



HAL
open science

Temporally Constrained Neural Networks (TCNN): A framework for semi-supervised video semantic segmentation

Deepak Alapatt, Pietro Mascagni, Armine Vardazaryan, Alain Garcia, Nariaki Okamoto, Didier Mutter, Jacques Marescaux, Guido Costamagna, Bernard Dallemagne, Nicolas Padoy

► To cite this version:

Deepak Alapatt, Pietro Mascagni, Armine Vardazaryan, Alain Garcia, Nariaki Okamoto, et al.. Temporally Constrained Neural Networks (TCNN): A framework for semi-supervised video semantic segmentation. 2021. hal-04212987

HAL Id: hal-04212987

<https://hal.science/hal-04212987v1>

Preprint submitted on 21 Sep 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Temporally Constrained Neural Networks (TCNN): A framework for semi-supervised video semantic segmentation

Deepak Alapatt, Pietro Mascagni, Armine Vardazaryan, Alain Garcia, Nariaki Okamoto, Didier Mutter, Jacques Marescaux, Guido Costamagna, Bernard Dallemagne and Nicolas Padoy

Abstract—A major obstacle to building models for effective semantic segmentation, and particularly video semantic segmentation, is a lack of large and well annotated datasets. This bottleneck is particularly prohibitive in highly specialized and regulated fields such as medicine and surgery, where video semantic segmentation could have important applications but data and expert annotations are scarce. In these settings, temporal clues and anatomical constraints could be leveraged during training to improve performance. Here, we present Temporally Constrained Neural Networks (TCNN), a semi-supervised framework used for video semantic segmentation of surgical videos. In this work, we show that autoencoder networks can be used to efficiently provide both spatial and temporal supervisory signals to train deep learning models. We test our method on a newly introduced video dataset of laparoscopic cholecystectomy procedures, Endoscapes, and an adaptation of a public dataset of cataract surgeries, CaDIS. We demonstrate that lower-dimensional representations of predicted masks can be leveraged to provide a consistent improvement on both sparsely labeled datasets with no additional computational cost at inference time. Further, the TCNN framework is model-agnostic and can be used in conjunction with other model design choices with minimal additional complexity.

Index Terms—video segmentation, autoencoders, surgery, convolutional neural networks, semi-supervision

I. INTRODUCTION

IMAGE-GUIDED interventions are becoming the mainstay of surgery. The analysis of the videos natively guiding such minimally invasive interventions holds great potential to

This work was partially supported by French State Funds managed by the “Agence Nationale de la Recherche (ANR)” through the “Investissements d’Avenir” (Investments for the Future) Program under Grant ANR-10-IAHU-02 (IHU-Strasbourg) and through the National AI Chair program under Grant ANR-20-CHIA-0029-01 (Chair AI4ORSafety). This work was granted access to the HPC resources of IDRIS under the allocation 2021-AD011011640R1 made by GENCI.

Deepak Alapatt, and Nicolas Padoy are affiliated with ICube, University of Strasbourg, CNRS, France (email: {alapatt, npadoy}@unistra.fr)

Pietro Mascagni and Guido Costamagna are affiliated with Fondazione Policlinico Universitario A. Gemelli IRCCS, Rome, Italy

Pietro Mascagni, Armine Vardazaryan, Alain Garcia, Didier Mutter and Nicolas Padoy are affiliated with IHU-Strasbourg, Institute of Image-Guided Surgery, Strasbourg, France

Nariaki Okamoto, Jacques Marescaux and Bernard Dallemagne are affiliated with Institute for Research against Digestive Cancer (IRCAD), Strasbourg, France

study and assist surgical procedures [1]. In this context, video analysis for surgical scene understanding is a key enabler towards context-aware computer-assisted interventions [2], post-operative analysis of surgical procedures for documentation [3], research and education [4], and eventually, robot-assisted surgery [5]. The task of image semantic segmentation presents a straightforward path to scene understanding through the categorization of every pixel of an image based on the structure it constitutes. However, a major limitation is the significant annotation effort and opportunity cost that annotating segmentation datasets entails for medical professionals, particularly when annotating fine-grained structures with high surgical semantics. A natural consequence is a lack of large and well-annotated surgical segmentation datasets compared to the computer vision community at large. As a result, there has been an increasing focus on optimally using the limited and precious datasets that are available.

Overall, the extension of image semantic segmentation methods to video-based applications, or video semantic segmentation, remains a challenging and open problem. One shortcoming of applying methods designed for image semantic segmentation to videos is that they fail to leverage the rich and abundant relationships between frames of videos. Moreover, proposed approaches for video semantic segmentation are often complex and computationally demanding.

Further, traditional approaches treat semantic segmentation as a purely local problem, often training models using loss functions that penalize errors only at a pixel level. This approach has proven effective when working with large and well labeled datasets. However, intrinsic regularities in the data, such as the spatial relationship between anatomical structures in surgical videos, could be leveraged in the output space to boost performance even on smaller datasets.

To address these challenges, we propose a framework for semi-supervised video semantic segmentation, Temporally Constrained Neural Networks (TCNN). To validate our approach, we benchmark our method on two challenging surgical video datasets for both tool and anatomy segmentation. We believe this is the first work explicitly designed on the video semantic segmentation of anatomical classes in surgical videos.

Our contributions can be summarized as follows:

- We generate a new dataset for the segmentation of fine-

grained anatomical structures in surgical videos.

- We demonstrate that autoencoder networks can be used to regularize models on highly variable and dynamic surgical videos.
- We propose a new framework for semi-supervised video semantic segmentation that is independent of model design choice and incurs no additional inference computational cost.
- We present results and analysis of our method in the context of two challenging surgical video datasets. We pick two datasets representing distinct surgical procedures, laparoscopic cholecystectomy and microscopic cataract surgery, that differ greatly in appearance and dynamism in order to illustrate the generalizability of our approach.

II. CONTEXT

A. Image and Video Semantic Segmentation

Image semantic segmentation is an active research field that has seen significant progress since the pioneering work applying fully convolutional networks for the task [6]. Subsequent methods have focused on high quality [7] [8] [9] [10] [11] [12] and/or efficient [13] [14] [15] [16] design choices. More recently, the design of models for video semantic segmentation has received increasing attention. These methods largely fall into two categories: (1) approaches that reuse high-level features inferred from key frames for inference at surrounding frames [17] [18] [19] [20]; (2) approaches that run deep neural networks on every frame and use additional network layers to aggregate temporal information [21] [22] [23] [24] [25]. Due to a lack of densely labeled datasets, most methods in both categories use unlabeled frames to incorporate temporal context in their model training in a semi-supervised manner. However, the former category of methods primarily does so to minimize inference time, often at the cost of prediction performance [19] [20], while the latter category focuses on using temporal information to improve predictions. Our work falls into the second group but unlike most of the presented methods, our work is not restricted to a particular model design choice but rather presents a model-agnostic framework for spatio-temporal regularization of deep neural networks. For video semantic segmentation, in general, the use of optical flow appears to be the dominant strategy. However, flow-based methods often incur significant additional computational costs and are prone to failure when dealing with large motions, non-textured regions, and occlusions. These conditions are common in surgical videos and are further exacerbated by the highly similar appearance of anatomical structures. In the context of video semantic segmentation on surgical data, the topic remains largely unexplored barring a few works that have utilized optical flow for segmenting surgical instruments [26] [27]. To the best of our knowledge, no previously published work has been designed explicitly for the video semantic segmentation of anatomical classes. This may be because of the limited publicly available datasets or the additional complexity of the task.

B. Autoencoders and Semantic Segmentation

Most methods designed for semantic segmentation rely on loss functions such as cross entropy or dice loss to supervise the training process by penalizing errors made at a pixel level. As stated in [28], this treatment is suboptimal, particularly in constrained settings such as surgical video segmentation, since non-local features such as smoothness and topology are not explicitly taken into consideration. Moreover, learning constraints in the shape, position, and interaction between classes implicitly can be challenging without sizeable and representative datasets, as is common in the medical space. Our work closely relates to a group of recently proposed methods that address these limitations using autoencoder networks. [29] first introduced a cascaded architecture of a segmentation network followed by a denoising autoencoder, which was trained to map segmentation masks into a shape space before reconstructing them back. By incorporating an additional loss term that ensures that the encoded representations for the ground truth mask and the predicted mask lie close by in the shape space, they could train the network to learn non-local regularities in the output space such as shape. Along the same lines, [28] use a similar cascaded design except using a frozen encoder to prevent over-regularization. [30] [31] utilize a similar mask denoising autoencoder as a post-processing step to "filter out" irregularities in the output space such as isolated blobs. Different from these methods, we demonstrate that an autoencoder trained to map segmentation masks to a lower-dimensional shape space could be further utilized to provide temporal supervision using the surrounding unlabeled frames in a sequence. In addition, we test the proposed framework on surgical videos, much more dynamic and unconstrained than the volumetric radiological imaging used to evaluate the methods described in this subsection.

C. Video Semantic Segmentation Datasets

As previously stated, generating large and representative datasets for video semantic segmentation can be extremely challenging. Most available datasets either limit the number of sequences being annotated (≤ 10) [32] or provide sparse annotations ($> 1\text{Hz}$) [33]. While a few publicly available video semantic segmentation datasets do exist in the computer vision community, similar datasets in the surgical domain are lacking. The few video datasets that do exist provide only labels for instrument classes [34] [35] or label very few sequences [36]. Recently, CholecSeg8k [37] was introduced, providing anatomy and instrument labels for 101 selected short sequences (~ 3 seconds) from 17 publicly available laparoscopic cholecystectomy (LC) procedure recordings [38]. In this paper, we generate and use Endoscapes, a new dataset consisting of 201 endoscopic video recordings of LC procedures annotated with segmentation masks of fine-grained anatomical structures. Compared to CholecSeg8k, Endoscapes comprises many more surgical procedures (201 vs 17), annotated classes (29 vs 13) and provides sparse annotations, every 30 seconds, rather than densely labeled short sequences. Additionally, we adapt a public image semantic segmentation dataset for cataract surgery, CaDIS [39], to generate temporal sequences. We

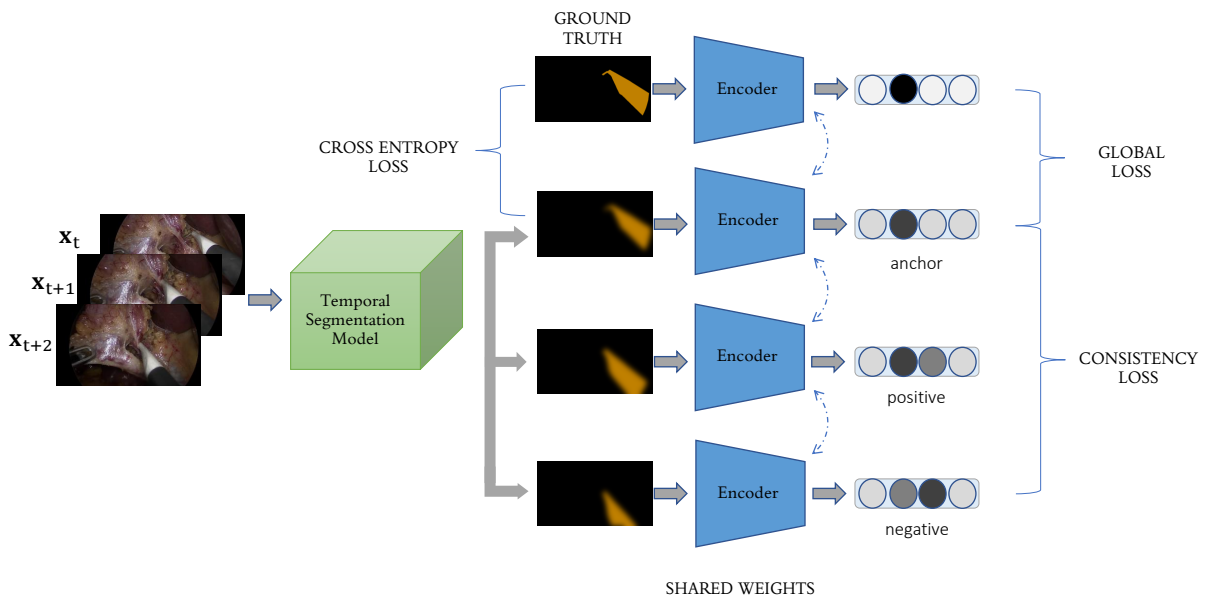


Fig. 1: An overview of the proposed TCNN framework.

hope that this work will promote research into video semantic segmentation for surgical scenes, which comes with a unique set of characteristics and challenges that warrant independent and further study; for example:

- Surgical images may present a number of unique confounding factors such as smoke, bleeding, unstable camera motion which may not be as common in natural images.
- There may be an extremely high inter-class similarity between certain anatomical structures.
- In contrast to most video datasets, surgical datasets typically depict an evolving scene rather than a changing one, i.e. the same environment and structures are usually in focus.

III. ENDOSCOPES

Cholecystectomy is a very common abdominal surgical procedure almost ubiquitously performed with a laparoscopic approach, hence guided by an endoscopic video. Deep learning models for LC video analysis have been developed with the aim of assisting surgeons during interventions [40] [41], improving staff awareness and readiness, and facilitating post-operative documentation and research [3]. However, datasets and models for video semantic segmentation of LC are lacking. Recognizing fine-grained hepatocystic anatomy through semantic segmentation could help surgeons better assess the critical view of safety (CVS), a universally recommended technique consisting in well exposing anatomical landmarks to prevent bile duct injuries [42]. Additionally, segmentation masks of hepatocystic structures could be leveraged by deep learning models for automatic assessment of CVS [40] and surgical action recognition [43] to improve their performance. We believe that generating a dataset for video semantic segmentation of hepatocystic anatomy will promote surgical data science research and accelerate the development

of applications for surgical safety. To generate a representative dataset, consecutive endoscopic videos of LC performed at Nouvel Hôpital Civil (Strasbourg, France) were collected. Non-endoscopic, i.e., out-of-body, video frames were blacked-out to comply with European data protection regulations. A frame every 30 seconds was sampled from the portion of the endoscopic video showing the hepatocystic anatomy being dissected, the most critical phase of the surgical procedure, and when surgeons should achieve the CVS. Such unselected and regularly spaced video frames were manually annotated with pixel-wise semantic annotations of anatomical and surgical instances, such as the cystic artery and the dissection of the hepatocystic triangle, respectively. Annotations were performed in double by specifically trained computer scientists and surgeons. For reproducibility and transparency, an in-depth description of the annotation process can be found in [44]. Overall, 1933 regularly spaced video frames from 201 LC videos were annotated with segmentation mask for 29 classes. A more detailed description of the annotated classes and their representation in Endoscapes is presented in Table 1.

IV. TEMPORALLY CONSTRAINED NEURAL NETWORKS

In this section, we present a framework for semi-supervised semantic segmentation, which we term Temporally Constrained Neural Networks (TCNN). Our work is inspired by Anatomically Constrained Neural Networks [28], a recently proposed framework utilizing autoencoder [45] networks to penalize global features in addition to local penalties such as cross entropy which are a standard for semantic segmentation. We extend this work to the much more dynamic and variable domain of surgical videos and further utilize the autoencoder network to supervise network training using the global and temporal context of the prediction. A brief overview of our framework is presented in Fig 1.

TABLE I: Class Distribution for Endoscopes.

Category	Class Name	Endoscopes (All Images)	Training Set	Validation Set	Test Set	
		Instances per Class (%)	Instances per Class (%)	Instances per Class (%)	Instances per Class (%)	
Anatomy	Background	1933 (100%)	1212 (100%)	409 (100%)	312 (100%)	
	Cystic Plate	701 (36.26%)	433 (35.73%)	138 (33.74%)	130 (41.67%)	
	HC Triangle (dissection)	683 (35.33%)	426 (35.15%)	133 (32.52%)	124 (39.74%)	
	Cystic Artery	1020 (52.77%)	636 (52.48%)	192 (46.94%)	192 (61.54%)	
	Cystic Duct	1477 (76.41%)	952 (78.55%)	267 (65.28%)	258 (82.69%)	
	Gallbladder	1812 (94.20%)	1174 (96.86%)	360 (88.01%)	287 (91.99%)	
Instrument	Grasper Tip	1386 (71.70%)	879 (72.52%)	277 (67.73%)	230 (73.72%)	
	Grasper Shaft	332 (17.18%)	210 (17.33%)	73 (17.85%)	49 (15.70%)	
	Bipolar Tip	138 (7.14%)	91 (7.51%)	37 (9.05%)	10 (3.20%)	
	Bipolar Shaft	109 (5.64%)	70 (5.78%)	30 (7.34%)	9 (2.88%)	
	Hook Tip	957 (49.51%)	611 (50.41%)	187 (45.72%)	159 (50.96%)	
	Hook Shaft	825 (42.68%)	536 (44.22%)	156 (38.14%)	133 (42.63%)	
	Scissors Tip	5 (0.26%)	3 (0.25%)	0 (0%)	2 (0.64%)	
	Scissors Shaft	4 (0.21%)	3 (0.25%)	0 (0%)	1 (0.32%)	
	Clipper Tip	211 (10.92%)	120 (9.90%)	48 (11.74%)	43 (13.78%)	
	Clipper Shaft	200 (10.35%)	117 (9.65%)	45 (11.00%)	38 (12.18%)	
	Plastic Clipper Tip	5 (0.26%)	5 (0.41%)	0 (0%)	0 (0%)	
	Plastic Clipper Shaft	1 (0.05%)	1 (0.08%)	0 (0%)	0 (0%)	
	Irrigator	72 (3.72%)	47 (3.88%)	11 (2.69%)	14 (4.49%)	
	Retractor	6 (0.31%)	4 (0.33%)	0 (0%)	2 (0.64%)	
	Bipolar Variant Tip	2 (0.10%)	2 (0.16%)	0 (0%)	0 (0%)	
	Bipolar Shaft Variant	1 (0.05%)	1 (0.08%)	0 (0%)	0 (0%)	
	Hook Variant Tip	101 (5.22%)	71 (5.86%)	15 (3.67%)	15 (4.81%)	
	Hook Shaft Variant	98 (5.07%)	67 (5.53%)	15 (3.67%)	16 (5.13%)	
	Clipper Variant Tip	1 (0.05%)	1 (0.08%)	0 (0%)	0 (0%)	
	Clipper Shaft Variant	1 (0.05%)	1 (0.08%)	0 (0%)	0 (0%)	
	Other	Lymph Node	453 (23.44%)	314 (25.91%)	98 (23.96%)	41 (13.14%)
		Undissected Gallbladder	12 (0.62%)	6 (0.50%)	6 (1.47%)	0 (0%)
Anatomical Variant		88 (4.55%)	73 (6.02%)	15 (3.67%)	0 (0%)	




Fig. 2: Augmentations strategies to train the autoencoder network. The autoencoder to reconstruct ground truth masks (top row), augmented masks generated by overlaying convex envelopes of anatomical classes (middle row) and generated using affine transforms and morphological operations (bottom row).

A. Methodology

Given a fully annotated semantic segmentation dataset $\mathbf{D} = \{(\mathbf{x}, \mathbf{y}); \mathbf{x} \in \mathbf{X}, \mathbf{y} \in \mathbf{Y}\}$, for each input image $\mathbf{x} = \{x_i; x_i \in \mathbb{R}, i \in \mathbb{S}\}$ a segmentation model (ϕ) is trained to estimate the corresponding ground truth mask $\mathbf{y} = \{y_i; y_i \in \mathbb{C}, i \in \mathbb{S}\}$. Conventionally, convolutional neural net-

works (CNN) are trained to do this by minimizing a loss function L_{seg} penalizing errors in predicting the ground truth mask. In most cases, L_{seg} is chosen to supervise the model at a pixel level using functions such as cross entropy, dice loss, etc. However, these loss functions do not fully exploit global constraints and regularities in the output space. As described in [28], this approach may be suboptimal for datasets depicting objects that are highly regular in their appearance or position such as in the case of anatomical structures in medical images. In [28], the authors make use of an autoencoder network to penalize global features in addition to a cross entropy loss term. To do this, they first train an autoencoder network, comprising an encoder and decoder, to be able to reconstruct segmentation masks. Internally, the encoder (ψ) learns a hidden representation, $\psi(\mathbf{y})$, which the decoder uses to reconstruct the original mask. If the hidden representation is undercomplete, i.e. of lower dimensionality than the input mask, the autoencoder learns to represent only the most salient, or global, features of the input mask in the encoding while discarding more granular, or local, features. They then balance their local cross entropy loss function with a loss term penalizing more global errors in the prediction:

$$L_{global} = d(\psi(\phi(\mathbf{x})), \psi(\mathbf{y})). \quad (1)$$

Here, d is a metric function quantifying the distance between the encodings for the predicted and ground truth mask. This loss term ensures that the prediction is close to the ground truth mask in the same low dimensional space, for example on a shape manifold. We further hypothesize that if such a low dimensional space can be learned, then the

predicted masks should move continuously across that space over time. We incorporate this into the training scheme in order to provide temporal supervision using an additional contrastive loss term. For three temporally consecutive predicted masks $\phi(\mathbf{x}_t)$ (anchor), $\phi(\mathbf{x}_{t+1})$ (positive), and $\phi(\mathbf{x}_{t+2})$ (negative), we use a triplet loss constraining the three corresponding encodings e_a , e_p , and e_n , respectively.

$$L_{consistency} = \max(0, d(e_a, e_p) - d(e_a, e_n)). \quad (2)$$

Using the additional unlabeled frames, this term penalizes abrupt changes in the encodings of predicted masks at consecutive timesteps, and consequently, abrupt changes in global characteristics of predicted masks over time. In this semi-supervised setup, we then balance the standard pixel-wise loss functions with the consistency loss and global loss as described below

$$Loss = L_{seg} + \lambda_g L_{global} + \lambda_c L_{consistency}. \quad (3)$$

We would like to emphasize that this framework is both model-agnostic and incurs no additional computational cost at inference time.

B. Training Strategies

Our proposed approach is trained in two stages. The aim of the first stage is to learn a lower-dimensional parameterization of a predicted or ground truth segmentation mask that captures the structure and configuration of the represented classes [28] [46]. To this end, we train an autoencoder network to generate an undercomplete representation, i.e. of lower dimensionality than the input, by learning to reconstruct ground truth segmentation masks. This stage of training is crucial to learning a compact and expressive representation space that can effectively supervise the segmentation network training. The fact that segmentation masks are significantly less detailed and more regular in appearance than surgical images allows for new and more robust augmentation strategies to support this stage of training. Using this property, we propose a number of strategies to simulate segmentation masks and inflate the dataset size: (1) We generate plausible artificial masks by overlaying the convex closure of each anatomical class in a ground truth segmentation mask. We do not modify the tool classes as these correspond to non-deformable classes with characteristic shapes; (2) We perform affine transforms such as rotation and rescaling; (3) Following [30], we use erosion and dilation with variable kernels to degrade the input in order to force the network to ignore minor local variations in the input. Empirically, we observe that these additional augmentation strategies and a deeper CNN design choice for the autoencoder were critical to getting reliable reconstruction accuracy, particularly on highly dynamic data such as laparoscopic videos. A brief overview of the first stage of training and a depiction of the augmentation strategies used is shown in Fig. 2.

In the second stage, we train a video segmentation model in a semi-supervised manner and use the encoding phase of the trained autoencoder to provide additional supervision as described previously. We freeze the autoencoder weights

TABLE II: Autoencoder Structure: Under the column Layer: Conv - Convolutional Layer, ConvT - Transpose Convolutional layer, FC - Fully Connected layer. All convolutional and transpose convolutional layers use a 3x3 kernel with number of filters, stride and activation described by the Num Outputs, Stride and Activation columns, respectively. The second FC layer regenerates the same number of outputs as the number of features entering the first FC layer, which is dependent on input size.

Layer	Num Outputs	Stride	Activation
Conv	16	2	ReLU
Conv	16	1	ReLU
Conv	32	2	ReLU
Conv	32	1	ReLU
Conv	64	2	ReLU
Conv	64	1	ReLU
Conv	64	2	ReLU
Conv	64	1	ReLU
FC	1024	-	None
FC	-	-	None
ConvT	64	2	ReLU
Conv	64	1	ReLU
ConvT	64	2	ReLU
Conv	64	1	ReLU
ConvT	64	2	ReLU
Conv	64	1	ReLU
Conv	32	2	ReLU
ConvT	32	1	ReLU
Conv	# classes	1	None

TABLE III: Results Endoscopes: From top to bottom, the F1 scores are presented per class followed by the mean F1, IoU and Pixel Accuracy over all considered classes. The per class results are ordered from the least represented class in the dataset to the most represented.

Class Name	Model	Image	Video	Ours
	Predictor	Predictor	Predictor	
F1 HC Triangle	60.72 %	62.63 %		64.39 %
F1 Cystic Plate	39.71 %	41.74 %		48.34 %
F1 Cystic Artery	47.74 %	46.31 %		50.79 %
F1 Cystic Duct	61.06 %	62.06 %		62.55 %
F1 Instruments	94.87 %	94.41 %		94.90 %
F1 Gallbladder	92.16 %	92.06 %		92.99 %
F1 Background	97.22 %	97.16 %		97.43 %
Mean F1	70.50 %	70.91 %		73.06 %
Mean IoU	59.14 %	59.47 %		61.58 %
Mean Pixel Accuracy	94.78 %	94.73 %		95.22 %

during this phase of training to avoid the network collapsing to a trivial solution to satisfy the encoding space constraints.

V. EXPERIMENTAL SETUP

We benchmark our algorithm on two surgical video datasets depicting sequences from laparoscopic cholecystectomies and cataract surgeries. We establish two baselines, Deeplab-v3+, a state-of-the-art method for image semantic segmentation used in this setting in [40], and an adapted version of the same model designed to propagate temporal information to make temporally coherent predictions. Hereafter, we refer to these

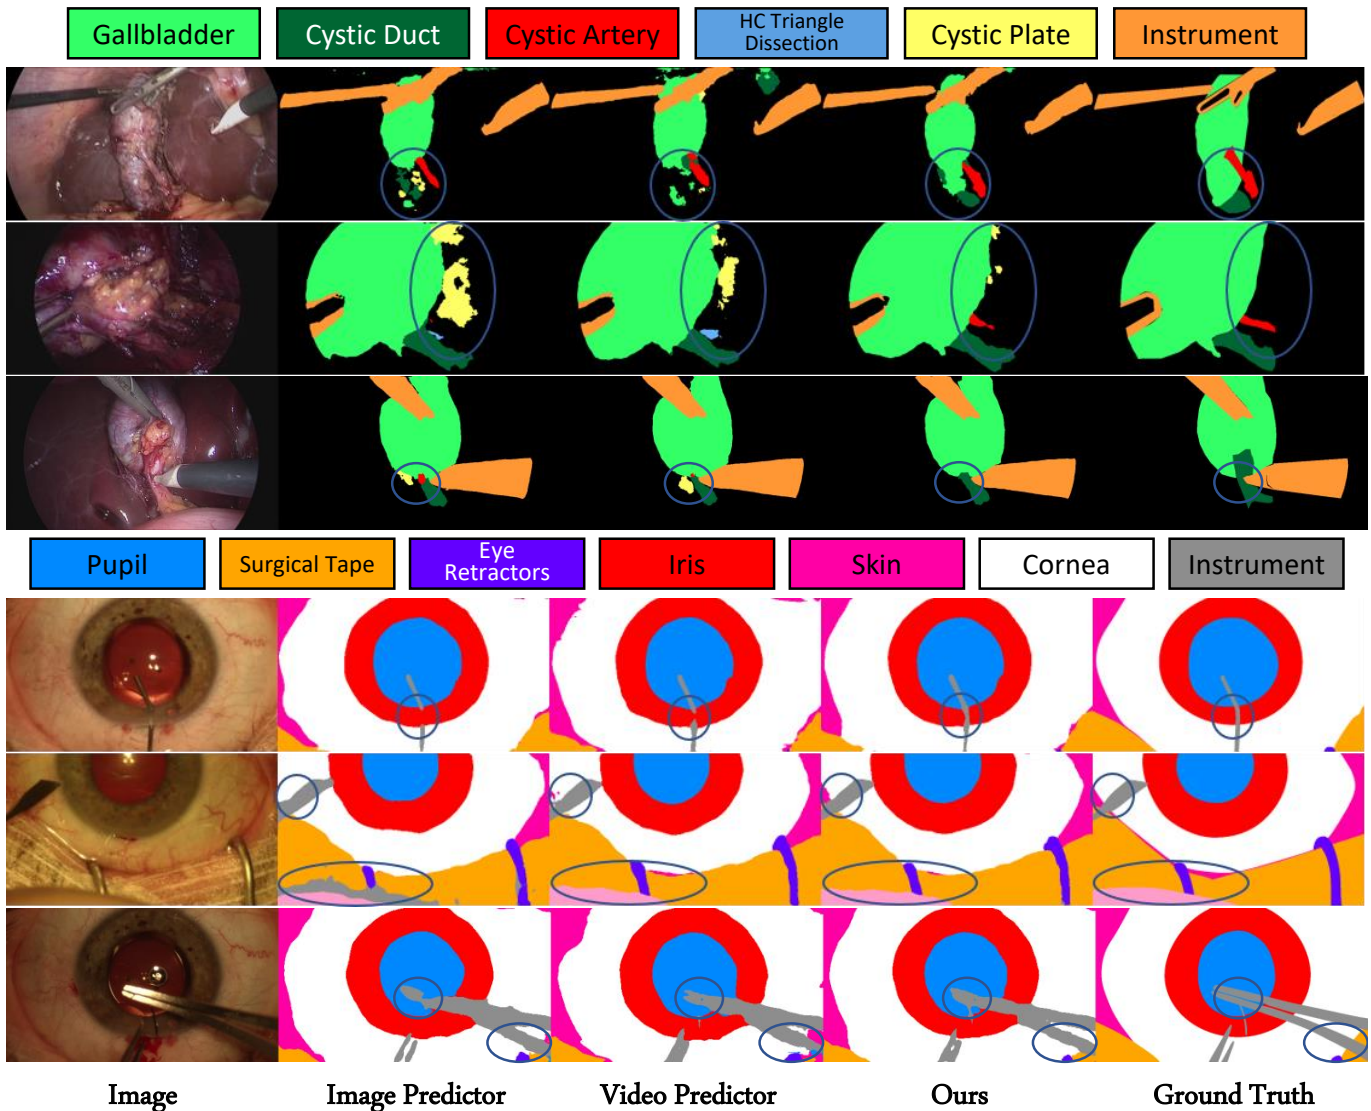


Fig. 3: Qualitative Results: The additional supervision incorporated into the training leads to a significant decrease in disconnected blobs, irregular boundaries and missing classes as highlighted in the circled regions.

TABLE IV: Results CaDIS: From top to bottom, the F1 scores are presented per class followed by the mean F1, IoU and Pixel Accuracy over all considered classes. The per class results are ordered from the least represented class in the dataset to the most represented.

Class Name	Model	Image Predictor	Video Predictor	Ours
F1 Hand		94.59 %	92.73 %	95.29 %
F1 Eye Retractors		80.31 %	82.40 %	83.58 %
F1 Surgical Tape		88.82 %	89.52 %	90.06 %
F1 Instrument		79.68 %	81.72 %	83.22 %
F1 Pupil		96.51 %	96.26 %	96.55 %
F1 Skin		86.75 %	90.81 %	90.77 %
F1 Iris		89.80 %	89.98 %	90.22 %
F1 Cornea		93.58 %	95.17 %	95.21 %
Mean F1		88.76 %	89.82 %	90.61 %
Mean IoU		80.28 %	81.90 %	83.18 %
Mean Pixel Accuracy		91.62 %	93.11 %	93.30 %

models as Image Predictor and Video Predictor, respectively. On both datasets, we analyze the results per class and observe a consistent improvement over the baseline methods in performance when training using the proposed approach.

A. Study Design

To design the Video Predictor, we adapt Deeplabv3+ [7] to propagate temporal information using a 3x3 ConvLSTM [47] module placed between the encoder and decoder. The base architecture for Deeplabv3+ is a modified resnet-50 [48] initialized with Imagenet pretrained weights. The autoencoder is a 18 layer convolutional neural network. Given the design choice for the Video and Image Predictor, this corresponds to $\sim 5\%$ additional network capacity during training. To note, in the inference phase, the autoencoder network can be discarded and inference is performed without any additional complexity. The autoencoder network architecture is detailed in Table 2.

TABLE V: Ablation Study.

Dataset	Temporal Module	Global loss	Consistency loss	mean IoU	mean F1	mean Pixel Accuracy
Endoscapes	\times	\times	\times	59.14 %	70.5 %	94.78 %
	\times	\checkmark	\times	59.25 %	70.42 %	95.04 %
	\checkmark	\times	\times	59.47 %	70.91 %	94.73 %
	\checkmark	\times	\checkmark	60.5 %	71.96 %	94.94 %
	\checkmark	\checkmark	\times	60.92 %	72.39 %	95.09 %
	\checkmark	\checkmark	\checkmark	61.58 %	73.06 %	95.22 %
CaDIS	\times	\times	\times	80.28 %	88.76 %	91.62 %
	\times	\checkmark	\times	79.68 %	88.45 %	91.22 %
	\checkmark	\times	\times	81.9 %	89.82 %	93.11 %
	\checkmark	\times	\checkmark	82.35 %	90.12 %	92.9 %
	\checkmark	\checkmark	\times	82.37 %	90.14 %	92.74 %
	\checkmark	\checkmark	\checkmark	83.18 %	90.61 %	93.3 %

To train each model, we feed a temporally regularly spaced sequence of 3 unlabeled frames and the following labeled frame. To find the optimal weighting for the loss terms for each dataset, we explore combinations of λ_c in $\{0.0001, 0.001, 0.01, 0.1\}$ and λ_g in $\{0.5, 1, 1.5, 2\}$ using L2 and mean square distance as the distance function for $L_{consistency}$ and L_{global} , respectively.

B. Datasets

To evaluate the model in a variety of surgical scenarios, we test the TCNN framework on two surgical video datasets with greatly differing appearance and motion patterns, Endoscapes and CaDIS. The former comprises recordings of laparoscopic cholecystectomy procedures and the latter recordings of microscopic cataract surgeries. Since cataract procedures are performed using a mounted microscope focused on the patient’s eye, the appearance of the scene is much more stable in terms of scale, orientation and motion during and between procedures compared to laparoscopic surgeries.

1) *Endoscapes*: The Endoscapes dataset introduced in Section 3 was distributed as follows: 120 videos, 41 videos, and 40 for training, validation, and testing, respectively. A detailed description of the dataset splits is presented in Table 1.

To mitigate the label imbalance issue, which is not the primary focus of this work, we take a number of steps.

- We ignore all classes that appear in under 5% of the images in any of the training, validation, or test splits during both the training and evaluation phase.
- We group all the remaining instrument tips and shafts into a single class.
- We ignore the lymph node class due to its highly variable appearance, position, and ambiguous boundaries.

We sample frames of resolution 854x480 at 1 frame per second (fps) to generate training sequences.

2) *CaDIS*: CaDIS [39] is an image semantic segmentation dataset containing fully annotated labels for 25 videos or 4671 images. Of this, 3550 frames (19 videos) were used for training, 534 frames (3 videos) for validation, and 587 frames (3 videos) for testing following [39]. This dataset was generated by labeling a subset of frames from the training videos of CATARCTS [49] dataset. In order to generate

temporal sequences, we merge preceding frames available in the CATARCTS dataset with labels from CaDIS.

We follow the first experimental setup in [39] to select class labels and downsize the images to a resolution of 270x480 during both training and evaluation. In this setup, 4 classes correspond to anatomical structures, 1 to instruments and 3 to miscellaneous structures. We sample frames at 15 fps in order to generate training sequences.

C. Evaluation Metrics

Dice Similarity coefficient (F-score), Intersection over Union (IoU), and Pixel Accuracy were used to evaluate the models. Pixels belonging to an ignored class were excluded from evaluation.

D. Implementation Details

All models were trained on a single NVIDIA V100 for 200 epochs using a momentum optimizer with momentum=0.9, using the polynomial learning rate schedule described in [7] with the initial and final learning rates set as $7e-3$ and $1-6$, respectively. The best performing model was selected during training based on mIoU on the validation set. To train the segmentation network, we set $\lambda_c = 0.01$ for Endoscapes, $\lambda_c = 0.001$ for CaDIS and $\lambda_g = 1.5$ and in all presented results based on performance on the validation set.

VI. RESULTS

In Table 3 and Table 4, we present comparative results for the Image Predictor, Video Predictor, and our approach using the TCNN framework. Specifically, we observe an average improvement of 2.6% and 1.9% F1 over the Image Predictor on Endoscapes and CaDIS, respectively. Of note, while there is an improvement in the average performance over all classes, similar to [50] [51], not all classes benefit from the addition of the temporal module by itself. This may be because the increased network capacity leads the model to overfit to the training data, indicating that additional supervision could be necessary. This is reinforced by the fact that there is a marked and consistent improvement when using the global and consistency losses. On both datasets, we observe that the least represented classes benefit the most from

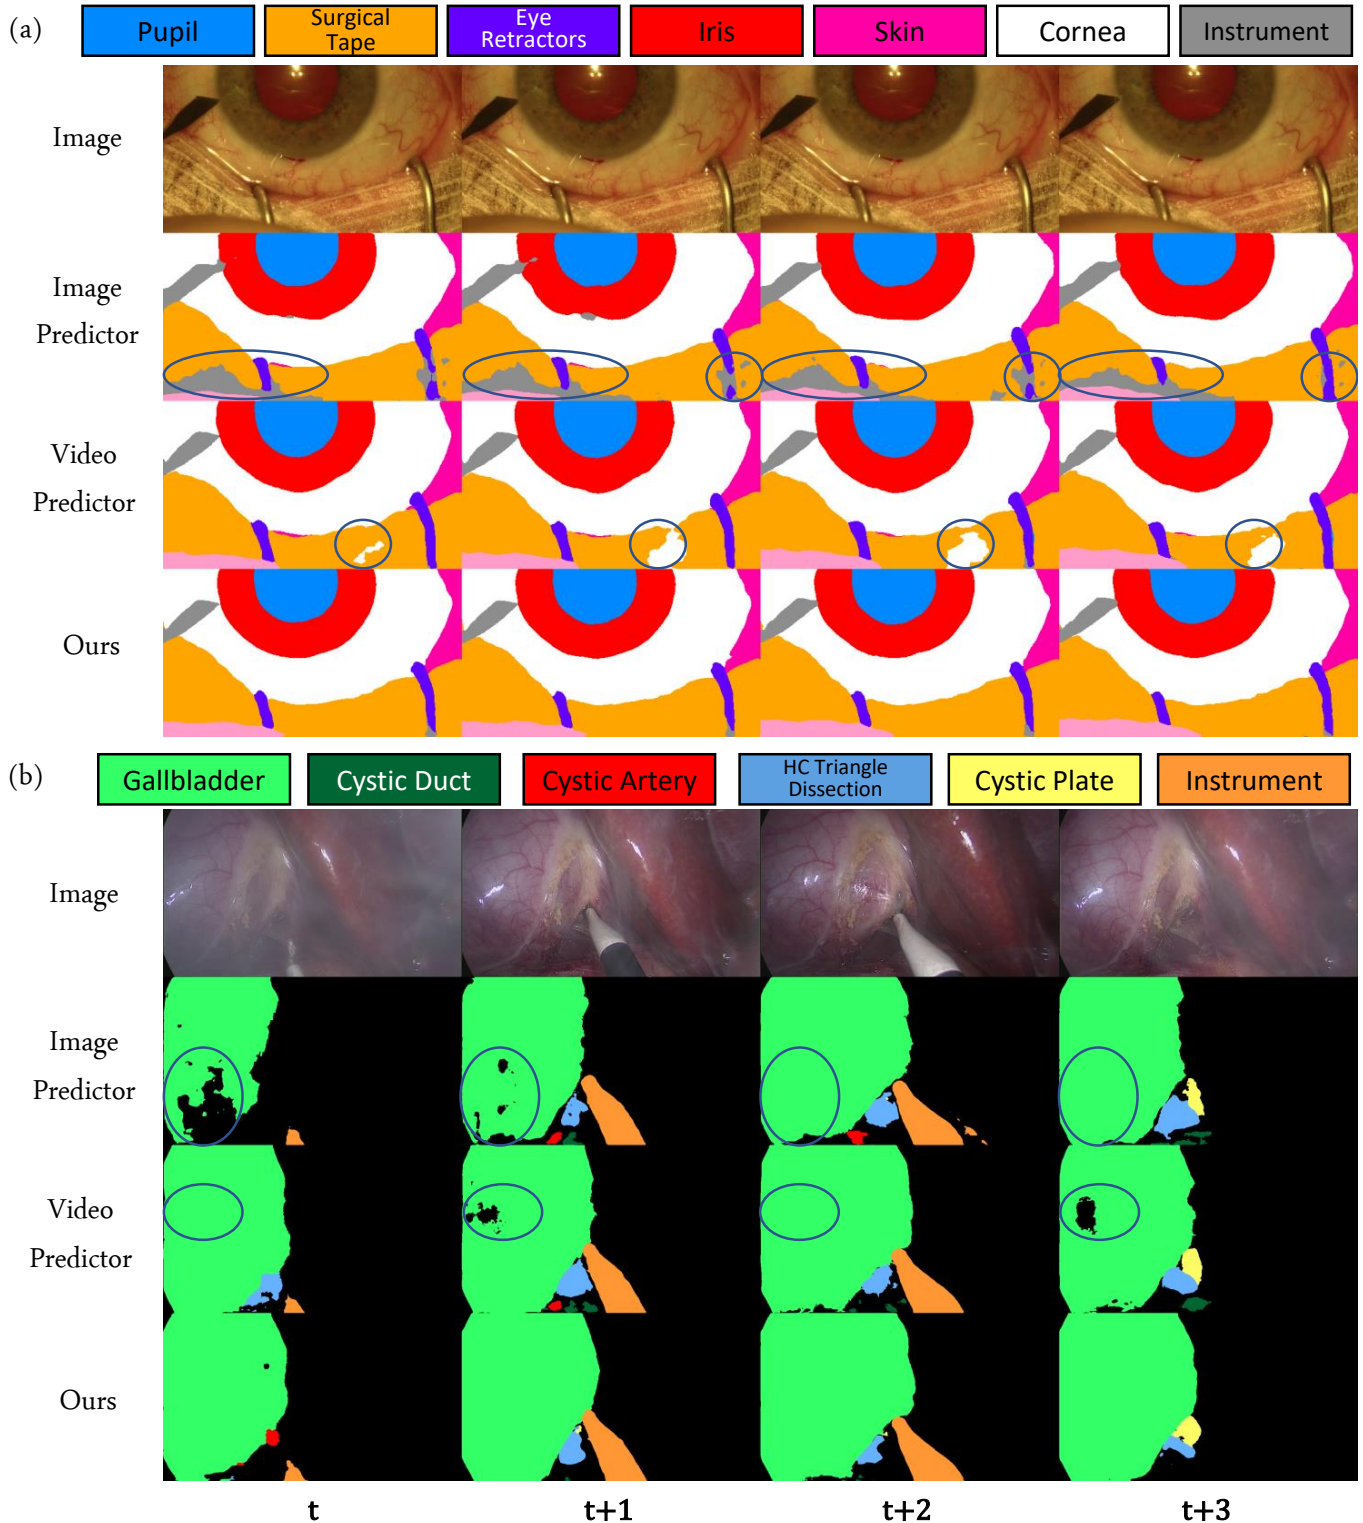


Fig. 4: Example results showing the progression of predictions over consecutive timesteps. Case (a) depicts a sequence from the CaDIS test set showing how the addition of the temporal module in the Video Predictor allows for a more consistent prediction over time and further how the additional supervision leads to a more logical one. Case (b) shows a sequence from the Endoscapes test set showing how the TCNN framework makes the model less susceptible to error when dealing with smoke, a common confounding during surgical procedures.

the additional supervision. Using the additional loss terms, we obtain a mean improvement of 4.3% and 1.4% over the Video Predictor on the 3 least represented classes in Endoscapes and CaDIS, respectively, at the exact same computational cost. On the Endoscapes dataset, there is a substantial performance improvement in the identification of the Cystic Plate class with a gain of 6.6% F1 over the Video Predictor. While this class does not have a characteristic appearance because it is defined surgically by the quality of the dissection, it does have more particular global characteristics, such as position, that may be better captured using our proposed framework. We provide some qualitative results in Fig. 3 to illustrate the benefits of the additional supervision. We see that the TCNN framework helps minimize the occurrence of improbable predictions such as isolated blobs and anatomical structures at unrealistic positions. In Fig. 4, we further demonstrate how the additional loss terms help make more consistent predictions over consecutive timesteps, making the network more robust to confounding factors such as smoke and occlusions.

Finally, we present the results of the ablation study in Table 5. Interestingly, the global loss term proves more effective when used in conjunction with the temporal module. Specifically, global loss contributes to an increase of 1.48% and 0.32% F1 for the Video Predictor on Endoscapes and CaDIS, respectively. In contrast, it results in a decrease in performance of the image predictor by 0.08% and 0.31% F1 on the two datasets, respectively. Note, that this case corresponds to the method proposed in [28]. We hypothesize that the additional temporal context provided through the temporal module allows the network to better rectify global discrepancies in the predicted mask. Finally, on both datasets we observe that the consistency loss provides a consistent improvement and works well in combination with the global loss.

VII. CONCLUSION

In this work, we present a framework, TCNN, for semi-supervised video semantic segmentation utilizing autoencoder networks to provide spatio-temporal regularization for convolutional neural networks. Our approach is easy to implement and is independent of the model design, providing additional supervisory signals to arbitrary video semantic segmentation models. With minimal additional computational cost during training, the TCNN framework shows consistent improvement in segmentation performance on multiple datasets with exactly the same runtime and computational cost during inference. We hope the work on this framework and presented datasets will enhance video semantic segmentation in settings with limited data and intrinsic topological constraints and promote research into surgical video segmentation.

Ethical approval The surgical videos were recorded and collected in an anonymized manner following the informed consent of patients in compliance with the local Institutional Review Board (IRB) requirements.

REFERENCES

- [1] P. Mascagni and N. Padoy, "Or black box and surgical control tower: Recording and streaming data and analytics to improve surgical care," *Journal of Visceral Surgery*, 2021.
- [2] T. Vercauteren, M. Unberath, N. Padoy, and N. Navab, "Cai4cai: the rise of contextual artificial intelligence in computer-assisted interventions," *Proceedings of the IEEE*, vol. 108, no. 1, pp. 198–214, 2019.
- [3] P. Mascagni, D. Alapatt, T. Urade, A. Vardazaryan, D. Mutter, J. Marescaux, G. Costamagna, B. Dallemagne, and N. Padoy, "A computer vision platform to automatically locate critical events in surgical videos: documenting safety in laparoscopic cholecystectomy," *Annals of Surgery*, vol. 274, no. 1, pp. e93–e95, 2021.
- [4] T. M. Ward, P. Mascagni, A. Madani, N. Padoy, S. Perretta, and D. A. Hashimoto, "Surgical data science and artificial intelligence for surgical education," *Journal of Surgical Oncology*, vol. 124, no. 2, pp. 221–230, 2021.
- [5] L. R. Kennedy-Metz, P. Mascagni, A. Torralba, R. D. Dias, P. Perona, J. A. Shah, N. Padoy, and M. A. Zenati, "Computer vision in the operating room: Opportunities and caveats," *IEEE transactions on medical robotics and bionics*, vol. 3, no. 1, pp. 2–10, 2020.
- [6] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 3431–3440.
- [7] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 801–818.
- [8] C. Peng, X. Zhang, G. Yu, G. Luo, and J. Sun, "Large kernel matters—improve semantic segmentation by global convolutional network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4353–4361.
- [9] J. Wang, K. Sun, T. Cheng, B. Jiang, C. Deng, Y. Zhao, D. Liu, Y. Mu, M. Tan, X. Wang *et al.*, "Deep high-resolution representation learning for visual recognition," *IEEE transactions on pattern analysis and machine intelligence*, 2020.
- [10] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, "Pyramid scene parsing network," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2881–2890.
- [11] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 12, pp. 2481–2495, 2017.
- [12] G. Lin, A. Milan, C. Shen, and I. Reid, "Refinenet: Multi-path refinement networks for high-resolution semantic segmentation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 1925–1934.
- [13] C. Yu, J. Wang, C. Peng, C. Gao, G. Yu, and N. Sang, "Bisenet: Bilateral segmentation network for real-time semantic segmentation," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 325–341.
- [14] H. Zhao, X. Qi, X. Shen, J. Shi, and J. Jia, "Icnet for real-time semantic segmentation on high-resolution images," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 405–420.
- [15] A. Paszke, A. Chaurasia, S. Kim, and E. Cukurciello, "Enet: A deep neural network architecture for real-time semantic segmentation," *arXiv preprint arXiv:1606.02147*, 2016.
- [16] A. Chaurasia and E. Cukurciello, "Linknet: Exploiting encoder representations for efficient semantic segmentation," in *2017 IEEE Visual Communications and Image Processing (VCIP)*. IEEE, 2017, pp. 1–4.
- [17] S. Jain, X. Wang, and J. E. Gonzalez, "Accel: A corrective fusion network for efficient semantic segmentation on video," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 8866–8875.
- [18] Y. Li, J. Shi, and D. Lin, "Low-latency video semantic segmentation," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 5997–6005.
- [19] E. Shelhamer, K. Rakelly, J. Hoffman, and T. Darrell, "Clockwork convnets for video semantic segmentation," in *European Conference on Computer Vision*. Springer, 2016, pp. 852–868.
- [20] X. Zhu, Y. Xiong, J. Dai, L. Yuan, and Y. Wei, "Deep feature flow for video recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 2349–2358.
- [21] D. Nilsson and C. Sminchisescu, "Semantic video segmentation by gated recurrent flow propagation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 6819–6828.

- [22] R. Gadde, V. Jampani, and P. V. Gehler, "Semantic video cnns through representation warping," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 4453–4462.
- [23] X. Jin, X. Li, H. Xiao, X. Shen, Z. Lin, J. Yang, Y. Chen, J. Dong, L. Liu, Z. Jie *et al.*, "Video scene parsing with predictive feature learning," in *Proceedings of the IEEE International Conference on Computer Vision*, 2017, pp. 5580–5588.
- [24] M. Ding, Z. Wang, B. Zhou, J. Shi, Z. Lu, and P. Luo, "Every frame counts: joint learning of video segmentation and optical flow," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 07, 2020, pp. 10 713–10 720.
- [25] B. Wang, L. Li, Y. Nakashima, R. Kawasaki, H. Nagahara, and Y. Yagi, "Noisy-lstm: Improving temporal awareness for video semantic segmentation," *IEEE Access*, vol. 9, pp. 46 810–46 820, 2021.
- [26] Z. Zhao, Y. Jin, X. Gao, Q. Dou, and P.-A. Heng, "Learning motion flows for semi-supervised instrument segmentation from robotic surgical video," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2020, pp. 679–689.
- [27] Y. Jin, K. Cheng, Q. Dou, and P.-A. Heng, "Incorporating temporal prior from motion flow for instrument segmentation in minimally invasive surgery video," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 440–448.
- [28] O. Oktay, E. Ferrante, K. Kamnitsas, M. Heinrich, W. Bai, J. Caballero, S. A. Cook, A. De Marvao, T. Dawes, D. P. O'Regan *et al.*, "Anatomically constrained neural networks (acnns): application to cardiac image enhancement and segmentation," *IEEE transactions on medical imaging*, vol. 37, no. 2, pp. 384–395, 2017.
- [29] H. Ravishankar, R. Venkataramani, S. Thiruvankadam, P. Sudhakar, and V. Vaidya, "Learning and incorporating shape models for semantic segmentation," in *International conference on medical image computing and computer-assisted intervention*. Springer, 2017, pp. 203–211.
- [30] A. J. Larrazabal, C. Martínez, B. Glocker, and E. Ferrante, "Post-dae: Anatomically plausible segmentation via post-processing with denoising autoencoders," *IEEE Transactions on Medical Imaging*, vol. 39, no. 12, pp. 3813–3820, 2020.
- [31] A. J. Larrazabal, C. Martinez, and E. Ferrante, "Anatomical priors for image segmentation via post-processing with denoising autoencoders," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2019, pp. 585–593.
- [32] G. J. Brostow, J. Fauqueur, and R. Cipolla, "Semantic object classes in video: A high-definition ground truth database," *Pattern Recognition Letters*, vol. 30, no. 2, pp. 88–97, 2009.
- [33] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3213–3223.
- [34] M. Allan, A. Shvets, T. Kurmann, Z. Zhang, R. Duggal, Y.-H. Su, N. Rieke, I. Laina, N. Kalavakonda, S. Bodenstedt *et al.*, "2017 robotic instrument segmentation challenge," *arXiv preprint arXiv:1902.06426*, 2019.
- [35] Z.-L. Ni, G.-B. Bian, X.-H. Zhou, Z.-G. Hou, X.-L. Xie, C. Wang, Y.-J. Zhou, R.-Q. Li, and Z. Li, "Raunet: Residual attention u-net for semantic segmentation of cataract surgical instruments," in *International Conference on Neural Information Processing*. Springer, 2019, pp. 139–149.
- [36] M. Allan, S. Kondo, S. Bodenstedt, S. Leger, R. Kadkhodamohammadi, I. Luengo, F. Fuentes, E. Flouty, A. Mohammed, M. Pedersen *et al.*, "2018 robotic scene segmentation challenge," *arXiv preprint arXiv:2001.11190*, 2020.
- [37] 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 arXiv:2012.12453*, 2020.
- [38] A. P. Twinanda, S. Shehata, D. Mutter, J. Marescaux, M. De Mathelin, and N. Padoy, "Endonet: a deep architecture for recognition tasks on laparoscopic videos," *IEEE transactions on medical imaging*, vol. 36, no. 1, pp. 86–97, 2016.
- [39] M. Grammatikopoulou, E. Flouty, A. Kadkhodamohammadi, G. Quellec, A. Chow, J. Nehme, I. Luengo, and D. Stoyanov, "Cadis: Cataract dataset for surgical rgb-image segmentation," *Medical Image Analysis*, vol. 71, p. 102053, 2021.
- [40] P. Mascagni, A. Vardazaryan, D. Alapatt, T. Urade, T. Emre, C. Fiorillo, P. Pessaux, D. Mutter, J. Marescaux, G. Costamagna *et al.*, "Artificial intelligence for surgical safety: automatic assessment of the critical view of safety in laparoscopic cholecystectomy using deep learning," *Annals of Surgery*, 2021.
- [41] A. Madani, B. Namazi, M. S. Altieri, D. A. Hashimoto, A. M. Rivera, P. H. Pucher, A. Navarrete-Welton, G. Sankaranarayanan, L. M. Brunt, A. Okrainec *et al.*, "Artificial intelligence for intraoperative guidance: using semantic segmentation to identify surgical anatomy during laparoscopic cholecystectomy," *Annals of Surgery*, 2021.
- [42] L. M. Brunt, D. J. Deziel, D. A. Telem, S. M. Strasberg, R. Aggarwal, H. Asbun, J. Bonjer, M. McDonald, A. Alseidi, M. Ujiki *et al.*, "Safe cholecystectomy multi-society practice guideline and state of the art consensus conference on prevention of bile duct injury during cholecystectomy," *Annals of surgery*, vol. 272, no. 1, pp. 3–23, 2020.
- [43] C. I. Nwoye, C. Gonzalez, T. Yu, P. Mascagni, D. Mutter, J. Marescaux, and N. Padoy, "Recognition of instrument-tissue interactions in endoscopic videos via action triplets," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2020, pp. 364–374.
- [44] P. Mascagni, D. Alapatt, A. Garcia, N. Okamoto, A. Vardazaryan, G. Costamagna, B. Dallemagne, and N. Padoy, "Surgical data science for safe cholecystectomy: a protocol for segmentation of hepatocystic anatomy and assessment of the critical view of safety," *arXiv preprint arXiv:2106.10916*, 2021.
- [45] P. Vincent, H. Larochelle, I. Lajoie, Y. Bengio, P.-A. Manzagol, and L. Bottou, "Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion," *Journal of machine learning research*, vol. 11, no. 12, 2010.
- [46] A. Sharma, O. Grau, and M. Fritz, "Vconv-dae: Deep volumetric shape learning without object labels," in *European Conference on Computer Vision*. Springer, 2016, pp. 236–250.
- [47] S. Xingjian, Z. Chen, H. Wang, D.-Y. Yeung, W.-K. Wong, and W.-c. Woo, "Convolutional lstm network: A machine learning approach for precipitation nowcasting," in *Advances in neural information processing systems*, 2015, pp. 802–810.
- [48] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.
- [49] H. Al Hajj, M. Lamard, P.-H. Conze, S. Roychowdhury, X. Hu, G. Maršalkaitė, O. Zisimopoulos, M. A. Dedmari, F. Zhao, J. Prellberg *et al.*, "Cataracts: Challenge on automatic tool annotation for cataract surgery," *Medical image analysis*, vol. 52, pp. 24–41, 2019.
- [50] A. Pfeuffer, K. Schulz, and K. Dietmayer, "Semantic segmentation of video sequences with convolutional lstms," in *2019 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2019, pp. 1441–1447.
- [51] R. Sibeichi, O. Booi, N. Baka, and P. Bloem, "Exploiting temporality for semi-supervised video segmentation," in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019, pp. 0–0.