels are provided. Research in this area often incorporates motion constraints to enhance prediction accuracy, as demonstrated in studies such as [28, 31, 53]. [33] employs LSTM to reason across time for detection and tracking in surgical videos. [14] adapts a 3D CNN to incorporate temporal dependencies for detection in videos.

The second type of WSVOS (Type II) has object labels available only at the video-level, thus not explicitly determining object presence in a particular frame. However, a common assumption in the literature in the WSVOS-COP

setting is that labeled objects are present throughout the entire video, with frames missing these objects treated as noise. This scenario is exemplified by well-known datasets such as YouTube-Objects [18]. This setup is close to a Type I problem as frame content can be inferred from video-level annotations. Recent work on TCAM [1] extracts the most reliable pseudo-semantic supervision signal in video to train a subsequent network for detection in individual images, efficiently addressing object detection for this case, while still showing limited robustness to smaller objects and temporal inconsistencies.

In the third type (Type III), labels merely indicate the general presence of objects throughout an entire video, despite each object possibly appearing only temporarily. This is depicted in the right column of Figure 1. Consequently, the presence label assigned to a video may not accurately reflect a large subset of frames, introducing higher uncertainty that leads to challenge in learning generalized representations. Using motion cues could not resolve such an issue when including scenarios where the camera can occasionally move while the objects stay still.

Existing semantic segmentation methods that use sparse or coarse labels to segment surgical video only focus on the WSSS problem with presence labels for individual frames available [26, 33, 51]. Zero-shot approaches rely on SAM and a prompting strategy [54]. However, no attempts have been made to apply WSVOS to surgical videos or similar TOP scenarios using a single class presence label per video. To the best of our knowledge, our work is the first to address this type of TOP data and train an end-to-end network to predict segmentation masks for the video as a whole.

2.3. Network reasoning across time and knowledge distillation

3D CNN [43], recurrent neural networks (RNN) [37], and temporal Transformers [49] are widely used architectures that capture time dependency and have been applied to various video and volumetric imaging tasks [29, 32, 59]. However, due to the lack of direct supervision signals, learning spatio-temporal features for weakly supervised segmentation is difficult.

Knowledge Distillation (KD) [11] using a teacher-student network was originally designed for model compression and can gradually improve network learning capacity [6]. This property allows KD to enhance learning in weakly or semi-supervised tasks. L2G [15] applied KD to extend CAM from local to global. SCD [48] uses KD to incorporate feature correspondence for CAM refinement in images. Seco [50] employs KD to decouple occurrences for refining CAM in images. Our method employs KD to address the challenges of ill-posed video-level weak supervision, which causes activation conflicts of features that are not associated with the correct class.

We use a constrained teacher to allow the network not only to produce correct activation, but also to refine the backbone to learn generalized features for new domain-specific data. Decoupled KD [57] has proven to be a versatile method for flexibly distilling meaningful knowledge. Our design, which could be categorized as being based on semi-decoupled KD, is tailored for video weakly supervised segmentation. Rather than relying on soft labels from the teacher, we enable the teacher’s predictions while utilizing ground truth to form gates for student supervision. This approach allows the student to learn effectively from an imperfect teacher while still gaining temporal reasoning abilities.

3. Methods

Figure 2 (left) illustrates the proposed VDST-Net framework. The core of the framework consists of a teacher-student network pair designed to disentangle spatial and temporal knowledge. Both utilize a spatio-temporal class activation map (ST-CAM) and share input from ViT [4] encoded frames and video-level class labels as the supervision signal. The primary distinction between the teacher and the student lies in their upper-layer feature extraction modules: the teacher employs a Multilayer Perceptron (MLP), while the student employs a temporal convolutional network (TCN). This design intentionally restricts the teacher from temporal interactions, reducing the potential for activation conflicts between frames, while progressively increasing the student’s ability to reason across time.

We utilize a semi-decoupled online distillation approach, where both teacher and student modules are trained using the same video labels, and efficient knowledge transfer is facilitated through gated activation maps. While the teacher module has a powerful backbone capable of generating highly accurate activation maps, it still exhibits detection gaps and cannot consistently provide fully reliable knowledge for distillation. To resolve these limitations, the semi-decoupled distillation allows the student module to selectively learn from the teacher’s gated outputs under the constraints of video-level ground truth, while enhancing temporal consistency.

3.1. Video encoder

A ViT encoder is used to encode patches from individual frames of the input video $V \in \mathbb{R}^{W \times H \times D}$, producing a stack of token embeddings reshaped as $X \in \mathbb{R}^{w \times h \times D \times C_e}$, where $w = W/P$, $h = H/P$, and P denotes the width/height of a patch, and C_e is the embedding dimension. X captures features with strict temporal correspondence (maintaining the temporal dimension D) and high spatial correspondence.

3.2. Teacher module

As depicted in the blue block of Fig 2, video embedding tensor X is further processed by a MLP block that lin-

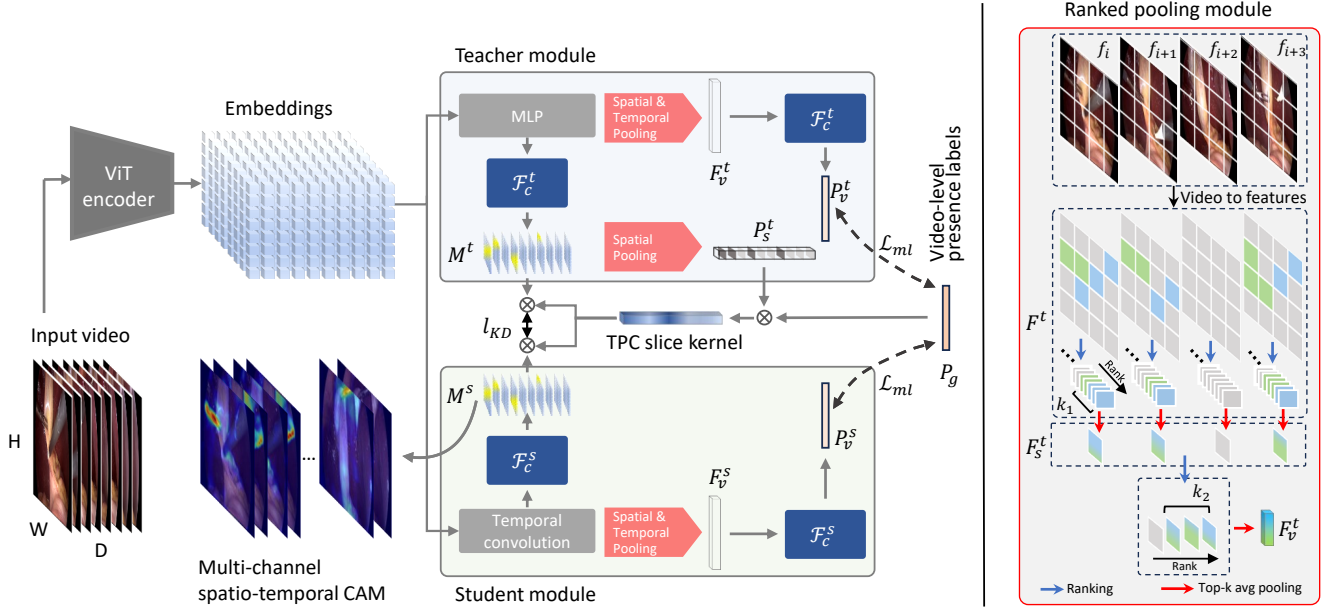


Figure 2. Left: Our VDST-Net approach deploying knowledge distillation to disentangle spatial and temporal information for weakly supervised learning with video-level labels. Ground truth video presence labels P_g provide supervision to both teacher and student, while knowledge in activation map M^t is transferred from teacher to student. Right: Ranked spatial and temporal pooling captures information about multiple objects in a frame while filtering out spurious information from frames where a target object is missing.

early projects X to a feature map $F^t \in \mathbb{R}^{w \times h \times D \times C}$ with C channels, maintaining spatio-temporal correspondence. Following a line of work on ViT-based 2D CAM [7, 47], spatial pooling is applied to F^t to obtain a slice-wise feature map $F_s^t \in \mathbb{R}^{1 \times 1 \times D \times C}$. Then, temporal pooling is applied to acquire a final feature map for video classification $F_v^t \in \mathbb{R}^{1 \times 1 \times 1 \times C}$. Inspired by weakly supervised temporal scene anomaly localization [40] we use a ranked top- k average pooling:

$$F_s^{t,i} = \frac{1}{k_1} \sum_{j=1}^{k_1} F^{t,r_j^i} \quad (1)$$

where $i \in [1, D]$ is the temporal index of feature F_s^t . $r_j^i \in [1, H \times W]$ is the ranked spatial index of F^t at a given temporal slice. This approach first ranks all the spatial patch features by magnitude, and then selects the top- k_1 patches for average pooling. Following a similar strategy, ranked temporal pooling is applied as:

$$F_v^t = \frac{1}{k_2} \sum_{l=1}^{k_2} F_s^{t,r_l} \quad (2)$$

where $r_l \in [1, D]$ is the ranked temporal index of F_s^t . k_2 is used to select top temporal features. The ranked spatio-temporal pooling process is illustrated in the Fig. 2 (right). Maximum pooling or global average pooling can be seen as extreme cases of ranked pooling (i.e. $k_2=1$ or $k_2=D$), which encourage activating either on the most salient feature location or whole video, causing low accuracy in the

surgical video TOP scenario. Ranked pooling design can mitigate such issues and force the ST-CAM to activate on the proper temporal slices and area of interest. Finally, with a fully connected classification layer \mathcal{F}_c^t , the final video prediction $P_v^t \in \mathbb{R}^{1 \times 1 \times 1 \times N}$ is obtained by $P_v^t = \mathcal{F}_c^t(F_v^t)$, where N is the number of classes. Plugging \mathcal{F}_c^t channel-wise to feature map F^t results in a multi-channel ST-CAM $M^t \in \mathbb{R}^{w \times h \times D \times N}$. To help the backbone learn generalized features at different scales, after each MLP layer, features are randomly downsampled during training with a probability of 0.5.

3.3. Student module

Employing an MLP for encoding of embeddings X , the teacher module processes each frame in isolation, thus eliminating temporal conflicts. However, this fails to exploit the temporal coherence to enhance prediction robustness. We thus designate a temporal convolution-based network in the student module. As depicted in Fig. 2, the student network adopts the same spatio-temporal pooling mechanism from the teacher and trains an additional fully connected classification layer \mathcal{F}_c^s . This layer produces the student's video prediction $P_v^s \in \mathbb{R}^{1 \times 1 \times 1 \times N}$ and the student ST-CAM $M^s \in \mathbb{R}^{w \times h \times D \times N}$. The student's TCN can infer associations between frames, yet will introduce errors due to the absence of frame-level annotation. To mitigate this issue, we utilize a semi-decoupled KD technique that refines

a student network with higher capacity through efficiently selected supervision signals.

3.4. Training loss and semi-decoupled knowledge distillation

The teacher module optimizes on a multi-label soft margin classification loss $\mathcal{L}_{ml}(P_v^t, P_g)$ utilizing ground truth video labels $P_g \in \mathbb{R}^{1 \times N}$. The student module optimizes the same classification loss and a KD loss that further refines the student’s activation map. We employ a semi-decoupled [57] approach to knowledge distillation where weak video ground truth labels supervise both teacher and student networks’ predictions while applying KD through activation maps. KD was used to transfer teacher’s prediction logits as soft labels [6, 11], or attention [8, 55]; our approach is closer to latter, as it uses attention maps to guide the student [8]. KD loss is computed by considering only the true positive channel (TPC) and positive slice predictions from the teacher network.

Initially, a teacher slice prediction $P_s^t \in \mathbb{R}^{1 \times 1 \times D \times N}$ is obtained by applying maximum spatial pooling to the teacher ST-CAM. Subsequently, we create a TPC slice kernel $K_{tp} \in \mathbb{R}^{1 \times 1 \times D \times N}$ through the operation $K_{tp} = P_s^t \otimes r(P_g)$, where $r(\cdot)$ represents a broadcast operation to match dimensions for element-wise multiplication \otimes . A final gated KD loss is calculated by:

$$l_{KD} = \mathcal{L}_{mse}(M^s \otimes r(K_{tp}), M^t \otimes r(K_{tp})) \quad (3)$$

where \mathcal{L}_{mse} is a Mean Squared Error loss function, and $r(\cdot)$ here matches the dimension of K_{tp} to M^t and M^s . In analogy, this operation focusing exclusively on frames where true positive pseudo masks are predicted, akin to singling out specific instances for loss calculation. The final loss used for optimizing the student module is given by:

$$l_s = \mathcal{L}_{ml}(P_v^s, P_g) + \alpha l_{KD} \quad (4)$$

In Eq. (4) α is a scaling factor for the distillation loss. Finally, we apply a ReLU layer and normalization to the ST-CAMs extracted from classification networks.

4. Experiments

4.1. Datasets

To investigate the characteristics of our proposed method, we utilize two distinct video dataset collections: Cholec80 [45] dataset and YouTube-Objects video dataset (v2.2) [18]. The Cholec80 dataset consists of 80 cholecystectomy surgery videos, each downsampled to 1 FPS, with presence labels of seven surgical tools ($N = 7$) in each frame. We create a variant TOP video dataset based on Cholec80. First these videos are split into 30-second clips without overlap, merging the tool labels across frames to

generate a video-level tool presence binary vector label for each clip. The test set is based on CholecSeg8K [12] which contains 8K images from 17 videos in Cholec80 with segmentation labels for instruments and anatomy. From this dataset we derive 100 video clips containing 3K frames. To prevent contamination between testing and training data, 5,296 training clips are sampled from 63 Cholec80 videos not included in CholecSeg8k. We calculate tool frequencies across clips using frame per clip (FPC) which denotes the percentage of frames in a clip with the tool present.

We also test our method on YouTube-Objects to demonstrate the generalizability and robustness of our proposed method across different types of videos. YouTube-Objects dataset is a well-known benchmark in the computer vision community which contains 106 videos for training, and 49 videos for testing, where ten categories ($N = 10$) of objects are covered. Each video is split into short-duration clips, and in each clip a few frames are annotated with a bounding box for evaluation.

YouTube-Objects dataset has nearly 100% FPC for each clip and thus represents the second case (Type II – WSVOS-COP) demonstrated in Figure 1. Most categories in Cholec80 based dataset appear on average in less than 60% of the frames of a particular video labeled with that category, which makes it a much more challenging dataset that represents a WSVOS-TOP (Type III) problem.

4.2. Experimental Setup and Evaluation Metrics

We refrained from data augmentation and present the results trained only with original images in this paper. Video frames were resized to 256×256 pixels. Adam optimizer with a learning rate of 1×10^{-4} and a weight decay of 1×10^{-4} was used for training. For the initial nine epochs, only the teacher module and backbone were optimized. Subsequently, the student module was activated alongside the teacher. The k_1 value is calculated as 10% of the number of patches. We fine-tune the k_2 value for the ranked-pooling based on data, and the final results reported in this paper are under $k_2 = 67\%L_v$ (L_v is video length) for the YouTube-Objects data, and $k_2 = 40\%L_v$ for the surgical data. We assessed the classification accuracy on frame and video level. For all the models, final segmentations were refined using dense CRF [23]. For the surgical dataset, segmentation performance was evaluated using Intersection over Union (IoU), Dice score and Hausdorff distance (HD). For localization performance we use CorLoc (IoU > 50%) rate [1].

4.3. Surgical video tool segmentation performance

On surgical videos, we compare our method to frame-level weakly supervised methods MCTformer [47], LayerCAM [16], and XGradCAM [5], all trained with frame-level labels. Models were adapted to our dataset using their

Table 1. Comparison to state-of-the-art weakly supervised methods with either frame-level supervision or video-level supervision. Video (V) and Frame (F) accuracy (AC).

Method	V-AC[%] ↑	F-AC[%] ↑	IoU[%] ↑	HD[pixel] ↓
MCTformer [47]	98.27	97.63	43.14	50.07
LayerCAM [16]	95.67	94.87	31.97	75.59
XGradCAM [5]	92.21	90.21	32.13	53.53
TCAM [1]	97.08	95.10	28.22	54.60
T only (ours)	98.26	97.56	47.50	48.72
T-S fusion (ours)	97.11	97.23	48.64	39.61
VDST-Net (ours)	96.82	98.23	61.80	28.10




Figure 3. Activation maps of surgical video clips from different methods (MCTformer, TCAM, and our VDST-Net).

open-source code implementations. We also compare our method to the video-level supervised method TCAM [1], where both our method and TCAM are trained with more challenging video-level class labels. We test our model using teacher (T) only, teacher-student (T-S) fusion, and full VDST-Net variations. For the surgical video data we utilize pre-trained Masked Auto-Encoder (MAE) ViT backbone [10] which is tuned on a large scale data for segmentation [19] before it is adapted to the new domain (Cholec80).

As presented in Table 1, MCTformer shows the best

performance among frame-level supervision methods. Our teacher-only method, also using a transformer encoder but trained on video-level labels, achieves comparable performance thanks to MLP and our spatio-temporal pooling. Full VDST-Net significantly enhances the segmentation performance (61.80% IoU), with a slight trade-off in video classification score due to potential conflicts between KD loss and classification loss. Figure 3 presents activation maps predicted from different methods, VDST-Net produces activation maps capable of discerning object shapes more effectively than MCTformer.

TCAM achieved commendable video and frame accuracy of 97.08% and 95.10%, respectively. However, in contrast to detection results on the YouTube-Objects (see below), its segmentation performance drops compared to our method, indicating a challenge in learning representations for WSVOS-TOP problems (i.e. in surgical videos).

4.4. Performance on YouTube-Objects

As shown in Table 2, among state-of-the-art methods, TCAM [1] leads in performance on the YouTube-Objects dataset for localization. However, qualitatively, it tends to underfit the object contours and over-activate on background, resulting in lower segmentation quality compared to our method (Figure 4). Our method achieves an overall accuracy improvement from 72.2% to 75.6%. Segmentation scores are not reported due to the lack of ground truth masks in the official dataset.

Originally, we utilized the pre-trained DINO ViT [2] as backbone and fine-tuned it on domain-specific data. Here we also measure performance using a ResNet34 backbone by converting features to the form of ViT. We investigated different optimization strategies for the backbone on new domain data: back-propagation through the teacher (denoted as “refined by T”), through the student (“refined by S”), or keeping the backbone frozen. We also present detection results from either Teacher (T) or final student (full), or extracted from fused CAM (T-S fusion).

Results show that even a frozen ResNet captures a certain level of generalized features. Through our semi-decoupled KD approach, we improve performance from 56.6% to 69.1%. Interestingly, refining the backbone through the student tends to degrade the final output compared to MLP teacher (68.4% to 67.4%), likely due to the challenge of learning generalized features from video-level labels. Optimizing the backbone through the teacher proves to be the effective approach for both ResNet and ViT backbones leveraging well-learned features. For ResNet, we observe that the T-S fusion has a better localization score than the final student (72.2% vs. 70.6%), possibly because ResNet features are less generalized and lean more toward the MLP from the teacher module, while the activations learned by the student complement those of the teacher.

Table 2. Detection performance on YouTube-Objects [34] dataset under CorLoc Metrics. S=Student, T=Teacher

Method		Aero	Bird	Boat	Car	Cat	Cow	Dog	Horse	Mbike	Train	avg
(Joulin et al., 2014) [17]		25.1	31.2	27.8	38.5	41.2	28.4	33.9	35.6	23.1	25.0	31.0
(Kwak et al., 2015) [24]		56.5	66.4	58.0	76.8	39.9	69.3	50.4	56.3	53.0	31.0	55.8
(Rochan et al., 2016) [36]		60.8	54.6	34.7	57.4	19.2	42.1	35.8	30.4	11.7	11.4	35.8
(Tokmakov et al., 2016) [42]		71.5	74.0	44.8	72.3	52.0	46.4	71.9	54.6	45.9	32.1	56.6
(Koh et al., 2016) [21]		64.3	63.2	73.3	68.9	44.4	62.5	71.4	52.3	78.6	23.1	60.2
(Tsai et al., 2016) [44]		66.1	59.8	63.1	72.5	54.0	64.9	66.2	50.6	39.3	42.5	57.9
(Haller and Leordeanu, 2017) [9]		76.3	71.4	65.0	58.9	68.0	55.9	70.6	33.3	69.7	42.4	61.2
(Croitoru et al., 2019) [3]		81.7	51.5	54.1	62.5	49.7	68.8	55.9	50.4	33.3	57.0	56.5
XGradCAM(Fu et al., 2020) [5]		68.2	44.5	45.8	64.0	46.8	86.4	44.0	57.0	44.9	45.0	54.6
LayerCAM(Jiang et al., 2021) [16]		80.0	84.5	47.2	73.5	55.3	83.6	71.3	60.8	55.7	48.1	66.0
TCAM (Belharbi et al., 2023) [1]		79.4	94.9	75.7	61.7	68.8	87.1	75.0	62.4	72.1	45.0	72.2
VDST-Net + Res34 frozen	T only	70.6	97.8	44.7	78.0	31.3	57.7	53.8	42.5	73.9	15.6	56.6
	T-S fusion	81.6	97.7	74.4	73.9	66.7	71.4	71.8	51.6	59.5	25.0	67.4
	full	91.5	93.5	67.4	82.5	63.4	80.6	64.1	56.3	55.3	36.2	69.1
VDST-Net + Res34 refined by S	T only	88.0	95.7	74.6	69.7	78.5	80.7	57.1	62.4	55.8	20.9	68.4
	T-S fusion	87.7	95.1	76.1	70.5	81.6	81.2	58.7	61.3	58.0	22.9	69.3
	full	86.4	95.1	77.8	71.8	76.2	77.2	54.2	55.1	57.7	22.1	67.4
VDST-Net + Res34 refined by T	T only	86.1	94.3	75.6	72.9	76.7	76.2	57.9	60.4	61.6	23.1	68.5
	T-S fusion	90.0	93.3	77.2	71.8	84.8	80.7	60.9	68.8	58.6	36.4	72.2
	full	85.4	90.9	73.2	69.0	82.5	75.2	55.8	73.8	51.5	48.4	70.6
VDST-Net + DINO ViT (refined by T)	T only	83.8	88.6	69.5	59.3	80.7	76.6	43.2	64.9	48.7	17.8	63.3
	T-S fusion	91.0	88.2	72.7	68.7	85.9	79.2	52.4	70.4	55.9	37.0	70.1
	full	92.3	91.8	72.3	70.1	82.2	82.5	73.6	72.6	67.8	50.7	75.6

In conclusion, we recommend using ViT within our framework and advise against direct student-driven backbone refinements.

4.5. Ablation Study

We evaluated the effects of architecture modifications, pooling variations, and training techniques on model performance (Table 3). Figure 5 shows activation maps from different conditions. Additionally, we demonstrate a further refinement of segmentation generated by using activation maps as prompts for the SAM [19], showcasing a way of generating semantic pseudo-labels.

Importance of knowledge distillation: The results highlight the significance of KD through teacher and student modules for enhanced segmentation accuracy. A standalone student module, lacking frame-level supervision, shows a notable decline in segmentation performance (Dice score from 67.80% to 47.64%), because it predicts error activation by taking information from wrong frames or features, pointed out by red arrowheads in Figure 5. The teacher module has no such issue but lacks the ability to generate activation maps with good connectivity, as indicated by red asterisks in the figure.

Effect of ranked top- k pooling: We compare the ranked top- k pooling to global average and max pooling. Specifically, with average pooling, we noted a decrease in Dice score by 11.49%, and with max pooling, the decrease was 11.71%. Furthermore, there was a considerable decrease in frame classification accuracy, amounting to 3.72% with av-




Figure 4. Segmentation and detection performance on YouTube-Objects. The second and third columns are the activation maps and post-processed binary masks of our method. In the last column results are taken from [1].

erage pooling and 3.04% with max pooling.

TPC decoupling strategy: Additionally, the proposed TPC slice kernel for activations gating, enhances the student’s learning. The absence of this component leads to a decline in segmentation performance by 4.16% in Dice

Table 3. Ablation study with modifications of overall architecture, spatio-temporal pooling and training strategies. W/O denotes a specific module is removed.

	Conditions	Video AC[%]↑	Frame AC[%]↑	IoU[%]↑	Dice[%]↑	HD[pixel]↓
	VDST-Net full	96.82	98.23	61.80	67.80	28.10
Structural ablation	Teacher only	98.41	97.56	47.50	53.07	48.72
	Student only	94.37	94.76	42.83	47.78	69.45
Training	Avg pooling	94.37	94.51	53.16	56.31	41.65
	Max strategy	94.95	95.19	52.47	56.09	40.05
	W/O TPC kernel	96.54	97.10	60.46	63.64	40.75

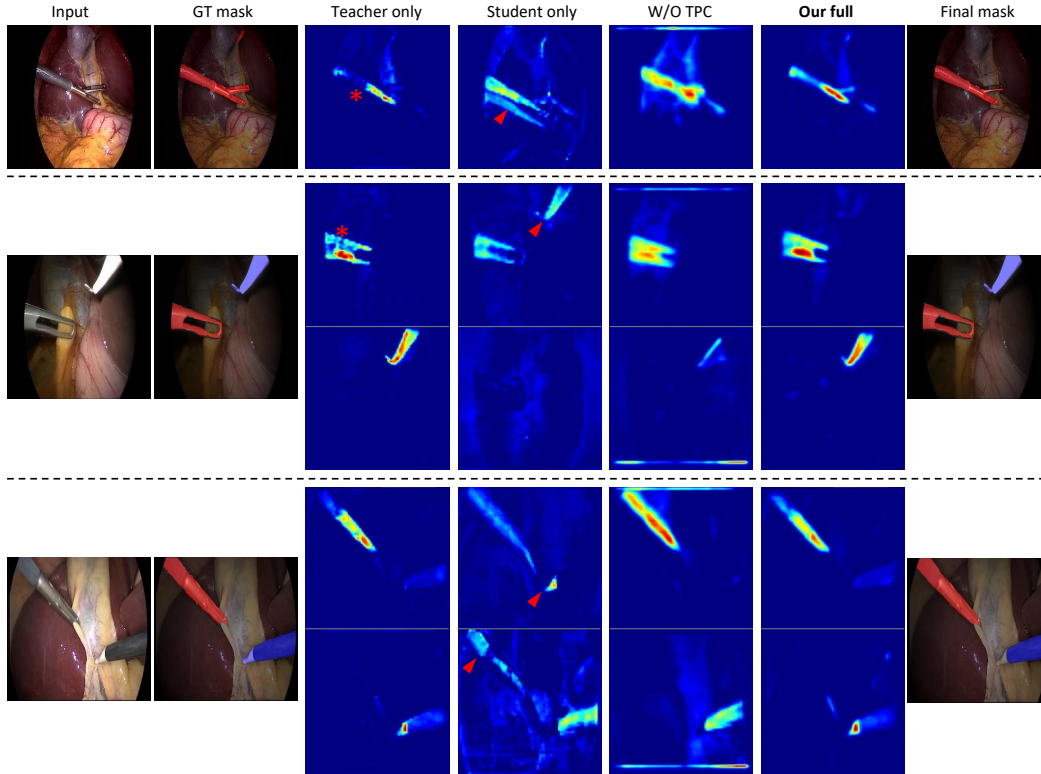


Figure 5. Qualitative results of ablation study and final segmentation masks.

and 1.34% in IoU, as demonstrated in Table 3, due to over-introduced loss from non-target class knowledge [57].

5. Conclusion

We presented a novel framework named VDST-Net for WSVOS in surgical videos, which is end-to-end trained with coarse video labels. This method uses a teacher-student architecture and KD to disentangle spatial and temporal knowledge, minimizing temporal interference for the teacher while enhancing the student’s learning capacity. Experiments on both COP and TOP datasets demonstrate our method’s efficacy in generating accurate segmentation maps, outperforming state-of-the-art approaches.

Future work could involve using pseudo masks generated by the student network to train a subsequent lightweight student network (e.g., U-net) for real-time in-

ference. Additionally, our method could be applied to other video datasets beyond the surgical domain, and exploring spatial knowledge captured by other backbones or foundation models for WSVOS using our method would be interesting. We also suggest that, although the student should not directly update the backbone in a shared backbone design, employing separate backbones combined with continual learning strategies such as Exponential Moving Average (EMA) [41] or Elastic Weight Consolidation (EWC) [20] for updating weights could be beneficial.

6. Acknowledgements.

This work was supported by the University of Pennsylvania Thomas B. McCabe and Jeannette E. Laws McCabe Fellow Award and the Linda Pechenik Montague Investigator Award.

References

- [1] Soufiane Belharbi, Ismail Ben Ayed, Luke McCaffrey, and Eric Granger. TCAM: Temporal class activation maps for object localization in weakly-labeled unconstrained videos. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023. 1, 3, 5, 6, 7
- [2] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, 2021*. 6
- [3] Ioana Croitoru, Simion-Vlad Bogolin, and Marius Leordeanu. Unsupervised learning of foreground object segmentation. *International Journal of Computer Vision*, 2019. 7
- [4] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, 2021*. 3
- [5] Ruigang Fu, Qingyong Hu, Xiaohu Dong, Yulan Guo, Yinghui Gao, and Biao Li. Axiom-based Grad-CAM: Towards accurate visualization and explanation of CNNs. In *31st British Machine Vision Conference 2020, BMVC, 2020*. 5, 6, 7
- [6] Tommaso Furlanello, Zachary Chase Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born-again neural networks. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, Proceedings of Machine Learning Research*, 2018. 3, 5
- [7] Wei Gao, Fang Wan, Xingjia Pan, Zhiliang Peng, Qi Tian, Zhenjun Han, Bolei Zhou, and Qixiang Ye. TS-CAM: Token semantic coupled attention map for weakly supervised object localization. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, 2021*. 1, 2, 4
- [8] Ziyao Guo, Haonan Yan, Hui Li, and Xiaodong Lin. Class attention transfer based knowledge distillation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, 2023*. 5
- [9] Emanuela Haller and Marius Leordeanu. Unsupervised object segmentation in video by efficient selection of highly probable positive features. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, 2017*. 7
- [10] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, 2022*. 6
- [11] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *ArXiv preprint*, 2015. 3, 5
- [12] W-Y Hong, C-L Kao, Y-H Kuo, J-R Wang, W-L Chang, and C-S Shih. CholecSeg8k: A semantic segmentation dataset for laparoscopic cholecystectomy based on Cholec80. *ArXiv preprint*, 2020. 5
- [13] Qibin Hou, Peng-Tao Jiang, Yunchao Wei, and Ming-Ming Cheng. Self-erasing network for integral object attention. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada, 2018*. 2
- [14] Rui Hou, Chen Chen, and Mubarak Shah. Tube convolutional neural network (T-CNN) for action detection in videos. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, 2017*. 2
- [15] Peng-Tao Jiang, Yuqi Yang, Qibin Hou, and Yunchao Wei. L2G: A simple local-to-global knowledge transfer framework for weakly supervised semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, 2022*. 3
- [16] Peng-Tao Jiang, Chang-Bin Zhang, Qibin Hou, Ming-Ming Cheng, and Yunchao Wei. LayerCAM: Exploring hierarchical class activation maps for localization. *IEEE Transactions on Image Processing*, 2021. 5, 6, 7
- [17] Armand Joulin, Kevin Tang, and Li Fei-Fei. Efficient image and video co-localization with Frank-Wolfe algorithm. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, Proceedings, Part VI 13*. Springer, 2014. 7
- [18] Vicky Kalogeiton, Vittorio Ferrari, and Cordelia Schmid. Analysing domain shift factors between videos and images for object detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (11), 2016. 3, 5
- [19] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloé Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross B. Girshick. Segment anything. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, 2023*. 6, 7
- [20] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. In *Proceedings of the National Academy of Sciences*, number 13. National Acad Sciences, 2017. 8
- [21] Yeong Jun Koh, Won-Dong Jang, and Chang-Su Kim. POD: Discovering primary objects in videos based on evolutionary refinement of object recurrence, background, and primary object models. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, 2016*. 7
- [22] Alexander Kolesnikov and Christoph H Lampert. Seed, expand and constrain: Three principles for weakly-supervised image segmentation. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, Proceedings, Part IV 14*. Springer, 2016. 2
- [23] Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected CRFs with Gaussian edge potentials. In *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011, Granada, Spain, 2011*. 5

- [24] Suha Kwak, Minsu Cho, Ivan Laptev, Jean Ponce, and Cordelia Schmid. Unsupervised object discovery and tracking in video collections. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, 2015*. 7
- [25] Hyeokjun Kweon and Kuk-Jin Yoon. From SAM to CAMs: Exploring Segment anything model for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2024*. 1, 2
- [26] Eung-Joo Lee, William Plishker, Xinyang Liu, Shuvra S Bhattacharyya, and Raj Shekhar. Weakly supervised segmentation for real-time surgical tool tracking. *Healthcare Technology Letters*, (6), 2019. 3
- [27] Jungbeom Lee, Eunji Kim, Sungmin Lee, Jangho Lee, and Sungroh Yoon. FickleNet: Weakly and semi-supervised semantic image segmentation using stochastic inference. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, 2019*. 2
- [28] Haofeng Li, Guanqi Chen, Guanbin Li, and Yizhou Yu. Motion guided attention for video salient object detection. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), 2019*. 2
- [29] Jianing Li, Shiliang Zhang, and Tiejun Huang. Multi-scale 3D convolution network for video based person re-identification. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, 2019*. 3
- [30] Yuqi Lin, Minghao Chen, Wenxiao Wang, Boxi Wu, Ke Li, Binbin Lin, Haifeng Liu, and Xiaofei He. CLIP is also an efficient segmenter: A text-driven approach for weakly supervised semantic segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, 2023*. 2
- [31] Qing Liu, Vignesh Ramanathan, Dhruv Mahajan, Alan L. Yuille, and Zhenheng Yang. Weakly supervised instance segmentation for videos with temporal mask consistency. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2021*. 1, 2
- [32] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2022*. 3
- [33] Chinedu Innocent Nwoye, Didier Mutter, Jacques Marescaux, and Nicolas Padoy. Weakly supervised convolutional LSTM approach for tool tracking in laparoscopic videos. *International Journal of Computer Assisted Radiology and Surgery*, 2019. 1, 2, 3
- [34] Alessandro Prest, Christian Leistner, Javier Civera, Cordelia Schmid, and Vittorio Ferrari. Learning object class detectors from weakly annotated video. In *2012 IEEE Conference on Computer Vision and Pattern Recognition, Providence, RI, USA, 2012*. 7
- [35] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Aspell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML, Proceedings of Machine Learning Research, 2021*. 2
- [36] Mrigank Rochan, Shafin Rahman, Neil DB Bruce, and Yang Wang. Weakly supervised object localization and segmentation in videos. *Image and Vision Computing*, 2016. 7
- [37] Alex Sherstinsky. Fundamentals of recurrent neural network (RNN) and long short-term memory (LSTM) network. *Physica D: Nonlinear Phenomena*, 2020. 3
- [38] Krishna Kumar Singh and Yong Jae Lee. Hide-and-seek: Forcing a network to be meticulous for weakly-supervised object and action localization. In *IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, 2017*. 1, 2
- [39] Meng Tang, Federico Perazzi, Abdelaziz Djelouah, Ismail Ben Ayed, Christopher Schroers, and Yuri Boykov. On regularized losses for weakly-supervised CNN segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV), 2018*. 2
- [40] Yu Tian, Guansong Pang, Yuanhong Chen, Rajvinder Singh, Johan W. Verjans, and Gustavo Carneiro. Weakly-supervised video anomaly detection with robust temporal feature magnitude learning. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, 2021*. 4
- [41] Joshua Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and P. Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 23–30, 2017. 8
- [42] Pavel Tokmakov, Karteek Alahari, and Cordelia Schmid. Weakly-supervised semantic segmentation using motion cues. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, Proceedings, Part IV 14*. Springer, 2016. 1, 7
- [43] Du Tran, Lubomir D. Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3D convolutional networks. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, 2015*. 3
- [44] Yi-Hsuan Tsai, Guangyu Zhong, and Ming-Hsuan Yang. Semantic co-segmentation in videos. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, Proceedings, Part IV 14*. Springer, 2016. 7
- [45] Andru P Twinanda, Sherif Shehata, Didier Mutter, Jacques Marescaux, Michel De Mathelin, and Nicolas Padoy. EndoNet: A deep architecture for recognition tasks on laparoscopic videos. *IEEE Transactions on Medical Imaging*, (1), 2016. 2, 5
- [46] Mikhail Volkov, Daniel A Hashimoto, Guy Rosman, Ozanan R Meireles, and Daniela Rus. Machine learning and coresets for automated real-time video segmentation of laparoscopic and robot-assisted surgery. In *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017. 1

- [47] Lian Xu, Wanli Ouyang, Mohammed Bennamoun, Farid Boussaid, and Dan Xu. Multi-class token transformer for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 2, 4, 5, 6
- [48] Rongtao Xu, Changwei Wang, Jiayi Sun, Shibiao Xu, Weiliang Meng, and Xiaopeng Zhang. Self correspondence distillation for end-to-end weakly-supervised semantic segmentation. In *Thirty-Seventh AAAI Conference on Artificial Intelligence, AAAI 2023, Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI 2023, Thirteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2023, Washington, DC, USA, 2023*. 3
- [49] Bin Yan, Houwen Peng, Jianlong Fu, Dong Wang, and Huchuan Lu. Learning spatio-temporal transformer for visual tracking. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, 2021*. 3
- [50] Zhiwei Yang, Kexue Fu, Minghong Duan, Linhao Qu, Shuo Wang, and Zhijian Song. Separate and conquer: Decoupling co-occurrence via decomposition and representation for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2, 3
- [51] Zixin Yang, Richard Simon, and Cristian Linte. A weakly supervised learning approach for surgical instrument segmentation from laparoscopic video sequences. In *Medical Imaging 2022: Image-Guided Procedures, Robotic Interventions, and Modeling*. SPIE, 2022. 3
- [52] Yazhou Yao, Tao Chen, Guo-Sen Xie, Chuanyi Zhang, Fumin Shen, Qi Wu, Zhenmin Tang, and Jian Zhang. Non-salient region object mining for weakly supervised semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR, 2021*. 2
- [53] Novanto Yudistira, Muthu Subash Kavitha, and Takio Kurita. Weakly-supervised action localization, and action recognition using global-local attention of 3D CNN. *International Journal of Computer Vision*, (10), 2022. 1, 2
- [54] Wenxi Yue, Jing Zhang, Kun Hu, Yong Xia, Jiebo Luo, and Zhiyong Wang. SurgicalSAM: Efficient class promptable surgical instrument segmentation. In *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2024, Vancouver, Canada, 2024*. 3
- [55] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, Conference Track Proceedings, 2017*. 5
- [56] Dingwen Zhang, Junwei Han, Le Yang, and Dong Xu. SPFTN: A joint learning framework for localizing and segmenting objects in weakly labeled videos. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (02), 2020. 1
- [57] Borui Zhao, Quan Cui, Renjie Song, Yiyu Qiu, and Jiajun Liang. Decoupled knowledge distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 3, 5, 8
- [58] Bolei Zhou, Aditya Khosla, Àgata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, 2016*. 1, 2
- [59] Xiaoyu Zhu, Jeffrey Chen, Xiangrui Zeng, Junwei Liang, Chengqi Li, Sinuo Liu, Sima Behpour, and Min Xu. Weakly supervised 3D semantic segmentation using cross-image consensus and inter-voxel affinity relations. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, 2021*. 3