Exploring Segment-Level Semantics for Online Phase Recognition From Surgical Videos

Xinpeng Ding[®] and Xiaomeng Li[®], *Member, IEEE*

Abstract—Automatic surgical phase recognition plays a vital role in robot-assisted surgeries. Existing methods ignored a pivotal problem that surgical phases should be classified by learning segment-level semantics instead of solely relying on frame-wise information. This paper presents a segment-attentive hierarchical consistency network (SAHC) for surgical phase recognition from videos. The key idea is to extract hierarchical high-level semanticconsistent segments and use them to refine the erroneous predictions caused by ambiguous frames. To achieve it, we design a temporal hierarchical network to generate hierarchical high-level segments. Then, we introduce a hierarchical segment-frame attention module to capture relations between the low-level frames and high-level segments. By regularizing the predictions of frames and their corresponding segments via a consistency loss, the network can generate semantic-consistent segments and then rectify the misclassified predictions caused by ambiguous lowlevel frames. We validate SAHC on two public surgical video datasets, i.e., the M2CAI16 challenge dataset and the Cholec80 dataset. Experimental results show that our method outperforms previous state-of-the-arts and ablation studies prove the effectiveness of our proposed modules. Our code has been released at: https://github.com/xmedlab/SAHC.

Index Terms—Surgical video analysis, surgical phase recognition.

I. INTRODUCTION

COMPUTER-ASSISTED surgery systems can be used for pre-operative planning, surgical navigation and assist doctors in performing surgical procedures in modern operating rooms [1], [2]. Surgical phase recognition, one crucial and challenging task in computer-assisted surgery systems, aims to recognize the surgical activities from a surgery video [3]. It can monitor surgical processes and early alert the potential deviations and anomalies, advise the optimal arrangement, and provide clinical decision support [4]. Hence, developing an accurate online phase recognition algorithm, *i.e.*, predicts what

Xinpeng Ding is with the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong, SAR, China (e-mail: xdingaf@connect.ust.hk).

Xiaomeng Li is with the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong, SAR, China, and also with The Hong Kong University of Science and Technology Shenzhen Research Institute, Shenzhen 518057, China (e-mail: eexmli@ust.hk).

Digital Object Identifier 10.1109/TMI.2022.3182995

phase is occurring at the current frame without knowing the future information from surgery videos, is highly demanded in clinical practice.

One of the most critical challenges in automatically recognizing the surgical phase from videos is that the complicated surgical scenes usually have limited inter-phase variance and high intra-phase variance. As shown in Fig. 1(a), two frames from two different surgical phases, *e.g.*, "GallbladderPackaging" and "GallbladderRetraction", has a high visual similarity, which may easily lead to misclassification of the surgical phases. Hence, learning the long-range temporal dynamics by observing the neighboring frames is the key solution to this problem.

Recent studies solved this issue by capturing the long-term frame-wise relation between the current frame and previous frames [5], [6]. For example, Jin *et al.* [5] introduced TMRNet, which contains a memory bank to store long-range information to learn the relation between the current frame and previous frames. Another category of research [7], [8] employed a multi-stage architecture, including a predictor stage to generate a frame-wise prediction and a refinement stage to refine the previous prediction. For example, Czempiel *et al.* [7] introduced MS-TCN [9] into surgical phase recognition, and used causal temporal convolutional networks [10] for online prediction. To solve the insufficient training of the refinement stage, Yi *et al.* [8] proposed a non-end-to-end stage to train the refinement stage separately.

Although these methods attempted to learn long-range temporal dynamics, they generated phase predictions by learning inter-frame relations in a low-level fine-grained way. In this paper, we aim to explore that: would the frame-level phase recognition be improved by incorporating segment-level modeling? As shown in Fig. 1(b), two ambiguous frames (solid red boxes) are too similar, making the model hard to distinguish between two surgical phases. In contrast, their corresponding high-level segment (dashed boxes) can present discriminative semantics for surgical activities, contributing to the better recognition of surgical phase; see Fig. 5 for details.

To this end, we present a segment-attentive hierarchical consistency network (SAHC) for online phase recognition from surgical videos. *The key idea is to learn high-level segments from surgical videos and then adopt them to rectify the ambiguous semantics caused by low-level frames.* To achieve it, we first design a hierarchical temporal feature extractor to generate high-level segments by capturing the feature of the current frame and its multi-scale neighboring frames. Then, to rectify the ambiguous information in low-level frames, we develop a hierarchical segment-frame attention to capture the relations between the high-level temporal segments and

1558-254X © 2022 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

Manuscript received 30 March 2022; revised 30 May 2022; accepted 8 June 2022. Date of publication 14 June 2022; date of current version 27 October 2022. This work was supported in part by the Research Grant from Shenzhen Municipal Central Government Guides Local Science and Technology Development Special Funded Projects under Grant 2021Szvup139 and in part by a Research Grant from HKUST-BICI Exploratory Fund under Grant HCIC-004. (*Corresponding author: Xiaomeng Li.*)



Fig. 1. (a) Ambiguous frame information would make the model generate erroneous predictions for surgical phase recognition. (b) Segmentlevel difference. Compared to the frame-level information, segment-level information can present more discriminative features for surgical phase recognition.

low-level frames. By enhancing the consistency of the predictions from frames and their neighboring segments, we find that the features of ambiguous frames and their corresponding high-level segments would be pulled together, resulting in the effective refinement of the frame-wise ambiguity.

This paper has the following contributions:

- We introduce the importance of high-level segment information for surgical phase recognition, and leverage its semantics to refine the ambiguous low-level frames.
- We present a segment-attentive hierarchical consistency network (SAHC), which generates high-level semanticconsistent segments and then rectifies prediction errors via a segment-attentive module.
- We propose a hierarchical segment-frame attention module to learn relationships between frames and segments.
- Experiments on two public surgical phase recognition datasets show that our method achieves a significant improvement over the prior art, *e.g.*, over 3.8% on the M2CAI16 dataset.

II. RELATED WORK

A. Surgical Phase Recognition From Surgical Video

Early work for surgical phase recognition from surgical videos is mainly based on hand-crafted features, such as pixel values and intensity gradients [11], spatial-temporal features [12] and features consisting of color, texture, and shape [13]. Concurrently, there are also other works that using linear statistical models to capture the temporal information of surgical videos, *e.g.*, left-right HMM [13], Hidden semi-Markov Model [14], hierarchical HMM [3], conditional Random Fields [15]–[17] and Dynamic Time Warping [11]. However, their performance is limited by the empirically designed low-level features. In recent years, neural networks have extracted spatial and temporal features for surgical phase recognition from surgical videos. These methods can be broadly classified into two categories.

One category aims at modeling spatial and temporal features with frame-wise labels. For example, Twinanda et al. [3] employed ResNet [18] to extract video-level features and demonstrated its effectiveness for surgical phase recognition. Jin et al. [19] introduced SV-RCNet, a unified framework that integrates ResNet and LSTM module to sequentially learn spatial-temporal features for surgical phase recognition. To effectively capture the long-range temporal dynamics, Jin et al. [5] developed TMRNet, which consists of a memory bank to fuse long-range and multi-scale temporal features for surgical phase recognition. In addition to learning temporal information, Gao et al. [6] used a hybrid embedding aggregation transformer to enhance the importance of spatial features for phase recognition. Some work employed multistage architecture, e.g., a predictor stage and a refinement stage, in surgical phase recognition, where the misclassification in the predictor stage can be well rectified during the refinement stage. For example, TeCNO [7] adapted the multistage temporal convolution network (MS-TCN) [9], [20] to online surgery scenario via causal and dilated convolutions. Yi et al. [8] found that directly using MS-TCN brings little improvement and then proposed a new non end-to-end training strategy.

Another category leverages additional information, e.g., designing multi-task learning to improve the performance for surgical phase recognition. For instance, Twinanda et al. [3] trained a shared work for feature extraction in a multi-task way consisting of tool presence detection and phase recognition. Zisimopoulos et al. [21] used a ResNet [18] to predict the binary predictions for the tool presence and then combine the predictions and features for phase recognition. MTRCNet-CL [22] introduced a novel correlation loss to explicitly model the relations between tool presence and phase classification. In addition to the tool information, Nakawala et al. [23] leveraged more cues, such as management tools, ontology and production rules to improve the performance. Some works [24] also extracted the optical flows and leverage the motion information to enhance the learning of the model. These methods suffer from extra annotation cost for multitasking or bring additional computation overhead to obtain other modalities, e.g., optical flows.

Our method belongs to the first category. Most existing methods aim to capture frame-wise relations by learning temporal dynamics while ignoring the high-level semantic information and its influences for low-level frames. There are some works [16] using conditional Random Fields to combine frame-level information for high-level surgical tasks. However, in surgical phase recognition, we need detailed frame-wise classification. In this paper, we combine the high segment-level and low frame-level semantics for surgical phase recognition. By the proposed segment-frame attention, our model uses high-level semantics to refine low-level errors to improve the frame-level performance.

B. Temporal Pyramid Learning in Videos

Our method is also related to the temporal pyramid learning in video action recognition tasks. Temporal pyramid methods



Fig. 2. Illustration of our proposed Segment-Attentive Hierarchical Consistency Network (SAHC). The frame-wise feature **H** is extracted after feeding video frames into the frame-wise feature encoder. Then, several stacked residual causal dilated temporal convolution layers (RCDL), defined in Section III-A and Fig. 3(b), takes **H** as the input to obtain \mathbf{F}^0 . Since the frames in different phases show very similar appearance, *i.e.*, ambiguous (see red hollow boxes), the model would extract some ambiguous features (*i.e.*, red solid circles), which would result in ambiguous predictions (see red solid boxes). To alleviate this, based on the frame features, we introduce *M* segment-wise feature extractors (SFE), defined in Section III-B.1 and Fig. 3, to extract a set of hierarchical segment features { $\mathbf{F}^i_{j=1}^M$, which contains more discriminative information (see green and orange triangles). After adding the positional encoding (PE) to { $\mathbf{F}^i_{j=1}^M$, we obtain { \mathbf{F}^i_{pos} } and then use a hierarchical segment-frame attention (SFA), defined in Section III-C and Fig. 4, to capture the relations between frames and segments. Finally, we regularize the network to encourage all prediction from low-level to high-level, *i.e.*, { \hat{y}_p^i }, to be consistent, which uses correct high-level predictions to refine the low-level ones.

aim to process the variant duration of actions by extracting multi-scale temporal information from videos.

Early approaches generated a fixed multi-scale sliding windows, which act as proposals for temporal action localization [25]–[27] or video grounding [28]–[31]. Recently, researchers sampled frames at different temporal rates to construct an input-level frame pyramid. And frames in each level of the pyramid were extracted by separate networks to obtain the corresponding mid-level features, which were then fused for final prediction [32], [33]. However, these methods required the additional networks, which may be computationally expensive [34]. Motivated by the feature pyramid network (FPN) [35], Yang *et al.* [34] captured visual tempos in multi-scale feature levels with only a single input. Specifically, they utilized a feature pyramid network that temporally downsamples features to obtain different temporal scales.

Compared with existing temporal pyramid networks [31], [34] for video analysis in computer vision, our method has the following differences. (1) Different video tasks and bottlenecks. Existing methods [31], [34] proposed to learn multi-temporal scales to tackle the variant duration of the action instances for natural video action recognition. In contrast, our goal is to use segment-level information to *refine the erroneous predictions caused by ambiguous frame-level information in surgical videos*. (2) Network design with consistency regularization. We design a hierarchical network with consistency regularization to rectify the erroneous prediction caused by ambiguous frame-level information. (3) Attention. We further introduce an attention module to capture the relationship between frame- and segment-level information. These two contributions do not exist in current temporal pyramid methods.

C. Long-Term Video Understanding

To capture long-term information of videos, Farha et al. [9], [20] propose a multi-stage temporal convolution network (MS-TCN) for the temporal action segmentation task, which enlarges the receptive fields to capture long-range temporal information by cascaded dilated 1D convolutions. To reduce the annotation cost in long videos, Zhukov et al. [36] introduces a long-range temporal order verification to to isolate actions from their background in a self-supervised manner. Wu et al. [37] propose a long-term feature bank to contain supportive information from the entire video, which augments the state-of-the-art video models that otherwise would only view short clips of 2-5 seconds. Recently, Object Transformer [38] is proposed to use transformer [39] to model the long-term relations. In this paper, we follow previous works [7], [8], [40] to use cascaded TCN to capture long-range information for surgical videos.

III. METHODOLOGY

Fig. 2 illustrates our proposed segment-attentive hierarchical consistency network (SAHC). The video frames are firstly fed into (a) a spatial-temporal feature extractor to obtain the features with temporal frame-level information, followed by (**b**) a segment-level hierarchical network consists of three segment-wise feature extractors (SFE) to capture the multi-scale segment-level semantics. After that, we introduce (**c**) a segment-frame attention (SFA) module to jointly learn the relations between frames and segments. Finally, (**d**) the segment-frame hierarchical consistency loss is developed to enhance the consistent prediction from low-level frames and their neighbouring segments, such that the erroneous frame-wise prediction can be rectified by segments. In the following sections, we describe our method in detail.

A. Spatial-Temporal Feature Extractor

We denote a video as $\mathbf{V} = {\{\mathbf{v}_t\}}_{t=0}^{T-1}$, where *T* is the number of frames and $\mathbf{v}_t \in \mathbb{R}^{H \times W \times 3}$ is a frame with height *H*, weight *W* and three channels. Let *C* denote a set of surgical phases, where $C = \{0, ..., C - 1\}$ and *C* is the number of phase categories. Our goal is to learn a deep network f_{θ} that maps the input \mathbf{v}_t to a phase label $y_t \in \mathbb{R}^C$, which is a one-hot vector of phase label $c \in C$.

In order to capture the spatial-temporal information of the videos, we first extract the frame-wise feature, and then use the temporal convolutional network to model their temporal relations. Specifically, we first feed the frames **V** into the the spatial encoder, *i.e.*, ResNet-50 [18], to extract the spatial feature of each frame, denoted as $\mathbf{H} = {\{\mathbf{h}_t\}}_{t=0}^{T-1}$, where $\mathbf{h}_t \in \mathbb{R}^D$ is the feature of the \mathbf{v}_t and D the dimension of the feature.

Then, we feed $\mathbf{H} \in \mathbb{R}^{T \times D}$ into several stacked residual causal dilated temporal convolution layers (RCDL) (shown in Fig. 3(b)) to capture the frame-wise relation and obtain the corresponding features $\mathbf{F}^0 \in \mathbb{R}^{T \times D}$. The operations of each RCDL can be formally as follows:

$$\mathbf{Z}_{l} = ReLU\left(\mathbf{W}_{1,l} * \mathbf{F}_{l-1} + \mathbf{b}_{1,l}\right), \qquad (1)$$

$$\mathbf{F}_{l} = \mathbf{F}_{l-1} + \mathbf{W}_{2,l} * \mathbf{Z}_{l} + \mathbf{b}_{2,l}, \qquad (2)$$

where \mathbf{F}_l is the output of the layer l, * denotes the convolution operator, $\mathbf{W}_{1,l}$ is the dilated 1D convolution kernel [7], $\mathbf{W}_{2,l}$ is the weights of a 1 × 1 convolution and $\mathbf{b}_{1,l}$, $\mathbf{b}_{2,l}$ are bias vectors. The dimension of \mathbf{F}^i is set to 64 in all RCDL. Following [7], to predict the label of frame \mathbf{f}_t , casual dilated convolution only relies on the current and previous frames, *i.e.*, ($\mathbf{f}_{t-n}, \ldots, \mathbf{f}_t$), which allows for the online recognition. Before the residual addition in Eq. 2, we adopt the dropout [41] to avoid over-fitting. In this paper, we set L to be 11 as same in [20]. To achieve online recognition of surgical activities, we use the casual dilated convolution that only relies on the current and previous frames, *i.e.*, ($\mathbf{f}_{t-n}, \ldots, \mathbf{f}_t$).

B. Segment-Level Hierarchical Network

1) Segment-Level Feature Extractor (SFE): After obtaining the frame-level temporal features \mathbf{F}^0 , we then use them to extract the information of segments by a segment-level feature extractor. Each segment-level feature extractor consists of a temporal fusion layer (see Fig. 3(a)), followed by several *L* stacked RCDL (see details in Section III-A and Fig. 3(b)). The goal of the temporal fusion layer is to aggregate the features of a frame and its neighbouring frames to obtain



Fig. 3. Illustration of SFE. It consists of two main components: (a) the temporal fusion layer and (b) L stacked residual causal dilated temporal convolution layers (RCDL). After feeding low-level features into the temporal fusion layer, we obtain the high-level features. Then, we use RCDL to capture the relations between high-level features.

their corresponding segment features, which can be formally defined as:

$$\mathbf{f}_p^s = \mathrm{TFL}(\mathbf{f}_t^0, \dots, \mathbf{f}_{t+k}^0), \tag{3}$$

where TFL refers to the operation of the temporal fusion layer, \mathbf{f}_p^s is the segment feature and k indicates the number of frames that each segment contains. We have several choices to devise the temporal fusion layer, *e.g.*, a convolution layer, a max-pooling layer or an average pooling layer with kernel size and stride of k. For examples shown in Fig. 3, given a sequence low-level input features $\{\mathbf{f}_i\}_{i=0}^5$ and a temporal fusion layer with kernel size and stride of 2, we will then obtain the high-level output features of $\{\mathbf{f}_i^s\}_{i=0}^2$. The segment feature \mathbf{f}_0^s aggregates \mathbf{f}_0 and \mathbf{f}_1 , and \mathbf{f}_1^s aggregates \mathbf{f}_2 and \mathbf{f}_3 , and so on. Hence, after feeding \mathbf{F}^0 into the temporal fusion layer,

Hence, after feeding \mathbf{F}^0 into the temporal fusion layer, we obtain the segment-level features, defined as $\mathbf{F}^s = {\{\mathbf{f}_p^s\}}_{p=0}^{T^s-1} \in \mathbb{R}^{T^s \times D}$. $T^s = \lfloor T/k \rfloor$ and $\lfloor \cdot \rfloor$ indicates the round down. We can achieve this by several methods, *e.g.*, the convolution, the max-pooling and the average-pooling with setting both the kernel size and stride to *k*. We will discuss these different kinds of temporal fusion layers in Section IV-C.4. Similar in frame-level, we also expect the model to capture the temporal relations between segments. To this end, we input \mathbf{F}^s into RCDL to use temporal convolution to capture relations in \mathbf{F}^s and finally obtain the output \mathbf{F}^1 .

The implementation of our temporal fusion layer, *i.e.*, fusing frames to obtain the segments, may generate the ambiguous segments, which contains frames belonging to different phases. In surgical videos, the category of frames in the same phases are consistent. Hence, only segments in the boundaries



Fig. 4. Illustration of (a) the hierarchical SFA (described in Fig. 2(c)) and (b) the transformer (Trans) layer.

between two phases would contain frames with different class labels. That is to say, the number of ambiguous segments is too few and can be ignored.

2) Hierarchical Network: The duration of surgical videos and phases is various [5]; hence we need to extract the multi-scale temporal information for segments. To achieve it, we introduce a hierarchical network, consisting of a sequence of segment-level feature extractor modules, to obtain a set of features, *i.e.*, $\{\mathbf{F}^i \in \mathbb{R}^{T_i \times D}\}_{i=1}^M$, where *M* is the number of different temporal scales. In this paper, we set M = 3 empirically. \mathbf{F}^i is the output of the *i*-th segment-level feature extractor sequence, and T^i is its temporal duration and $T^{i+1} = |T^i/k|$.

C. Segment-Frame Attention (SFA) Module

The proposed RCDL in Section III-B.1 can only model the relations in each scale, *i.e.*, the relations in \mathbf{F}^i . To use the high-level segment information to refine the erroneous predictions in low-level frames, we need to capture the relations between features in different scales, *e.g.*, frames \mathbf{F}^0 and its corresponding segments \mathbf{F}^1 and \mathbf{F}^2 and \mathbf{F}^3 , respectively. Recently, Transformer [39], a kind of attention layer, shows promising results in learning attentive weights/relationships with applications in images, words, and videos [39], [42]–[48]. Here, we adopt the transformer to learn the relationships between segments and frames.

To this end, we design a segment-frame attention (SFA) module, where frames are used as the queries, and hierarchical segments are served as keys and values. In this way, the query, *i.e.*, a frame, can find its relation with segments overall the video across different phases, and use the features of its high-level segments to refine the ambiguous frames, *i.e.*, pushing the frame-level features to be close with the segment-level ones. See qualitative analysis in Fig. 6. Since there is no frame sequence information in the attention module, we need to embed position encoding additionally, which is formulated as:

$$\mathbf{F}^{i}_{pos} = \mathbf{F}^{i} + \mathbf{E}, \quad i \in \{0, \dots, M\}$$
(4)

where $\mathbf{E} \in \mathbb{R}^{T \times D}$ is the learned positional embedding, *M* is the number of different temporal scales (defined in Section III-B.2) and we set M = 3 empirically. Fig. 4(a) shows the details of the hierarchical SFA. Concretely, \mathbf{F}_{pos}^{0} and $\{\mathbf{F}_{pos}^{i}\}_{i=1}^{3}$ are fed into three shared transformer layers (Trans layer) and then generate a set of outputs, defined as $\{\hat{\mathbf{F}}^{i}\}_{i=1}^{3}$. We finally concatenate the multi-scale aggregated features $\{\hat{\mathbf{F}}^{i}\}_{i=1}^{3}$ to obtain the final feature $\overline{\mathbf{F}} = \{\overline{\mathbf{f}}_{t}\}_{t=0}^{T-1}$. The details of the Trans layer are shown in Fig. 4(b). The core component of the Trans layer is the multi-head attention, and the intuitive idea is that each token can interact with other tokens and can learn to aggregate useful semantics more effectively. Given \mathbf{F}_{pos}^{0} and \mathbf{F}_{pos}^{i} , each head of the multi-head attention can be formally defined as follows:

Attention
$$\left(\mathbf{Q}; \mathbf{K}^{i}; \mathbf{V}^{i}\right) = \operatorname{softmax}\left(\frac{\mathbf{Q}{\mathbf{K}^{i}}}{\sqrt{D}}\right) \mathbf{V}^{i}$$
 (5)

where $\mathbf{Q} = \mathbf{F}_{pos}^{0} \mathbf{W}_{i}^{q}$, $\mathbf{K}^{i} = \mathbf{F}_{pos}^{i} \mathbf{W}_{i}^{k}$ and $\mathbf{V}^{i} = \mathbf{F}_{pos}^{i} \mathbf{W}_{i}^{v}$ are linear layers, and \mathbf{W}_{i}^{q} , \mathbf{W}_{i}^{k} , $\mathbf{W}_{i}^{v} \in \mathbb{R}^{D \times \frac{D}{N_{head}}}$. N_{head} is the number of heads and \sqrt{D} controls the effect of growing magnitude of dot-product with larger D [39]. Subsequently, the output of all heads are concatenated and fed into a linear layer, followed by a layer normalization [49]. Then, it is followed by a feed-forward network with ReLU activation. The residual connection [18] and layer normalization [49] are also applied as in the multi-head attention. Finally, we obtain the output of each Trans layer, $\hat{\mathbf{F}}^{i}$. We then concatenate the multi-scale aggregated features $\{\hat{\mathbf{F}}^{i}\}_{i=1}^{3}$ to obtain the final feature $\overline{\mathbf{F}} = \{\overline{\mathbf{f}}_{i}\}_{i=0}^{T-1}$.

D. Consistency Between Frames and Segments

Based on the final obtained features of frames \mathbf{F}^0 and multi-scale segments $\{\mathbf{F}^i\}_{i=1}^3$, we get their corresponding predictions \hat{y}_t^0 and $\hat{\mathcal{Y}}_t = \{\hat{y}_p^i\}_{i=1}^3$ by a shared prediction network, where $\{\hat{y}_t^0\} \in \mathbb{R}^{C \times T}$ and $\{\hat{y}_p^i \in \mathbb{R}^{C \times T^i}\}_{i=1}^3$. The shared prediction network is a convolution layer. The kernel size, stride, input dimension and the output dimension are 1, 1, 256, 7, respectively. The network simply follows previous state-of-the-art methods (Yi and Jiang 2021; Czempiel *et al.* 2020). Then, the classification losses of the frames (\mathcal{L}_{frame}) and their corresponding neighbouring segments ($\mathcal{L}_{segment}$) can be defined as follows:

$$\mathcal{L}_{frame} = -\frac{1}{T} \sum_{t=0}^{T-1} \sum_{c=0}^{C-1} y_{t,c}^0 \log(\hat{y}_{t,c}^0), \tag{6}$$

$$\mathcal{L}_{segment} = -\frac{\beta}{T^{i}} \sum_{i=1}^{3} \sum_{p=0}^{T^{i-1}} \sum_{c=0}^{C-1} y^{i}_{p,c} \log(\hat{y}^{i}_{p,c}), \qquad (7)$$

where $y_{p,c}^i \in \mathbb{R}^{C \times T^i}$ is the corresponding ground truth at *i*-th scale, which is generated by applying down-sampling to $y_{t,c}$. β is the hyper-parameter to control the weight of the segment-wise prediction at different scales. The multi-scale consistency loss can be the combination of \mathcal{L}_{frame} and $\mathcal{L}_{segment}$:

$$\mathcal{L}_{msc} = \mathcal{L}_{frame} + \mathcal{L}_{segment}.$$
 (8)

In this way, the model would encourage the consistency of the prediction of the frames and their segments at multiscale tempos. For predicting more smoothing results, a mean squared error over the classification probabilities of every two adjacent frames are used [9]:

$$\mathcal{L}_{smooth} = \frac{1}{TC} \sum_{c=0}^{C-1} \sum_{t=1}^{T-1} \left| \hat{y}_{t,c}^{0} - \hat{y}_{t+1,c}^{0} \right|^{2} + \frac{\beta}{T^{i}C} \sum_{i=1}^{3} \sum_{p=0}^{T^{i}-1} \sum_{c=0}^{C-1} \left| \hat{y}_{p,c}^{i} - \hat{y}_{p+1,c}^{i} \right|^{2}.$$
 (9)

Hence, the overall of the objective of our SAHC is:

$$\mathcal{L} = \mathcal{L}_{msc} + \lambda \mathcal{L}_{smooth},\tag{10}$$

where λ is a hyper-parameter to control the weight between two losses. We will discuss the effect of α and λ in Experiments.

IV. EXPERIMENTAL RESULTS

In this section, we evaluate our proposed SAHC on two benchmark datasets for surgical phase recognition.

A. Datasets, Metrics and Implementation Details

1) Datasets:

a) M2CAI16: The M2CAI16 dataset [51] consists of 41 laparoscopic videos at 25 fps and each frame is with a resolution of 1920 × 1080. The videos contain 8 surgical phases, *i.e.*, TrocarPlacement, Preparation, CalotTriangleDissection, ClippingCutting, GallbladderDissection, GallbladderRetraction, CleaningCoagulation, GallbladderPackaging, labeled by professional surgeons. More detailed information can be found in [19]. We use 27 videos for training, 7 videos for validation and 14 videos for testing, following the same evaluation protocols in [5]. Note that the validation set is used for model and hyper-parameters selection. All videos are sub-sampled to 5 fps. All frames of videos belong to one and only one phase label, and there are not any background frames.

b) Cholec80: The Cholec80 dataset [3] contains 80 videos of cholecystectomy surgeries, which are are recorded at 25 fps and are annotated into 7 surgical phases: Preparation, CalotTriangleDissection, ClippingCutting, GallbladderDissection, GallbladderRetraction, CleaningCoagulation, Gallblad*derPackaging*. The resolution of each frame is 1920×1080 or 854×480 . We split the dataset into 40 training videos, 8 validation videos and 40 testing videos, following the same setting in prior methods [5]. Similar as M2CAI16, we also use the validation set for model and hyper-parameters selection. All videos of both M2CAI16 and Cholec80 datasets are sub-sampled into 5 fps. All frames in two datasets belong to one and only one phase label, and there are not any background frames. The datasets (M2CAI16 and Cholec80) we employed in this work are the largest public benchmark datasets, which have been widely used in many published surgical workflow recognition papers [5], [7], [8], [19], [22]. The datasets we used are very challenging and hard to be



Fig. 5. Effect of learning with segment-level semantics. We visualize results, *i.e.*, Accuracy, precision, recall and jaccard (see definition in Section IV-A), on baseline model with different temporal scale, *i.e.*, 1, 3, 5, on the M2CAI16 dataset [51]. It is clear that segments with more frames achieve higher performance.

overfitting. The reasons are shown as follows: (a) The duration of each video is very long (around 30 minutes), and each of the dataset has around more than 170K frames, which is hard to over-fitting (note that our task is a frame-wise classification). (b) The variance for videos and different phases is very large. For example, the Cholec80 datasets is collected by 13 surgeons which is very diverse. Furthermore, the standard deviation of duration for phase "Gallbladder dissection" is 551s. Each video may not have all phases and there is no obvious prior between two phases.

2) Evaluation Metrics: We employ four commonly-used metrics, *i.e.*, accuracy (AC), precision (PR), recall (RE), and Jaccard (JA) to evaluate the phase prediction accuracy. AC represents a video-level evaluation, which is defined as the percentage of correctly classified frames in the entire video. Due to the imbalanced phases presented in videos, PR, RE, and JA refer to the phase-level evaluation, which is evaluated within each phase and then averaged over all the phases. Specifically, we first compute PR, RE and JA of each phase by $\frac{|GT \cap P|}{|P|}$, $\frac{|GT \cap P|}{|GT|}$ and $\frac{|GT \cap P|}{|GT \cup P|}$, where GT and P refer to the ground-truth and prediction set, respectively. Then, we obtain the mean and the standard deviation of these scores over all phases and obtain the performance of the entire video. The evaluation protocols are the same with previous methods [5], [7], [8], [19], [22].

3) Implementation Details: The model is built with Pytorch [52] and is trained by 1 NVIDIA 3090 GPU. We use Adam [53] optimizer with the learning rate of 5×10^{-4} , decayed by 30 epochs, and The model is totally trained for 100 epochs. We select the model with the highest performance on the validation, and reports its results on the test set. The number heads of segment-frame attention N_{head} is set to 4 and the positional encoding is learned, due to the variant duration of videos. In the frame-wise feature encoder, the number of the residual causal dilated temporal convolution layers (RCDL) is set to 11. In the segment-level hierarchical network, we set the number of RCDL to be 10. The dimension of the features, i.e., D, is set to 64. We use the max-pooling as the temporal fusion layer, and we will evaluate the effect of different methods, e.g., average-pooling or convolutions in Section IV-C.4. We set the sizes of the kernel and the stride k of the temporal fusion layers are both set to 7, which achieves the best performance. There

 TABLE I

 COMPARISONS WITH THE STATE-OF-THE-ARTS ON M2CAI16 DATASET

Methods	Accuracy	Precision	Recall	Jaccard
PhaseNet [50]	79.5 ± 12.1	-	-	64.1 ± 10.3
SV-RCNet [19]	81.7 ± 8.1	81.0 ± 8.3	81.6 ± 7.2	65.4 ± 8.9
OHFM [40]	85.2 ± 7.5	-	-	68.8 ± 10.5
TMRNet [5]	87.0 ± 8.6	87.8 ± 6.9	88.4 ± 5.3	75.1 ± 6.9
Not-End [8]	84.1 ± 9.6	-	88.3 ± 9.6	69.8 ± 10.7
Trans-SVNet [6]	87.2 ± 9.3	88.0 ± 6.7	87.5 ± 5.5	74.7 \pm 7.7
Ours	91.6 ± 7.8	$93.5~\pm~5.6$	92.9 ± 6.3	85.7 ± 7.7

TABLE II COMPARISONS WITH THE STATE-OF-THE-ARTS ON CHOLEC80 DATASET

Methods	Accuracy	Precision	Recall	Jaccard
PhaseNet [50]	78.8 ± 4.7	71.3 ± 15.6	76.6 ± 16.6	-
SV-RCNet [19]	85.3 ± 7.3	80.7 ± 7.0	83.5 ± 7.5	-
OHFM [40]	87.3 ± 5.7	-	-	67.0 ± 13.4
TeCNO [7]	88.6 ± 2.7	-	85.2 ± 10.6	-
TMRNet [5]	90.1 ± 7.6	90.3 ± 3.3	89.5 ± 5.0	79.1 ± 5.7
Not-End [8]	88.8 ± 7.1	-	84.9 ± 7.2	73.2 ± 9.8
Trans-SVNet [6]	90.3 ± 7.1	90.7 ± 5.0	88.8 ± 7.4	79.3 ± 6.6
Ours	91.8 ± 8.1	90.3 ± 6.4	$90.0~\pm~6.4$	$81.2~\pm~5.5$

is no overlapping when generating the segment-level features to avoid the same segment contains the frames from different phases. The ablation study of k is conducted in Section IV-C.4. We set α and β to be 1 for both M2CAI16 and Cholec80, the detailed analysis is shown in Fig. 7.

B. Comparison With the State-of-the-Arts

As shown in Table I and Table II, we compare the proposed SAHC with the state-of-the-art approaches on the M2CAI16 and Cholec80 dataset, including PhaseNet [50], SV-RCNet [19], OHFM [40], TMRNet [5], Trans-SVNet [6], TeCNO [7]. Note that "ours" in Table I and Table II indicate the model described in Fig. 2, *i.e.*, using three scales with the segment-frame attention, and the implementation of the temporal fusion layer is max-pooling. Our method achieves 4.8% and 3.0% improvements over the prior state-of-the-arts on the M2CAI16 dataset and the Cholec80 dataset, respectively. Notably, the improvements of our method are more significant in M2CAI16 than in Cholec80. This is because M2CAI16 contains more ambiguous frames, as shown in Fig. 1(b), demonstrating the effectiveness of our method to address the main challenge in surgical video recognition.

As described in Section IV-A, higher accuracy represents higher frame-level performance, while higher precision, recall, and Jaccard indicates higher phase-level accuracy. Since the imbalanced phases shown in videos, the phase-level performance is more reasonable for surgical video phase recognition. Among three phase-level metrics, compared with precision $(\frac{|GT \cap P|}{|P|})$ and recall $(\frac{|GT \cap P|}{|GT|})$, Jaccard $(\frac{|GT \cap P|}{|GT \cup P|})$ measures how close the predictions of phase set to the ground-truth set are, which is more accurate and stricter. From Table I and Table II, we find that our method can achieve higher improvement of Jaccard than other metrics, which indicates that our model can produce less misclassification prediction. As shown in Fig. 8(c). "GT" indicates ground-truth, "Base" refers to the baseline model. The predictions of the baseline and our method have the similar Accuracy score. Compared to the

TABLE III

Ablation Study on the Segment-Level Information and the Hierarchical Consistency Network on M2CAI16 Dataset

Methods	Accuracy	Precision	Recall	Jaccard
Base	87.1 ± 9.0	87.7 ± 7.1	88.8 ± 6.3	76.6 ± 8.9
$\{\mathbf{F}^1\}$	87.8 ± 7.3	89.3 ± 6.5	89.2 ± 5.5	82.5 ± 6.8
$\{\mathbf{F}^1,\mathbf{F}^2\}$	88.6 ± 7.8	90.9 ± 6.3	90.5 ± 5.2	84.2 ± 6.8
$\{{f F}^1,{f F}^2,{f F}^3\}$	90.2 ± 8.8	92.4 ± 6.2	91.3 ± 5.5	84.8 ± 6.8

baseline, the prediction of our method shows a higher Jaccard score, demonstrating that our model has a good capacity in refining erroneous predictions caused by ambiguous frames within a surgical phase.

C. Ablation Study

We ablate the proposed SAHC to evaluate the effectiveness of each component and analyze why they work. Our baseline model is the network in Fig. 2 without generating segment-level information and hierarchical segment-frame attention, denoted as "Base" for simplification. In other words, all output of the baseline model is at the same temporal scale, *i.e.*, frame-level. Note that all models in this section are optimized and trained on the training dataset with corresponding losses independently.

1) Effect of Learning With Segment-Level Semantics: In this ablation study, we prove that high-level segment feature outperforms than the frame-level one. As shown in Fig. 5, the performance of the model trained with segments, *i.e.*, "3-scale" and "5-scale", outperforms that of the model trained with frames, *i.e.*, "1-scale", on both four metrics.

2) Hierarchical Segment Information: Table III shows the importance of the segment-level information and the hierarchical consistency network on M2CAI16. $\{a, b, c\}$ indicates the consistency with **a**, **b** and **c**. Compared "Base" and " $\{\mathbf{F}^1\}$ ", it is clear that the segment-level prediction can improve the performance from 87.1% to 87.8% on accuracy. We also find that the hierarchical consistency between segments and frames can achieve considerable improvements. For example, with an additional scale of segment, e.g., \mathbf{F}^2 , the model achieves 1.5% improvements in terms of accuracy. With three-scale segments, the model reaches the best performance, *i.e.*, 90.2% accuracy. More scales of segments are not feasible in our dataset since the temporal size at \mathbf{F}^3 is too small to be downsampled. More specifically, the average length of training videos in M2CAI16 is 2502. After three times downsampling with k = 11, the length of the final feature is only 1.8, which is too small to cover whole phases in the video. Note that the number of phases in a video is generally from 3 to 6. Furthermore, in order to prove that the improvement comes from the proposed hierarchical segment information, we also compare the number of parameters, computation cost and running time of our method with that of the baseline model, which is reported in Table IV. "Param" indicates the number of parameters. It is clear that our proposed method outperforms the baseline model with a clear margin, i.e., 4.7%, while having the same number of parameters, i.e., 26.8M. Compared with the baseline, our model would bring extra 0.23G FLOPS.



Fig. 6. (a) Attention weight of the segment-frame attention (SFA) module. The horizontal axis represents the frame features, and the vertical axis represents the segment features. In this figure, there are a total of 12 frames (*i.e.*, $\{f_i\}_{i=1}^{12}$) from four different categories of phases. Different colors indicate different phases, *e.g.*, $\{f_i\}_{i=1}^{3}$ belong to the first phase (the blue one). The segment-level features (triangles with different colors) are generated by feeding three frames into SFE (defined in Section III-B.1), *i.e.*, $S_i = SFE(\{f_{3*i-2}, f_{3*i-1}, f_{3*i}\})$). It is clear that the segment representations show the high similarity with the frame ones that sharing the same phase class. (b) The cosine similarity matrices of features of frames. The horizontal and vertical axis are both represent the frame features. The features in the same phase become more similar after the refinement of our method. The bi-direction arrows with different colors indicate the different phases.

TABLE IV

COMPARISON OF NUMBER OF PARAMETERS, COMPUTATION COST AND RUNNING TIME OF DIFFERENT MODELS. "PARAM" INDICATES THE NUMBER OF PARAMETERS. "W/O SFA" INDICATES OUR SEGMENT-LEVEL MODEL WITHOUT USING SFA

Methods	Accuracy	Param	GFLOPS	Running time
Base	87.1 ± 9.0	26.8M	1.60G	0.27s
Ours w/o SFA	90.2 ± 8.8	26.8M	1.83G	0.29s
Ours	91.8 \pm 8.1	26.9M	2.02G	0.31s

TABLE V Ablation Study on the Hierarchical Segment-Frame Attention (SFA) on M2CAI16

Methods	Accuracy	Precision	Recall	Jaccard
Segment w/o SFA	90.2 ± 8.8	92.4 ± 6.2	91.3 ± 5.5	84.8 ± 6.8
$\{\mathbf{F}_{pos}^1\}$	91.0 ± 7.9	92.8 ± 5.9	92.2 ± 4.8	85.1 ± 8.4
$\{\mathbf{F}_{pos}^{1}, \mathbf{F}_{pos}^{2}\}$	91.6 ± 8.3	93.0 ± 4.9	92.5 ± 5.5	85.4 ± 5.8
$\{\mathbf{F}_{pos}^{1},\mathbf{F}_{pos}^{2},\mathbf{F}_{pos}^{3}\}$	91.6 ± 7.8	93.5 ± 5.6	92.9 ± 6.3	$\textbf{85.7} \pm \textbf{7.7}$

Furthermore, the running time of training/inference of our model is 0.31s for each video, only bringing extra 0.04s compared with the baseline model.

3) Segment-Frame Attention (SFA) Module: Table V shows the effectiveness of the hierarchical segment-frame attention (SFA) on M2CAI16 dataset. { \mathbf{F}_{pos}^1 , \mathbf{F}_{pos}^2 , \mathbf{F}_{pos}^3 } indicates SFA module between level 0 and 1, level 0 and 2, level 0 and 3, respectively. From comparison of "Segment w/o SFA "and "{ \mathbf{F}_{pos}^1 }", we can observe that using SFA between frames and one level of segment, the performance can achieve 0.8% improvement in terms of accuracy. Moreover, using SFA between frames and hierarchical segments can achieve better performance, demonstrating the effectiveness of the proposed hierarchical network and SFA.

To explore how segment-level information helps rectify erroneous predictions, we show the temporal attention weight

TABLE VI

COMPARISON WITH THE PERFORMANCE OF SFE WITH DIFFERENT TEMPORAL FUSION LAYERS ON M2CAI16. "CONV", "MP" AND "AP" INDICATE CONVOLUTION, MAX-POOLING AND AVERAGE-POOLING RESPECTIVELY

Methods	Accuracy	Precision	Recall	Jaccard
Base	87.1 ± 9.0	87.7 ± 7.1	88.8 ± 6.3	76.6 ± 8.9
Conv	90.0 ± 7.4	90.3 ± 6.5	91.0 ± 4.3	84.4 ± 7.5
MP	90.1 ± 7.5	91.8 ± 5.5	91.1 ± 4.4	84.7 ± 7.8
AP	90.5 ± 7.5	92.4 ± 5.6	91.6 ± 4.5	85.0 ± 7.7

of SFA in Fig. 6(a). Clearly, each high-level segment feature focuses on its neighboring frames, *i.e.*, the segment representations show the high similarity with the frame ones that sharing the same phase class. As a result, the features of frames (including ambiguous frames) belong to the same phase will be pulled together and refined to be consistent with the segment-level features. Fig. 6(b) visualizes the cosine similarity of pair-wise frames sampled from the test dataset. Notably, by using SFA, the features of frames from the same phases would be similar, and vice versa. In other words, with SFA, our method can better recognize surgical phases. For example, as shown in the right two heatmaps in Fig. 6(b), the frames in the same phase, *e.g.*, the purple arrows, are become more similar when using SFA.

4) Ablation on Parameters:

a) Comparison with different temporal fusion layers: As shown in Table VI, we ablate the performance of different temporal fusion layers, *i.e.*, convolution (Conv), maxpooling (MP) and average-pooling (AP). Note that in this ablation study, we just aim to compare the effect of different temporal fusion layers, and do not use the segment-frame attention (SFA) module. We find that AP is better than MP and Conv, *e.g.*, achieving 90.5% accuracy, 0.4% over than MP.



Fig. 7. Analysis of (a) β and (b) λ on M2CAI16 and Cholec80. We show the results of Accuracy, Precision, Recall and Jaccard of models with different β and λ .

TABLE VII ABLATION STUDY ON THE PERFORMANCE OF SFA WITH DIFFERENT MODELS ON M2CAI16

Methods	Accuracy	Precision	Recall	Jaccard
Base	87.1 ± 9.0	87.7 ± 7.1	88.8 ± 6.3	76.6 ± 8.9
SFA w/ Base	87.5 ± 7.0	88.2 ± 5.4	89.0 ± 6.0	77.0 ± 7.0
SFA w/ Conv	91.2 ± 7.3	92.5 ± 5.5	92.7 ± 4.5	85.3 ± 7.2
SFA w/ MP	91.3 ± 7.5	93.2 ± 5.5	92.9 ± 4.5	85.6 ± 7.7
SFA w/ AP	91.6 ± 7.8	93.5 ± 5.6	92.9 ± 6.3	$\textbf{85.7}~\pm~\textbf{7.7}$

This is because AP would encourage the overall frames in the kernel windows to be consistent with the high-level segment one. Hence, the features of the ambiguous frame would be pulled together with its corresponding high-level segments.

b) Comparison of SFA with different models: In Table VII, we compare the performance of SFA with Base and different temporal fusion layers, i.e., Conv MP and AP. "Conv", "MP" and "AP" indicate the convolution, max-pooling and averagepooling respectively. SFA is not very impressive without the segment-level information (Base), *i.e.*, only 0.6% improvement on Base. However, SFA can bring considerable improvement on the model with segment-level information. For example, the temporal fusion layer with Conv achieves 91.2% accuracy, outperforms 3.5% over the baseline model. The result demonstrates that SFA can capture the relations between high-level segments and low-level frames, which can help refine erroneous predictions in low-level frames. Furthermore, we also find that use AP as the temporal fusion layer bring the largest improvement for SFA, compared with Conv and MP. For instance, we obtain 91.6% accuracy on the temporal fusion layer with AP, 0.4% and 0.3% improvement over that with Conv and MP respectively.

c) Comparison with different kernel sizes k: In Table VIII, we analyze different kernel sizes k (see Eq. 3) in the temporal fusion layer. $1/\{1, 3, 5, 7, 9, 11\}$ indicates that we set k to be $\{1, 3, 5, 7, 9, 11\}$ respectively. Note that k = 1 indicates that the baseline model without the segment-level information. It is clear that with k > 1, *i.e.*, the segment-level information, the model achieve significant improvement. We can observe that too small and too large kernel size would hurt the performance. This is due to that too small kernel size may extract very local segment-level information, which can not capture high-level information to refine the ambiguous frames, *i.e.*, setting k to 3 only achieves 91.2% accuracy. On the other hand, too large kernel size may include the frames from different phases into one high-level segment, leading to mistakes in generating segments, *i.e.*, setting k to 11 only achieves 91.1% accuracy.

 TABLE VIII

 COMPARISON WITH DIFFERENT KERNEL SIZES k OF THE TEMPORAL

 FUSION LAYER ON M2CAI16

Size	Accuracy	Precision	Recall	Jaccard
1/1	87.1 ± 9.0	87.7 ± 7.1	88.8 ± 6.3	76.6 ± 8.9
1/3	91.2 ± 7.4	93.0 ± 5.9	92.5 ± 4.0	84.5 ± 7.8
1/5	91.5 ± 7.7	93.2 ± 6.5	92.5 ± 4.4	85.4 ± 7.4
1/7	91.6 ± 7.8	93.5 ± 5.6	92.9 ± 6.3	$85.7~\pm~7.7$
1/9	91.2 ± 7.7	93.2 ± 5.6	92.6 ± 6.5	85.3 ± 7.7
1/11	91.1 ± 7.7	93.3 ± 5.5	92.3 ± 6.4	85.2 ± 7.6

In our experiments, we find that 1/7 is the best reduction rate, which achieves 91.6% accuracy, 0.4% and 0.1% over that of 1/9 and 1/5 respectively.

d) Analysis on hyper-parameters β and λ : Fig. 7 shows the model performance with different values of β in Eq. 7 and λ in Eq. 10. β controls the weight of the frame-wise prediction and the segment-wise prediction. Setting β to zero indicates the model without the consistency between predictions of frames and segments, which shows the limited performance, *i.e.*, only obtaining the accuracy score of 87.7% on M2CAI16. In the experiments, we find that setting β to 1.0, our method achieves the best performance, *i.e.*, achieving 91.6% accuracy on M2CAI16, 3.9% over that setting β to 0.0. λ controls the importance of smoothing, which regularizes the model to predict smoothing results. It is clear that combining with the smoothing regularization, the model can achieve the better performance. For example, the accuracy of the model setting λ to 0.50 outperforms 0.8% over that of the model without \mathcal{L}_{smooth} , *i.e.*, $\lambda = 0.00$. We find that setting λ to 1.0 achieve the best performance on both M2CAI16 and Cholec80 datasets, *i.e.*,91.6% and 91.8% respectively.

D. Qualitative Analysis and Discussion

Fig. 8 shows the qualitative results of our method. We scale the temporal axes of three videos, *i.e.*, (a)-(c), for better visualization. As shown in Fig. 8 (a)-(b), our method predicts higher frame-wise accuracy, especially in the boundaries between two different phases. Furthermore, we can notice that "Base" makes many mistakes, *i.e.*, classifying ambiguous frames within different phases. For example, some frames within P3 are misclassified into P4 and some ones within P4 are misclassified into P3, as shown in Fig. 8(b). On the contrary, our proposed method can predict more robust and smooth results. Note that our methods predict higher Jaccard scores as shown in Fig. 8(c). Although the predictions of



Fig. 8. Prediction results of three video examples from M2CAI16 dataset. "P1-P5" indicates the phase 1 to phase 5. "GT" refers to the ground-truth and "Base" refers to baseline. The top row and the bottom row in Base and ours are the prediction results and the prediction scores, respectively. We scale the temporal axes of all videos for better visualization.

the baseline model and ours show the similar frame-wise accuracy, our methods present much higher phase-wise Jaccard scores, *i.e.*, more smooth results. From the comparison of Fig. 8(b) and (c), we can find that our method can solve the ambiguous prediction inside phases very well.

However, both the baseline model and our method generate the misalignment boundaries when transits on phase to another one, as shown in Fig. 8(c). We believe that this may be due to the noise annotations, since it is difficult to determine which frame is the precise boundary, and this process is subjective. In the future, we may use some uncertainty analysis methods [54]–[57] to alleviate this problem. Furthermore, in real life, surgeons may sometime have to redo a phase which leads to more complicated surgical activities than the current public datasets. In the future, we will collect the surgical videos by ourselves to develop algorithms for this situation. Moreover, we will develop a more flexible frame-level module and design an attention module to capture the relationships among different or recurrent phases.

V. CONCLUSION

This paper presents a novel segment-attentive hierarchical consistency network (SAHC) for surgical phase recognition from videos. Unlike previous methods, our key idea is to explore the segment-level semantics and use it to refine the erroneous predictions caused by ambiguous lowlevel frames. SAHC consists of two innovative modules: a segment-level hierarchical consistency network to generate high-level semantic-consistent segments and a segmentframe attention (SFA) module to better reflect high-level segment information to low-level frames. Our method achieves improved estimates of performance on two public surgical video recognition datasets. Ablation study demonstrates the effectiveness of the proposed segment-level hierarchical consistency network and SFA module.

REFERENCES

- [1] L. Maier-Hein *et al.*, "Surgical data science for next-generation interventions," *Nature Biomed. Eng.*, vol. 1, no. 9, pp. 691–696, 2017.
- [2] A. K. Tanwani, P. Sermanet, A. Yan, R. Anand, M. Phielipp, and K. Goldberg, "Motion2 Vec: Semi-supervised representation learning from surgical videos," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2020, pp. 2174–2181.

- [3] 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 Trans. Med. Imag.*, vol. 36, no. 1, pp. 86–97, Jan. 2017.
- [4] N. Padoy, "Machine and deep learning for workflow recognition during surgery," *Minimally Invasive Therapy Allied Technol.*, vol. 28, no. 2, pp. 82–90, 2019.
- [5] Y. Jin, Y. Long, C. Chen, Z. Zhao, Q. Dou, and P.-A. Heng, "Temporal memory relation network for workflow recognition from surgical video," *IEEE Trans. Med. Imag.*, vol. 40, no. 7, pp. 1911–1923, Jul. 2021.
- [6] X. Gao, Y. Jin, Y. Long, Q. Dou, and P.-A. Heng, "Trans-SVNet: Accurate phase recognition from surgical videos via hybrid embedding aggregation transformer," 2021, arXiv:2103.09712.
- [7] T. Czempiel et al., "Tecno: Surgical phase recognition with multi-stage temporal convolutional networks," in Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent. Lima, Peru: Springer, 2020, pp. 343–352.
- [8] F. Yi and T. Jiang, "Not end-to-end: Explore multi-stage architecture for online surgical phase recognition," 2021, arXiv:2107.04810.
- [9] Y. A. Farha and J. Gall, "MS-TCN: Multi-stage temporal convolutional network for action segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2019, pp. 3575–3584.
- [10] C. Lea, M. D. Flynn, R. Vidal, A. Reiter, and G. D. Hager, "Temporal convolutional networks for action segmentation and detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 156–165.
- [11] T. Blum, H. Feußner, and N. Navab, "Modeling and segmentation of surgical workflow from laparoscopic video," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Beijing, China: Springer, 2010, pp. 400–407.
- [12] L. Zappella, B. B. Haro, G. Hager, and R. Vidal, "Surgical gesture classification from video and kinematic data," *Med. Image Anal.*, vol. 17, no. 7, pp. 732–745, 2013.
- [13] F. Lalys, L. Riffaud, D. Bouget, and P. Jannin, "A framework for the recognition of high-level surgical tasks from video images for cataract surgeries," *IEEE Trans. Biomed. Eng.*, vol. 59, no. 4, pp. 966–976, Dec. 2011.
- [14] O. Dergachyova, D. Bouget, A. Huaulmé, X. Morandi, and P. Jannin, "Automatic data-driven real-time segmentation and recognition of surgical workflow," *Int. J. Comput. Assist. Radiol. Surg.*, vol. 11, pp. 1081–1089, Mar. 2016.
- [15] L. Tao, L. Zappella, G. D. Hager, and R. Vidal, "Surgical gesture segmentation and recognition," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Nagoya, Japan: Springer, 2013, pp. 339–346.
- [16] G. Quellec, M. Lamard, B. Cochener, and G. Cazuguel, "Real-time segmentation and recognition of surgical tasks in cataract surgery videos," *IEEE Trans. Med. Imag.*, vol. 33, no. 12, pp. 2352–2360, Dec. 2014.
- [17] C. Lea, G. D. Hager, and R. Vidal, "An improved model for segmentation and recognition of fine-grained activities with application to surgical training tasks," in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, Jan. 2015, pp. 1123–1129.
- [18] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. pattern Recognit.*, Jun. 2016, pp. 770–778.

- [19] Y. Jin *et al.*, "SV-RCNet: Workflow recognition from surgical videos using recurrent convolutional network," *IEEE Trans. Med. Imag.*, vol. 37, no. 5, pp. 1114–1126, May 2017.
- [20] S.-J. Li, Y. AbuFarha, Y. Liu, M.-M. Cheng, and J. Gall, "MS-TCN++: Multi-stage temporal convolutional network for action segmentation," *IEEE Trans. Pattern Anal. Mach. Intell.*, early access, Sep. 4, 2020, doi: 10.1109/TPAMI.2020.3021756.
- [21] O. Zisimopoulos et al., "DeepPhase: Surgical phase recognition in cataracts videos," in Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent. Granada, Spain: Springer, 2018, pp. 265–272.
- [22] Y. Jin *et al.*, "Multi-task recurrent convolutional network with correlation loss for surgical video analysis," *Med. Image Anal.*, vol. 59, Jan. 2020, Art. no. 101572.
- [23] H. Nakawala, R. Bianchi, L. E. Pescatori, O. De Cobelli, G. Ferrigno, and E. De Momi, "Deep-onto' network for surgical workflow and context recognition," *Int. J. Comput. Assist. Radiol. Surg.*, vol. 14, no. 4, pp. 685–696, 2019.
- [24] D. Sarikaya, K. A. Guru, and J. J. Corso, "Joint surgical gesture and task classification with multi-task and multimodal learning," 2018, arXiv:1805.00721.
- [25] Z. Shou, D. Wang, and S.-F. Chang, "Temporal action localization in untrimmed videos via multi-stage CNNs," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 1049–1058.
- [26] X. Ding, N. Wang, X. Gao, J. Li, X. Wang, and T. Liu, "Weakly supervised temporal action localization with segment-level labels," 2020, arXiv:2007.01598.
- [27] G. Li, J. Li, N. Wang, X. Ding, Z. Li, and X. Gao, "Multihierarchical category supervision for weakly-supervised temporal action localization," *IEEE Trans. Image Process.*, vol. 30, pp. 9332–9344, 2021.
- [28] J. Gao, C. Sun, Z. Yang, and R. Nevatia, "TALL: Temporal activity localization via language query," in *Proc. IEEE Int. Conf. Comput. Vis.*, Oct. 2017, pp. 5267–5275.
- [29] L. Anne Hendricks, O. Wang, E. Shechtman, J. Sivic, T. Darrell, and B. Russell, "Localizing moments in video with natural language," in *Proc. IEEE Int. Conf. Comput. Vis.*, Oct. 2017, pp. 5803–5812.
- [30] X. Ding et al., "Support-set based cross-supervision for video grounding," in Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV), Oct. 2021, pp. 11573–11582.
- [31] R. Zeng, H. Xu, W. Huang, P. Chen, M. Tan, and C. Gan, "Dense regression network for video grounding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2020, pp. 10287–10296.
- [32] B. Zhou, A. Andonian, A. Oliva, and A. Torralba, "Temporal relational reasoning in videos," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 803–818.
- [33] C. Feichtenhofer, H. Fan, J. Malik, and K. He, "Slowfast networks for video recognition," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, Oct. 2019, pp. 6202–6211.
- [34] C. Yang, Y. Xu, J. Shi, B. Dai, and B. Zhou, "Temporal pyramid network for action recognition," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2020, pp. 591–600.
- [35] T.-Y. Lin, P. Dollar, R. Girshick, K. He, B. Hariharan, and S. Belongie, "Feature pyramid networks for object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2117–2125.
- [36] D. Zhukov, J.-B. Alayrac, I. Laptev, and J. Sivic, "Learning actionness via long-range temporal order verification," in *Proc. Eur. Conf. Comput. Vis.* Glasgow, U.K.: Springer, 2020, pp. 470–487.

- [37] C.-Y. Wu, C. Feichtenhofer, H. Fan, K. He, P. Krahenbuhl, and R. Girshick, "Long-term feature banks for detailed video understanding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 284–293.
- [38] C.-Y. Wu and P. Krahenbuhl, "Towards long-form video understanding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 1884–1894.
- [39] A. Vaswani et al., "Attention is all you need," in Proc. Adv. Neural Inf. Process. Syst., 2017, pp. 5998–6008.
- [40] F. Yi and T. Jiang, "Hard frame detection and online mapping for surgical phase recognition," in *Proc. Int. Conf. Med. Image Comput. Comput.-Assist. Intervent.* Shenzhen, China: Springer, 2019, pp. 449–457.
- [41] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [42] D. Shao, Y. Zhao, B. Dai, and D. Lin, "Intra-and inter-action understanding via temporal action parsing," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 730–739.
- [43] A. Dosovitskiy et al., "An image is worth 16×16 words: Transformers for image recognition at scale," 2020, arXiv:2010.11929.
- [44] A. Arnab, M. Dehghani, G. Heigold, C. Sun, M. Lučič, and C. Schmid, "ViViT: A video vision transformer," 2021, arXiv:2103.15691.
- [45] J. Liang, J. Cao, G. Sun, K. Zhang, L. Van Gool, and R. Timofte, "SwinIR: Image restoration using Swin transformer," in *Proc. IEEE/CVF Int. Conf. Comput. Vis.*, Oct. 2021, pp. 1833–1844.
- [46] Z. Liu *et al.*, "Swin transformer: Hierarchical vision transformer using shifted Windows," 2021, arXiv:2103.14030.
- [47] H. Chen et al., "Pre-trained image processing transformer," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit., Jun. 2021, pp. 12299–12310.
- [48] Z. Sun, S. Cao, Y. Yang, and K. Kitani, "Rethinking transformerbased set prediction for object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 3611–3620.
- [49] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," 2016, arXiv:1607.06450.
- [50] A. P. Twinanda, D. Mutter, J. Marescaux, M. de Mathelin, and N. Padoy, "Single- and multi-task architectures for surgical workflow challenge at M2CAI 2016," 2016, arXiv:1610.08844.
- [51] R. Stauder, D. Ostler, M. Kranzfelder, S. Koller, H. Feußner, and N. Navab, "The TUM LapChole dataset for the M2CAI 2016 workflow challenge," 2016, arXiv:1610.09278.
- [52] A. Paszke, S. Gross, and S. Chintala, (2017). Pytorch Deep Learning Framework. [Online]. Available: http://pytorch.org/
- [53] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," 2014, arXiv:1412.6980.
- [54] J. Peterson, R. Battleday, T. Griffiths, and O. Russakovsky, "Human uncertainty makes classification more robust," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9617–9626.
- [55] W. J. Maddox, P. Izmailov, T. Garipov, D. P. Vetrov, and A. G. Wilson, "A simple baseline for Bayesian uncertainty in deep learning," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 32, 2019, pp. 13153–13164.
- [56] Y. Li, Y. Duan, Z. Kuang, Y. Chen, W. Zhang, and X. Li, "Uncertainty estimation via response scaling for pseudo-mask noise mitigation in weakly-supervised semantic segmentation," in *Proc. AAAI*, 2022, pp. 4556–4562.
- [57] Z. Wang, X. Ding, W. Zhao, and X. Li, "Less is more: Surgical phase recognition from timestamp supervision," 2022, arXiv:2202.08199.