

Temporal Memory Enhancement for Semantic Segmentation in Surgical Video

Zheyao Gao¹

Qian Wu¹

Yueyao Chen¹

Cheng Chen²

Hon Chi Yip³

Winnie Chiu Wing Chu⁴

Qi Dou¹

ZHEYAOGAO@CUHK.EDU.HK

QIANWU@LINK.CUHK.EDU.HK

CHENYUEYAO@LINK.CUHK.EDU.HK

CCHEN@EEE.HKU.HK

HCYIP@SURGERY.CUHK.EDU.HK

WINNIECHU@CUHK.EDU.HK

QIDOU@CUHK.EDU.HK

¹ CSE, The Chinese University of Hong Kong

² EEE, Hong Kong University

³ Department of Surgery, The Chinese University of Hong Kong

⁴ DIIR, The Chinese University of Hong Kong

Editors: Under Review for MIDL 2026

Abstract

Segmenting critical anatomical structures in surgical videos can enhance precision and patient safety by alerting surgeons to potential complications. While current methods that store features from past frames have advanced the performance in video segmentation, their reliance on a fixed-range local memory often fails to capture complex temporal contexts of surgical scenes. Specifically, the memory could fill with redundant features or omit informative frames due to non-uniform rate of operations by the surgeons. Besides, the image features in the same phase of the surgery share similar patterns, while local memory could not capture such long-term relationship. Therefore, we propose a memory enhancement method to enrich the local temporal context and incorporate global phase context for surgical video semantic segmentation. Concretely, we improve the local memory with a feature selection module based on Determinantal Point Process to choose past features that are diverse and relevant to the current feature. Besides, we introduce a global memory to store the common patterns of frames within each phase based on the conditional variational autoencoder with mixture of Gaussian prior. Experiments on endoscopic submucosal dissection (ESD) and laparoscopic cholecystectomy (CHO) video segmentation demonstrate that our method achieves superior performance over existing methods. Our code is available via https://github.com/key1589745/surgery_segmentation

Keywords: Surgery, video segmentation, temporal memory.

1. Introduction

The analysis of surgical video, such as from endoscopic submucosal dissection (ESD) or laparoscopic cholecystectomy (CHO), is essential for developing intra-operative guidance systems. In ESD, segmenting structures like the submucosa is key for safe lesion removal (Yang et al., 2023), while in LC, identifying critical anatomy like the cystic duct is paramount for preventing complications (Murali et al., 2023). However, a fundamental challenge across these domains is the high variability of surgical video—characterized by changes in camera

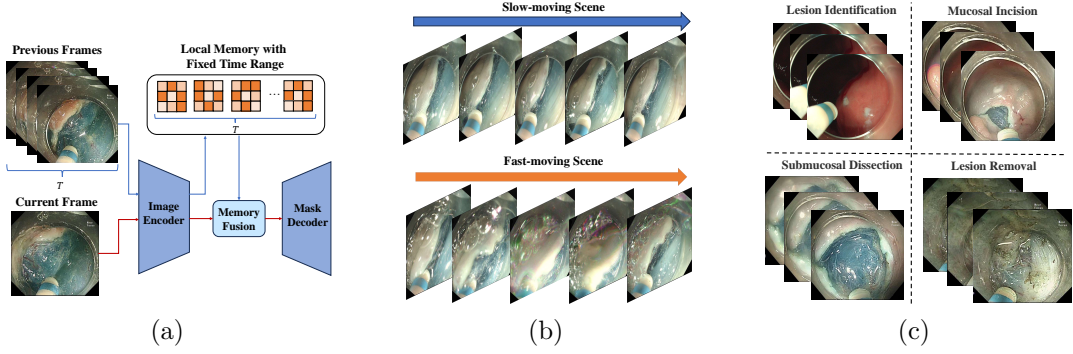


Figure 1: The typical memory-based video segmentation framework and its limitations in surgical video segmentation. (a) Video segmentation methods such as SAM2 (Ravi et al., 2024) apply a local memory bank to store the features of past few frames with a fixed time interval. (b) The temporal dynamics varies in surgical videos (c) Frames within each phase of the surgery share similar patterns.

motion, lighting, and tissue appearance—which often causes segmentation and detection models to produce inconsistent, frame-by-frame predictions (Mascagni et al., 2025).

Different from image segmentation, a crucial problem in video segmentation lies in how to effectively leverage temporal information to delineate regions of interest in current frames. To fully exploit the temporal dependency between frames, previous works (Wang et al., 2021; Zhuang et al., 2024, 2022) have proposed various feature extraction and memory mechanisms to capture the temporal relations and integrate the information from the past with current frames over the video sequence. Fig. 1(a) shows a typical framework of video segmentation with memory module, which incorporates temporal context by storing features of past frames and their segmentation masks as references. Recently, SAM2 (Ravi et al., 2024) and its variants (Liu et al., 2024) have leveraged the progress of such memory strategies to develop foundation models for video segmentation, which has achieved excellent performance in both natural and surgical scenes.

While video segmentation methods commonly employ a local memory bank to store features of recent frames, this can be insufficient for capturing long-term and complex temporal dependencies. To address this, advanced memory-based frameworks (Cheng and Schwing, 2022; Deng et al., 2024) incorporate both a local working memory and a long-term memory component, thereby extending the model’s effective temporal receptive field to hundreds or even thousands of frames. However, these existing methods remain suboptimal for endoscopic video segmentation due to two distinctive properties of surgical workflows. First, as illustrated in Fig. 1(b), the surgeon’s motion pace is highly variable. A static memory update strategy may store excessive redundant features during slow, deliberate phases, while missing critical reference frames during periods of rapid movement. Second, surgical procedures follow a structured workflow. As shown in Fig. 1(c), frames within the same surgical phase share similar visual patterns—such as tissue appearance and lighting conditions—that constitute a long-term, phase-specific context. Current memory mechanisms,

which primarily rely on feature similarity or recency, fail to explicitly model and leverage this high-level semantic consistency across the entire procedure.

To address the complex temporal context in surgical video segmentation, we propose a memory enhancement strategy comprising two key mechanisms. First, to reduce memory redundancy while preserving valuable past features, we design a feature selection module based on a greedy Determinantal Point Process (DPP) (Anari et al., 2025). This module efficiently selects a diverse and relevant subset from local memory to enhance the current feature representation. Second, to capture global phase context, we propose a conditional variational autoencoder with a mixture of Gaussian prior (CVAE-MoG) that encodes common patterns of each surgical phase. This module clusters frames by phase and integrates phase-specific context from the corresponding cluster to enhance current features.

Our contributions include: (1) we propose a temporal memory enhancement method for semantic segmentation in surgical video, which improves the segmentation performance by incorporating richer local context and long-term phase context; (2) we propose a local memory feature selection module based on a greedy DPP to address the varying movement rates in surgical videos. It selects diverse and relevant features from memory while filtering out redundancies; (3) we design a global memory module based on CVAE-MoG, which exploits common patterns of frames in each phase and provide the long-term phase context to enhance the segmentation performance. (4) we validated the effectiveness of the proposed method on two surgical video segmentation datasets. The results show that our method achieves superior performance over existing video segmentation methods.

2. Related Works

2.1. Video Semantic Segmentation

The core challenge in Video Semantic Segmentation (VSS) is maintaining temporal coherence to ensure consistent pixel classification across frames. Traditional approaches, such as optical flow-based propagation (Nilsson and Sminchisescu, 2017), 3D convolutional networks (He et al., 2017), and recurrent or attention-based mechanisms (Li et al., 2025), have faced challenges in generalizing to variable scenes while balancing robust temporal consistency with computational efficiency. Recently, the Segment Anything Model 2 (SAM2) (Ravi et al., 2024) has emerged as a promising foundation model. Pre-trained on large-scale datasets, SAM2 enables promptable video object segmentation and has demonstrated strong generalization across diverse scenes. This capability has prompted recent research (Syed Ariff et al., 2025) to adapt SAM2 for the VSS task, offering a new pathway toward achieving coherent, efficient, and generalizable video segmentation.

2.2. Surgical Scene Segmentation

Recent research in surgical video segmentation has rapidly evolved, moving from developing specialized models (Jin et al., 2022; Wang et al., 2024; Li et al., 2023) to adapting powerful foundation models like SAM2 (Ravi et al., 2024) for the surgical domain. Adaptations such as SurgiSAM2 demonstrates that fine-tuning the foundational model on multi-dataset surgical annotations significantly boosts anatomy segmentation performance (Devanish et al., 2025); Surgical SAM 2 introduces an efficient frame pruning (EFP) mechanism, reducing

computational cost (Liu et al., 2024); SAM2S establishes a new benchmark by integrating semantic learning and a diverse memory mechanism for long-term, class-aware instrument and tissue tracking (Liu et al., 2025a). While performance on larger structures such as surgical tools have become more robust, significant challenges remain in achieving reliable segmentation of small, critical structures like nerves and blood vessels.

2.3. Temporal Memory enhancement

Memory-based methods improve video segmentation by storing and retrieving past frame features, ensuring temporal consistency on various time scales. For long-term consistency, models like STMN (Oh et al., 2019) use an external memory bank, while XMem (Cheng and Schwing, 2022) adopts a multi-store architecture to retain information across thousands of frames. Specialized tools such as Memsam (Deng et al., 2024) further adapt this for domains like echocardiography analysis. For local refinement, methods enhance robustness and quality of the memory: SAMURAI (Yang et al., 2024) adds context-aware and occlusion-resilient memory to SAM2 and MA-SAM2 (Yin et al., 2025) adapts it for surgical videos; QDMN (Liu et al., 2025b) employs a dynamic memory bank that stores only high-quality masks to reduce error propagation.

3. Method

The aim of the proposed method is to improve the latent feature f learned by the image encoder of a segmentation network by incorporating both local and global phase context. Fig. 2 shows our framework.

During training, the image encoder extracts memory features $\{f_{-t}\}_{t=1}^T$ from past frames $\{x_{-t}\}_{t=1}^T$ and the current feature f from the frame x . The local memory evaluates similarities among memory features and with the current feature, selecting a subset \mathcal{B} that is both diverse and relevant as local context. The global memory concatenates the enhanced feature f' with phase context \tilde{z}^l , which is sampled from $q(z, c|x, y)$ estimated by post network $\beta(\cdot)$. The training objective combines segmentation loss ℓ_{CE} with KL divergence loss ℓ_{KL} , which aligns $p(z, c|x, y)$ estimated by the prior network $\alpha(\cdot)$ with $q(z, c|x, y)$.

During inference, the local memory only evaluates similarities between the current feature and memory features and stores the similarities for subset selection. The global memory generates \tilde{z}^l from $p(z, c|x)$ since the ground truth mask is not available.

We elaborate the details of the local memory and global memory in the following sections.

3.1. Local Memory with DPP selection

Given the current frame feature: $f \in \mathbb{R}^C$ and memory cache: $\mathcal{A} = \{f_{-t}\}_{t=1}^T$ where $f_{-t} \in \mathbb{R}^C$, we apply the DPP algorithm to choose M features $\mathcal{B} = \{f_{-t}\}_{t=1}^M$ where $M < T$. The selection should maximize relevance - features similar to current frame f and diversity - selected features are different from each other.

To prepare for the selection, the similarity matrix $S \in \mathbb{R}^{T \times T}$ and the relevance matrix $R \in \mathbb{R}^{T \times T}$ are calculated based on cosine similarity as follows,

$$S_{\mathcal{A}} = \bar{f}_{\mathcal{A}} \bar{f}_{\mathcal{A}}^{\top}, \quad R_{\mathcal{A}} = \text{diag} \left(\exp(\gamma \cdot \bar{f}_{\mathcal{A}} \bar{f}^{\top}) \right), \quad (1)$$




Figure 2: The framework of temporal memory enhancement method for video semantic segmentation. The local memory collects features of past T frames and applies greedy Determinantal Point Process (DPP) to select diverse and relevant features to enhance the current feature. The global memory models the common image and segmentation pattern within each phase as a Mixture of Gaussian (MoG) distribution. For each frame, it selects the most probable component and samples a phase context \tilde{z}_c^l , which is then concatenated with the feature f' .

where $\bar{f}_{\mathcal{A}} = [\frac{f_{-1}}{\|f_{-1}\|_2}, \frac{f_{-2}}{\|f_{-2}\|_2}, \dots, \frac{f_{-T}}{\|f_{-T}\|_2}] \in \mathbb{R}^{T \times C}$ and $\bar{f} = \frac{f}{\|f\|_2} \mathbb{R}^C$ are normalized memory and current features, respectively; $\gamma \in [0, 1]$ is a hyperparameter, balancing the weight of relevance score in the selection; $\text{diag}(\cdot)$ denotes the diagonal matrix formed by a vector. For streaming data, the cosine similarity $\bar{f}_{\mathcal{A}} \bar{f}^\top$ is calculated and cached in the memory during inference. The similarity matrix S could be derived by reusing the cached similarities, avoiding additional computation.

To find the optimal subset \mathcal{B}^* , we define the probability distribution $p(\mathcal{B})$ over all possible subsets of \mathcal{A} as follows,

$$p(\mathcal{B}) \propto \det(R_{\mathcal{B}} S_{\mathcal{B}} R_{\mathcal{B}}) = \left(\prod_{t \in \mathcal{B}} r_t \right)^2 \cdot \det(S_{\mathcal{B}}), \quad (2)$$

where $r_t = \exp(\gamma \cdot \bar{f}_t \bar{f}^\top)$ denotes a diagonal element in $R_{\mathcal{B}}$ and $\det(\cdot)$ denotes the determinant of a matrix. It could be observed based on Eq. (1) and (2) that the probability of selecting a subset correlates positively with similarities to the current frame (*i.e.*, diagonal elements of $R_{\mathcal{B}} S_{\mathcal{B}} R_{\mathcal{B}}$) and correlates negatively with the similarities between memory features (*i.e.*, off-diagonal elements). Therefore, the subset that contains diverse and relevant features could be derived through Maximum A Posteriori (MAP) inference,

$$\mathcal{B}^* = \arg \max_{\mathcal{B} \subseteq \mathcal{A}, |\mathcal{B}|=M} p(\mathcal{B}). \quad (3)$$

Finding the exact MAP (the subset with the absolute highest probability) by listing all combinations is compute-intensive. Therefore, we apply a greedy update algorithm (Gautier et al., 2019), which builds the subset one item at a time and always picks the memory feature maximizing the determinant increase. It reduces the time complexity to $o(TM^2)$, which is efficient when the size of working memory is small. The details of this algorithm are provided in Appendix A.

3.2. Global Memory based on CVAE-MoG

To incorporate the long-term memory information in the current feature f' from the preceding stage, we propose to represent the common pattern of the image x and segmentation y of each phase with the latent phase context z that comes from a MoG distribution $\sum_{c=1}^C \pi_c \mathcal{N}(\mu_c, \sigma_c^2 I)$, where $\mu_c, \sigma_c^2 \in \mathbb{R}^d$ are the parameters of the c th Gaussian component in the mixture distribution conditioned on the image and segmentation. $\pi_c \in [0, 1]$ is the mixing probability.

The MoG distribution is estimated based on the idea of conditional variational autoencoder (Kohl et al., 2018). Concretely, the objective is to find a variational distribution $q(z, c|x, y)$ to approach the true posterior by maximizing the lower bound of the likelihood $p(y|x)$ written as follows,

$$\begin{aligned} \log p(y|x) &= \log \int \sum_{c=1}^C p(y, z, c|x) dz \\ &\geq \mathbb{E}_{q(z, c|x, y)} [\log p(y|z, x)] - \mathcal{D}_{KL}[q(z, c|x, y) || p(z, c|x)], \end{aligned} \quad (4)$$

where \mathcal{D}_{KL} is the KL divergence. Since the lower bound could not be calculated explicitly, we introduce the prior network $\alpha(\cdot)$ and posterior network $\beta(\cdot)$ to estimate the parameters of $p(z, c|x)$ and $q(z, c|x, y)$, respectively, and maximize the lower bound through gradient descend. Formally, the distributions above are derived based on the following definitions,

$$p(z|c, x) = \mathcal{N}(\mu_c(x; \alpha), \sigma_c^2(x; \alpha)I), \quad p(c|x) = \text{Cat}(\pi^p(x; \alpha)), \quad (5a)$$

$$q(z|x, y) = \mathcal{N}(\mu(x, y; \beta), \sigma^2(x, y; \beta)I), \quad q(c|x, y) = \text{Cat}(\pi^q), \quad (5b)$$

where the mean field assumption is used in variational approximation, *i.e.*, $q(z, c|x, y) = q(z|x, y)q(c|x, y)$. Following the previous work (Jiang et al., 2017), the mixing probability π^q of the variational distribution is calculated as follows,

$$\pi_c^q = \mathbb{E}_{q(z|x, y)} [q(c|z)] \approx \frac{q(z = \tilde{z}^l | c, x, y) p(c)}{\sum_{c'=1}^C q(z = \tilde{z}^l | c', x, y) p(c')}, \quad (6)$$

where $\tilde{z}^l \sim q(z|x, y)$ and $p(c)$ is set to $\frac{1}{C}$ for each cluster. Based on the derivation of the above distributions, the first term in Eq. (4) could be maximized by minimizing the cross-entropy loss between the prediction of the mask decoder $\phi(\cdot)$ and the ground truth \hat{y} . Formally, the loss function for this term could be written as follows,

$$\ell_{seg} = \text{CE}(\phi(f' \oplus \tilde{z}_c^l), \hat{y}), \quad (7)$$

where \tilde{z}_c^l is sampled from $q(z|x, y)$ using the reparameterization trick (Kingma et al., 2013) such that the posterior encoder could be updated by backpropagation; \oplus denotes contention in the channel dimension by broadcasting \tilde{z}_c^l along the spatial dimension of f' .

The KL term could derive explicitly by using the definitions in Eq. (5a) and (5b). Thus, this term could be minimized as follows,

$$\ell_{KL} = \sum_{c=1}^C \pi_c^q \left(\log \frac{\sigma_c(x; \alpha)}{\sigma_c(x, y; \beta)} + \frac{\sigma_c^2(x; \alpha) + (\mu_c(x; \alpha) - \mu_c(x, y; \beta))^2}{2\sigma_c^2(x, y; \beta)} - \log \pi_c^p(x; \alpha) \right). \quad (8)$$

ℓ_{KL} forces $p(z, c|x)$ to be close with $q(z, c|x, y)$ that encodes the pattern of the segmentation mask. Therefore, the samples from $p(z|c, x)$ could be used as the phase context during inference when ground truth is not available. The detailed derivation for the lower bound and KL loss is provided in Appendix B.

4. Experiments

We validate our method on two surgical video datasets: ESD and CHO. In Section 4.2, we compare it against state-of-the-art video segmentation methods, followed by an ablation study and hyperparameter analysis in Section 4.3.

4.1. Dataset and Implementation details

Dataset. ESD dataset consisted of 1,201 annotated frames from 34 endoscopic surgical videos. For each video, submucosal, muscular tissue, and blood vessels in critical frames were annotated by expert surgeons. The dataset is available from the corresponding author upon reasonable request and non-commercial purposes. CHO dataset comes from the public dataset Endoscapes-Seg50 (Murali et al., 2023), containing 50 laparoscopic cholecystectomy videos, of which 493 frames are annotated with instance and semantic segmentation masks for six classes (five anatomical structures and one tool). Annotations are provided sparsely at a rate of one frame every 30 seconds, with splits of 30 training, 10 validation, and 10 test videos. For pre-processing, all the images were normalized to the same mean and variance with the training data in the previous work (Ravi et al., 2024) and resized to 256×256 as the network inputs.

Implementation details. We used the Hiera-large encoder pretrained by SAM2 (Ravi et al., 2024) as the backbone of the image encoder. Similar to the previous work (Chen et al., 2023), we froze most of the parameters in the encoder and only trained the adapters. In local memory, the memory size T and the subset size M were set to 16 and 8 for ESD dataset, 8 and 4 for CHO dataset. The time intervals between past frames were set to 0.5 and 1 based on the experience of previous works (Cao et al., 2023; Murali et al., 2023) for ESD dataset and CHO dataset, respectively. For the memory fusion module, we used the two layers of cross-attention as used in SAM2. In terms of the prior and posterior networks, we applied a similar configuration as in the previous work (Kohl et al., 2018) and set the number of components of the MoG as 4 for ESD dataset and 3 for CHO dataset in the comparison study. For the mask decoder, we applied a lightweight decoder network designed in the SegFormer (Xie et al., 2021) model. The framework was trained using Adam optimizer with an initial learning rate of $6e-4$ for 300 epochs, which decayed by a factor of

Table 1: Comparison results with state-of-the-art video segmentation methods.

Methods	CHO dataset			ESD dataset			FPS
	mIoU (%)	F-score (%)	wDice (%)	mIoU (%)	F-score (%)	wDice (%)	
Mask2Former	83.7 \pm 1.89	39.6 \pm 4.07	86.9 \pm 1.43	81.9 \pm 2.45	72.2 \pm 2.67	85.8 \pm 1.76	82.7
FRGM	82.9 \pm 1.89	39.2 \pm 4.07	86.2 \pm 1.43	81.1 \pm 2.45	71.5 \pm 2.67	85.0 \pm 1.76	111.9
QDMN++	84.5 \pm 1.55	42.9 \pm 3.71	89.0 \pm 0.94	83.4 \pm 2.07	72.8 \pm 2.59	87.1 \pm 1.28	78.4
MemSAM	84.9 \pm 1.38	44.5 \pm 3.53	89.7 \pm 0.70	84.1 \pm 1.88	73.1 \pm 2.55	87.8 \pm 1.04	57.8
SAM2	84.7 \pm 1.32	46.1 \pm 3.48	90.8 \pm 3.82	85.1 \pm 1.81	73.6 \pm 2.84	88.2 \pm 0.78	48.5
SurgSAM2	85.1 \pm 1.21	45.6 \pm 3.35	91.0 \pm 0.45	84.8 \pm 1.69	73.3 \pm 2.51	88.4 \pm 0.79	76.1
Ours	86.6 \pm 0.83	47.8 \pm 3.21	92.7 \pm 0.38	86.9 \pm 1.13	76.4 \pm 2.12	90.5 \pm 0.59	66.3

0.1 on the plateau of validation loss. The whole framework was implemented with PyTorch on RTX 4090 GPU.

Experiment settings. For experiments on ESD dataset, we randomly split the videos into the proportion of 60%, 20%, and 20% for training validation and testing, respectively. For experiments on CHO dataset, we follow the split of training, validation and testing in the previous work (Murali et al., 2023). The evaluation metrics include the average of Intersection over Union (mIoU), the F-score, and the weighted average of Dice score (wDice) according to the occurrence of each structure (Sudre et al., 2017). The results in all experiments were obtained by performing three independent runs.

4.2. Comparison Study

To demonstrate the effectiveness of the proposed memory enhancement method, we compared with widely used methods for video segmentation in natural scene: Mask2Former (Cheng et al., 2022) and SAM2 (Ravi et al., 2024), memory enhancement methods: MemSAM (Deng et al., 2024) that introduces a noise-resilient spatial temporal methods based on SAM, QDMN++ (Liu et al., 2025b) that introduces a quality control module filter low quality memory features, surgical video segmentation method: FRGM (Yang et al., 2023) which designs a calibration module based on an image segmentation model and SurgSAM2 (Liu et al., 2024), an adaptation of SAM2 to surgical videos.

Table 1 presents the comparison results. It could be observed that our proposed method achieves state-of-the-art performance across all segmentation accuracy metrics (mIoU, F-score, wDice) on both the CHO and ESD datasets. Among the compared methods, SAM2 and memory-enhanced methods (MemSAM, QDMN++) generally outperform the generic video segmentation model Mask2Former and the surgical-specific FRGM in terms of accuracy, underscoring the value of dedicated memory mechanisms and powerful foundation models. While our method does not achieve the highest FPS, its 66.3 FPS is sufficient for real-time processing and represents a favorable balance, being substantially faster than SAM2 and MemSAM while significantly outperforming them in accuracy. This result confirms that our memory enhancement method effectively leverages temporal information to boost segmentation precision and stability without sacrificing practical applicability.




Figure 3: The segmentation results achieved by the proposed method and compared methods on ESD dataset and CHO dataset. ESD dataset includes submucosal (green), muscular (blue), and blood vessel (red) and CHO dataset includes gallbladder (blue), cystic duct (yellow), cystic artery (green), calot’s triangle (orange), and tool (white).

4.3. Ablation Study

To validate the design of our proposed local and global memory modules, we compared them with different variants. Concretely, we first studied the importance of the similarity scores and relevance scores in the DPP selection by comparing with the ablated methods (w/o Diversity, w/o Relevance). Besides, the method that applies random selection (Random) was also compared as baseline. For the global memory, we implemented the method without phase context (w/o Cluster), the method that does not use segmentations mask to generate phase context (w/o Mask), and the method that uses K-means clustering to categorize the features into different clusters (KM-Cluster).

Table 2 presents the comparison results. For local memory, our method—which uses DPP selection with both relevance and diversity scores, achieves the best performance across all primary metrics (mIoU, F-score, wDice) on both datasets. It significantly outperforms the random baseline (e.g., +3.1% mIoU on CHO) and the ablated variants (“w/o Relevance” and “w/o Diversity”), confirming that both components are essential. For global memory, the proposed CVAE-MoG clustering method surpasses simplified variants (“w/o Cluster”, “w/o Mask”, “KM-Cluster”), demonstrating the value of using segmentation masks to create discriminative phase contexts. While the full method has a slightly lower FPS (66.3) than Random (78.1), this trade-off is justified by its substantially superior segmentation accuracy and robustness.

Table 2: Ablation study of the local and global memory.

Methods	CHO dataset			ESD dataset			FPS
	mIoU (%)	F-score (%)	wDice (%)	mIoU (%)	F-score (%)	wDice (%)	
Random	83.5 \pm 1.20	43.1 \pm 3.60	88.3 \pm 1.10	84.5 \pm 1.90	72.8 \pm 2.35	87.4 \pm 1.24	78.1
w/o Relevance	83.7 \pm 1.05	44.4 \pm 3.43	88.2 \pm 1.85	82.3 \pm 2.70	73.6 \pm 2.15	87.9 \pm 0.95	67.4
w/o Diversity	84.8 \pm 1.80	45.2 \pm 3.15	90.7 \pm 1.30	86.1 \pm 2.35	74.9 \pm 1.90	88.7 \pm 0.65	72.9
w/o Cluster	84.2 \pm 1.12	43.9 \pm 2.95	89.1 \pm 0.92	83.4 \pm 1.32	73.8 \pm 2.05	87.4 \pm 0.88	68.5
w/o Mask	84.8 \pm 1.06	45.7 \pm 2.80	89.9 \pm 0.85	85.8 \pm 1.26	75.3 \pm 1.95	88.4 \pm 0.80	66.7
KM-Cluster	85.4 \pm 0.99	46.1 \pm 2.65	91.5 \pm 0.78	85.7 \pm 1.20	75.8 \pm 1.85	88.6 \pm 0.74	66.4
Ours	86.6 \pm 0.83	47.8 \pm 3.21	92.7 \pm 0.38	86.9 \pm 1.13	76.4 \pm 2.12	90.5 \pm 0.59	66.3

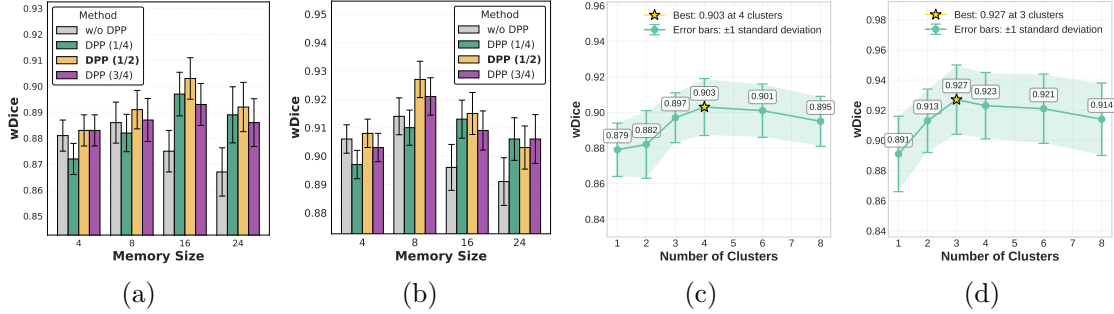


Figure 4: Analysis of the memory size on ESD dataset (a) and CHO dataset (b) and the cluster numbers on ESD dataset (c) and CHO dataset (d).

We further assessed the method’s robustness by testing different hyperparameters, specifically the ratio of features selected via DPP (1/4, 1/2, 3/4, and all features) from memories of size 4, 8, 16, and 24. Results in Fig. 4 (a-b) show that selecting half of the features performs best for most sizes. Performance without selection declines markedly with larger memories, suggesting the cross-attention module cannot fully utilize critical information from long sequences. We also studied the effects of cluster number in Fig. 4 (c-d). The performance did not improve with more clusters, peaking instead at 4 for the ESD dataset and 3 for the CHO dataset. These optimal numbers align with the actual number of surgical phases present in the videos.

5. Conclusion

In this paper, we propose a temporal memory enhancement method for surgical video semantic segmentation. The method comprises a local memory mechanism that dynamically selects diverse and relevant features via a greedy DPP, and a global memory module based on CVAE-MoG to model phase-specific patterns. By integrating both local and global temporal contexts, our approach effectively handles the varying motion rates and structured workflow inherent in surgical videos. Extensive evaluations on the ESD and CHO datasets show that our method outperforms state-of-the-art video segmentation methods, demonstrating its capability to manage complex temporal contexts in surgical video analysis.

References

- Nima Anari, Yang P Liu, and Thuy-Duong Vuong. Optimal sublinear sampling of spanning trees and determinantal point processes via average-case entropic independence. *SIAM Journal on Computing*, 54(5):FOCS22–93, 2025.
- Jianfeng Cao, Hon-Chi Yip, Yueyao Chen, Markus Scheppach, Xiaobei Luo, Hongzheng Yang, Ming Kit Cheng, Yonghao Long, Yueming Jin, Philip Wai-Yan Chiu, et al. Intelligent surgical workflow recognition for endoscopic submucosal dissection with real-time animal study. *Nature Communications*, 14(1):6676, 2023.
- Tianrun Chen, Lanyun Zhu, Chaotao Deng, Runlong Cao, Yan Wang, Shangzhan Zhang, Zejian Li, Lingyun Sun, Ying Zang, and Papa Mao. Sam-adapter: Adapting segment anything in underperformed scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3367–3375, 2023.
- Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022.
- Ho Kei Cheng and Alexander G. Schwing. XMem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European Conference on Computer Vision (ECCV)*, 2022.
- Xiaolong Deng, Huisi Wu, Runhao Zeng, and Jing Qin. Memsam: Taming segment anything model for echocardiography video segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9622–9631, 2024.
- Devanish et al. A fine-tuned foundational model SurgiSAM2 for surgical video anatomy segmentation and detection. *Scientific Reports*, 15(35961), 2025.
- Guillaume Gautier, Guillermo Polito, Rémi Bardenet, and Michal Valko. Dppy: Dpp sampling with python. *Journal of Machine Learning Research*, 20(180):1–7, 2019.
- Yihan He, Guanbin Li, Yizhou Yu, Jie Yang, and Ling Shao. An end-to-end 3d convolutional neural network for action detection and segmentation in videos. *arXiv preprint arXiv:1712.01111*, 2017.
- Zhuxi Jiang, Yin Zheng, Huachun Tan, Bangsheng Tang, and Hanning Zhou. Variational deep embedding: an unsupervised and generative approach to clustering. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence*, page 1965–1972, 2017.
- Yueming Jin, Yang Yu, Cheng Chen, Zixu Zhao, Pheng-Ann Heng, and Danail Stoyanov. Exploring intra- and inter-video relation for surgical semantic scene segmentation. *IEEE Transactions on Medical Imaging*, pages 1–1, 2022.
- Diederik P Kingma, Max Welling, et al. Auto-encoding variational bayes, 2013.

- Simon Kohl, Bernardino Romera-Paredes, Clemens Meyer, Jeffrey De Fauw, Joseph R Led-
sam, Klaus Maier-Hein, SM Eslami, Danilo Jimenez Rezende, and Olaf Ronneberger. A
probabilistic u-net for segmentation of ambiguous images. *Advances in neural information
processing systems*, 31, 2018.
- Derui Li, Yan Hu, Junyong Shen, Luoying Hao, and Jiang Liu. Semi-supervised surgical
video semantic segmentation with cross supervision of inter-frame. In *2023 IEEE 20th
International Symposium on Biomedical Imaging (ISBI)*, pages 1–5, 2023. doi: 10.1109/
ISBI53787.2023.10230375.
- Guorun Li, Lei Liu, Yuefeng Du, Peng Liu, Xiaoyu Li, Taiguo Qi, and Zhi Qiao. A recurrent-
attention mechanism for medical image segmentation. *Pattern Recognition*, 156:110818,
2025.
- Haofeng Liu, Erli Zhang, Junde Wu, Mingxuan Hong, and Yueming Jin. Surgical sam 2:
Real-time segment anything in surgical video by efficient frame pruning. *arXiv preprint
arXiv:2408.07931*, 2024.
- Haofeng Liu, Ziyue Wang, Sudhanshu Mishra, Mingqi Gao, Guanyi Qin, Chang Han Low,
Alex Y. W. Kong, and Yueming Jin. SAM2S: Segment anything in surgical videos via
semantic long-term tracking. *arXiv preprint*, 2025a.
- Yong Liu, Ran Yu, Fei Yin, Xinyuan Zhao, Wei Zhao, Weihao Xia, Jiahao Wang, Yitong
Wang, Yansong Tang, and Yujiu Yang. Learning high-quality dynamic memory for video
object segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*,
2025b.
- Pietro Mascagni, Deepak Alapatt, Aditya Murali, Armine Vardazaryan, Alain Garcia, Nar-
iaki Okamoto, Guido Costamagna, Didier Mutter, Jacques Marescaux, Bernard Dalle-
magne, et al. Endoscapes, a critical view of safety and surgical scene segmentation
dataset for laparoscopic cholecystectomy. *Scientific Data*, 12(1):1–8, 2025.
- Aditya Murali, Deepak Alapatt, Pietro Mascagni, Armine Vardazaryan, Alain Garcia, Nar-
iaki Okamoto, Guido Costamagna, Didier Mutter, Jacques Marescaux, Bernard Dalle-
magne, et al. The endoscapes dataset for surgical scene segmentation, object detection,
and critical view of safety assessment: Official splits and benchmark. *arXiv preprint
arXiv:2312.12429*, 2023.
- David Nilsson and Cristian Sminchisescu. Semantic video segmentation by gated recurrent
flow propagation. *arXiv preprint arXiv:1612.08871*, 2017.
- Seoung Wug Oh et al. Video object segmentation using space-time memory networks. In
Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV),
2019.
- Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma,
Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment
anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.

- Carole H Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support (MICCAI workshop)*, pages 240–248, 2017.
- Syed Hesham Syed Ariff, Yun Liu, Guolei Sun, Jing Yang, Henghui Ding, Xue Geng, and Xudong Jiang. Evaluating sam2 for video semantic segmentation. *arXiv preprint arXiv:2512.01774*, 2025.
- Hao Wang, Weining Wang, and Jing Liu. Temporal memory attention for video semantic segmentation. In *2021 IEEE International Conference on Image Processing (ICIP)*, pages 2254–2258. IEEE, 2021.
- Hongqiu Wang, Guang Yang, Shichen Zhang, Jing Qin, Yike Guo, Bo Xu, Yueming Jin, and Lei Zhu. Video-instrument synergistic network for referring video instrument segmentation in robotic surgery. *IEEE Transactions on Medical Imaging*, 43(12):4457–4469, 2024.
- Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in neural information processing systems*, 34:12077–12090, 2021.
- Cheng-Yen Yang, Hsiang-Wei Huang, Wenhao Chai, Zhongyu Jiang, and Jenq-Neng Hwang. Samurai: Adapting segment anything model for zero-shot visual tracking with motion-aware memory. *arXiv preprint arXiv:2411.11922*, 2024.
- Hongzheng Yang, Cheng Chen, Yueyao Chen, Hon Chi Yip, and DOU QI. Uncertainty estimation for safety-critical scene segmentation via fine-grained reward maximization. *Advances in neural information processing systems*, 36:36238–36249, 2023.
- Ming Yin et al. Memory-augmented sam2 for training-free surgical video segmentation. *arXiv preprint arXiv:2507.09577*, 2025. Accepted to MICCAI 2025.
- Jiafan Zhuang, Zilei Wang, and Yuan Gao. Semi-supervised video semantic segmentation with inter-frame feature reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3263–3271, 2022.
- Jiafan Zhuang, Zilei Wang, Yixin Zhang, and Zhun Fan. Infer from what you have seen before: Temporally-dependent classifier for semi-supervised video segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3575–3584, 2024.

Appendix A. Greedy Selection Algorithm

Algorithm 1 Greedy Selection (Orthogonal Matching Pursuit)

Initialize: selected set: $\mathcal{B} = \emptyset$, weighted features: $v_t = r_t \cdot \bar{f}_{-t}$, kernel matrix: $K_{ij} = v_i^\top v_j$ for $i, j = 1, \dots, T$

For $j = 1$ **to** M :

1. $t^* = \arg \max_{t \notin \mathcal{B}} K_{tt}$ (Select the item with the largest gain)
2. $\mathcal{B} \leftarrow \mathcal{B} \cup \{f_{-(t^*)}\}$
3. **If** $j < M$:
 $K \leftarrow K - \frac{1}{K_{t^*t^*}} K_{:,t^*} K_{:,t^*}^\top$ (Project onto orthogonal complement)

Return \mathcal{B}

Appendix B. Derivation of the evidence Lower Bound

B.1. Derivation of Equation (4)

We start with the log-likelihood:

$$\log p(y|x) = \log \int \sum_{c=1}^C p(y, z, c|x) dz.$$

Introduce a variational distribution $q(z, c|x, y)$ (approximate posterior) over latent variables z, c

$$\log p(y|x) = \log \int \sum_{c=1}^C q(z, c|x, y) \cdot \frac{p(y, z, c|x)}{q(z, c|x, y)}.$$

This can be written as:

$$\log p(y|x) = \log \mathbb{E}_{q(z, c|x, y)} \left[\frac{p(y, z, c|x)}{q(z, c|x, y)} \right].$$

By Jensen's inequality (since log is concave):

$$\log \mathbb{E}_{q(z, c|x, y)} \left[\frac{p(y, z, c|x)}{q(z, c|x, y)} \right] \geq \mathbb{E}_{q(z, c|x, y)} \left[\log \frac{p(y, z, c|x)}{q(z, c|x, y)} \right].$$

Therefore:

$$\log p(y|x) \geq \mathbb{E}_{q(z, c|x, y)} \left[\log \frac{p(y|z, x) p(z, c|x)}{q(z, c|x, y)} \right].$$

Factorize the joint distribution $p(y, z, c|x) = p(y, z|x) p(z, c|x)$ and split the LHS:

$$\log p(y|x) \geq \mathbb{E}_{q(z, c|x, y)} [\log p(y|z, x)] + \mathbb{E}_{q(z, c|x, y)} \left[\log \frac{p(z, c|x)}{q(z, c|x, y)} \right].$$

Since the second term is negative KL divergence, we have

$$\log p(y|x) \geq \mathbb{E}_{q(z, c|x, y)} [\log p(y|z, x)] - \mathcal{D}_{KL}[q(z, c|x, y) \| p(z, c|x)].$$

B.2. Derivation of KL Divergence in Equation (8)

Start with definition:

$$\mathcal{D}_{KL}[q(z, c|x, y)||p(z, c|x)] = \sum_{c=1}^C \int q(z, c|x, y) \log \frac{q(z, c|x, y)}{p(z, c|x)} dz..$$

Factorize the joint distributions:

$$= \sum_{c=1}^C \int q(z|x, y)q(c|x, y) \log \frac{q(z|x, y)q(c|x, y)}{p(z|c, x)p(c|x)} dz.$$

Substitute the defined distributions:

$$= \sum_{c=1}^C \pi_c^q \left[\int \mathcal{N}(z|\mu_q, \sigma_q^2 I) \log \frac{\mathcal{N}(z|\mu_q, \sigma_q^2 I)}{\mathcal{N}(z|\mu_c, \sigma_c^2 I)} dz + \log \frac{\pi_c^q}{\pi_c^p} \right].$$

Final expression:

$$= \sum_{c=1}^C \pi_c^q \left(\log \frac{\sigma_c}{\sigma_c} + \frac{\sigma_c^2 + (\mu_c - \mu_c)^2}{2\sigma_c^2} - \log \pi_c^p \right).$$