

Where It Moves, It Matters: Referring Surgical Instrument Segmentation via Motion

Meng Wei¹, Kun Yuan^{2,3,5*}, Shi Li³, Yue Zhou², Long Bai⁴, Nassir Navab², Hongliang Ren⁴,
Hong Joo Lee^{2,5*}, Tom Vercauteren^{1†}, Nicolas Padoy^{3†}

¹Biomedical Engineering and Imaging Sciences, King’s College London

²Technical University of Munich

³ICube, University of Strasbourg

⁴Department of Electrical Engineering, The Chinese University of Hong Kong

⁵Munich Center for Machine Learning

{meng.wei, tom.vercauteren}@kcl.ac.uk, {yue.zhou, nassir.navab, hongjoo.lee, kun.yuan}@tum.de
shi.li@ext.ihu-strasbourg.eu, {b.long@link, hlren@ee}.cuhk.edu.hk, npadoy@unistra.fr

Abstract

Enabling intuitive, language-driven interaction with surgical scenes is a critical step toward intelligent operating rooms and autonomous surgical robotic assistance. However, the task of referring segmentation, localizing surgical instruments based on natural language descriptions, remains underexplored in surgical videos, with existing approaches struggling to generalize due to reliance on static visual cues and predefined instrument names. In this work, we introduce *SurgRef*, a novel motion-guided framework that grounds free-form language expressions in instrument motion, capturing how tools move and interact across time, rather than what they look like. This allows models to understand and segment instruments even under occlusion, ambiguity, or unfamiliar terminology. To train and evaluate *SurgRef*, we present *Ref-IMotion*, a diverse, multi-institutional video dataset with dense spatiotemporal masks and rich motion-centric expressions. *SurgRef* achieves state-of-the-art accuracy and generalization across surgical procedures, setting a new benchmark for robust, language-driven surgical video segmentation.

Introduction

Surgical data science presents a critical and high-stakes application field for computer vision, where the model’s performance has a direct impact on patient safety and clinical outcomes. A core challenge in this field is the analysis of endoscopic video data, which requires recognizing surgical instruments (Twinanda et al. 2016; Allan et al. 2019, 2020), anatomical structures (Carstens et al. 2023), and procedural actions (Nwoye et al. 2022). Such capabilities enable real-time comprehension of surgical workflows, support skill assessment (Gupta et al. 2024), and facilitate complication prediction (Ma et al. 2022) and robotic navigation (Long et al. 2025). Central to these functions is the task of surgical video object segmentation, which provides precise, pixel-level identification of surgical instruments and anatomies.

*Corresponding Author.

†These authors contributed equally.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Despite the progress, recent studies have demonstrated that models pretrained on general-purpose datasets (Ding et al. 2023; Cheng et al. 2022; Huang et al. 2025a) often fail to generalize in surgical settings due to domain-specific complexities such as visual occlusion, similar-looking tools, and challenging imaging conditions, e.g., bleeding, smoke, and poor lighting (Nwoye et al. 2025; Sun et al. 2023a; Sun 2024). This highlights the importance of surgical domain-specific datasets and methods tailored to the unique visual and temporal characteristics of surgical scenes.

In this work, we move beyond conventional surgical segmentation (Wei et al. 2025) by tackling surgical referring segmentation, a task tailored to the surgical domain where instruments are localized based on static language cues (Wang et al. 2024a; Liu et al. 2025) such as appearance and spatial context (e.g., “the grasper on the left”), and dynamic motion-centric language expressions (e.g., “the tool grasping gallbladder”), as illustrated in Figure 1. This enables intuitive interaction between surgeons and AI systems. When integrated with AR, it supports surgical training by allowing users to query instruments or actions using language commands (Killeen et al. 2024). Intraoperatively, it acts as an intelligent assistant, highlighting key targets, answering verbal queries, and guiding surgeons (Seenivasan et al. 2022; Yuan et al. 2024; Huang et al. 2025b). Also, it supports natural human-robot collaboration through language-based procedural commands in robotic surgery (e.g., “cut where the grasper is holding”), allowing for fine-grained semantic control for surgery automation (Long et al. 2025).

However, enabling robust referring segmentation in surgical settings presents unique visual and linguistic challenges. Specifically, surgical scenes are complex and dynamic due to the visual occlusions, deformable anatomical structures, and the simultaneous presence of multiple similar-looking instruments. These factors introduce significant ambiguity when localizing target objects based solely on visual appearance. Linguistically, surgical language expressions are highly specialized and lack standardization across institutions and regions (Lavanchy et al. 2024). The same instru-

ment may be described in different ways depending on local surgical protocols, which limits the generalization capability of existing language-conditioned models (Yuan et al. 2025a,b). Existing methods often depend on static language expressions (Wang et al. 2024a), such as explicit instrument names (e.g., “Prograsp forceps”) for grounding, which restricts their robustness in handling unseen procedures or adapting to different clinical centers.

In contrast, instrument motion provides a more consistent and interpretable signal. Surgical procedures are defined by sequences of tool actions, including entry paths, retraction patterns, and tool-tissue interactions. These motions are important to both procedural understanding and surgical training. For example, the process of achieving the Critical View of Safety (CVS) in laparoscopic cholecystectomy (Murali et al. 2022) involves a well-defined series of instrument gestures that experienced surgeons can consistently recognize, regardless of where they were trained.

Therefore, we propose SurgRef, a motion-guided referring video segmentation framework that grounds natural language expressions in surgical video by explicitly modeling instrument motion. In addition to existing methods that rely on static appearance cues or explicit instrument names, SurgRef also interprets motion-centric expressions (e.g., “the tool entering from the right and retracting the gallbladder medially”) to produce fine-grained, temporally localized segmentations. To support this, we construct the Ref-IMotion dataset, which aggregates three public datasets to cover diverse surgical procedures and settings: laparoscopic cholecystectomy from Cholec80 (Twinanda et al. 2016), robot-assisted porcine surgery from EndoVis (Allan et al. 2019), and robotic-assisted laparoscopic radical prostatectomy from GraSP (Ayobi et al. 2024). Its composition spans both laparoscopic and robotic modalities, sourced from different institutions, enabling a comprehensive benchmark for evaluating model generalization across varied clinical domains and toolsets. For datasets lacking natural language supervision, we manually annotate high-quality motion-centric referring language expressions that describe instrument behavior such as entry trajectories, retraction patterns, and tool-tissue interactions (e.g., “the tool that enters from the top and pulls the gallbladder to the left”). The dataset includes 319 surgical video clips with a total of 21,350 annotated frames and 718 referring expressions, including 358 motion-based expressions, each paired with dense frame-level masks and temporally localized intervals.

In this work, SurgRef enhances temporal reasoning by our proposed key-frame attention module that adaptively selects a subset of expression-relevant frames by leveraging language-guided object semantics. Instead of uniformly processing all frames, the module computes frame-wise relevance scores from decoder-level object queries conditioned on the referring expression. By selecting the top-ranked frames for segmentation, this module enables the model to focus on temporally salient moments, suppressing redundancy and emphasizing frames with meaningful object interactions. This motion-centric design improves both temporal efficiency and the precision of spatial-temporal grounding in long surgical videos. SurgRef seg-

ments the spatial regions of referred objects while identifying expression-relevant frames through language-guided key-frame selection, achieving state-of-the-art performance in surgical video referring segmentation. Its motion-guided representations also exhibit strong generalization across diverse surgical procedures and toolsets. Our main contributions are summarized as follows:

- We construct the Ref-IMotion dataset by annotating multiple public datasets with diverse surgical procedures, and manually curating motion-based referring expressions paired with dense spatial-temporal masks.
- We propose SurgRef, a motion-guided referring video segmentation framework that grounds both static and motion-centric language expressions in surgical videos, enabling fine-grained spatial-temporal understanding.
- Extensive experiments demonstrate that SurgRef achieves state-of-the-art performance and strong generalization across varied surgical toolsets and procedures, setting a robust benchmark for motion-guided video segmentation in surgery.

Related Work

Video Segmentation. Video segmentation aims to assign pixel-wise labels to each frame in a video sequence, enabling the consistent tracking and delineation of objects or regions of interest over time (Oh et al. 2019; Yang, Wei, and Yang 2021; Cheng and Schwing 2022; Yang and Yang 2022). Unlike image segmentation, it accounts for temporal continuity, motion dynamics, occlusions, and appearance variations, making it a more challenging task. Since it plays a critical role in various applications, such as autonomous driving (Siam et al. 2018; Kim, Yim, and Kim 2020; Muhammad et al. 2022; Sun et al. 2023b) and surgical scene understanding (Ni et al. 2020; Ayobi et al. 2023; Wang et al. 2024b; Zhou et al. 2023), extensive studies have been conducted by exploiting spatiotemporal feature modeling, memory-based architectures, and attention mechanisms to improve temporal consistency and accuracy across frames. For example, Oh et al. (Oh et al. 2019) proposed STM. They introduced a memory network that retrieves useful information from past frames to guide segmentation in the current frame, improving temporal consistency. Also, Cheng et al. (Cheng and Schwing 2022) proposed XMem. It extends the STM by using a dynamic memory update strategy that allows scalable and efficient segmentation over long video sequences. Both methods highlight the effectiveness of memory-based architectures for robust video object segmentation.

Surgical Video Segmentation. Surgical video segmentation has been extensively studied as a core task in computer-assisted and robotic surgery. It serves as a foundation for a variety of downstream applications, including tool tracking, workflow analysis, surgical skill assessment, and autonomous navigation. Most prior work has focused on vision-based models trained with pixel-level supervision (Ni et al. 2020; Ayobi et al. 2023, 2024; Grammatikopoulou et al. 2024; Liao et al. 2025). For example, Ni et al. (Ni et al. 2020) introduced LWANet, an attention-guided lightweight

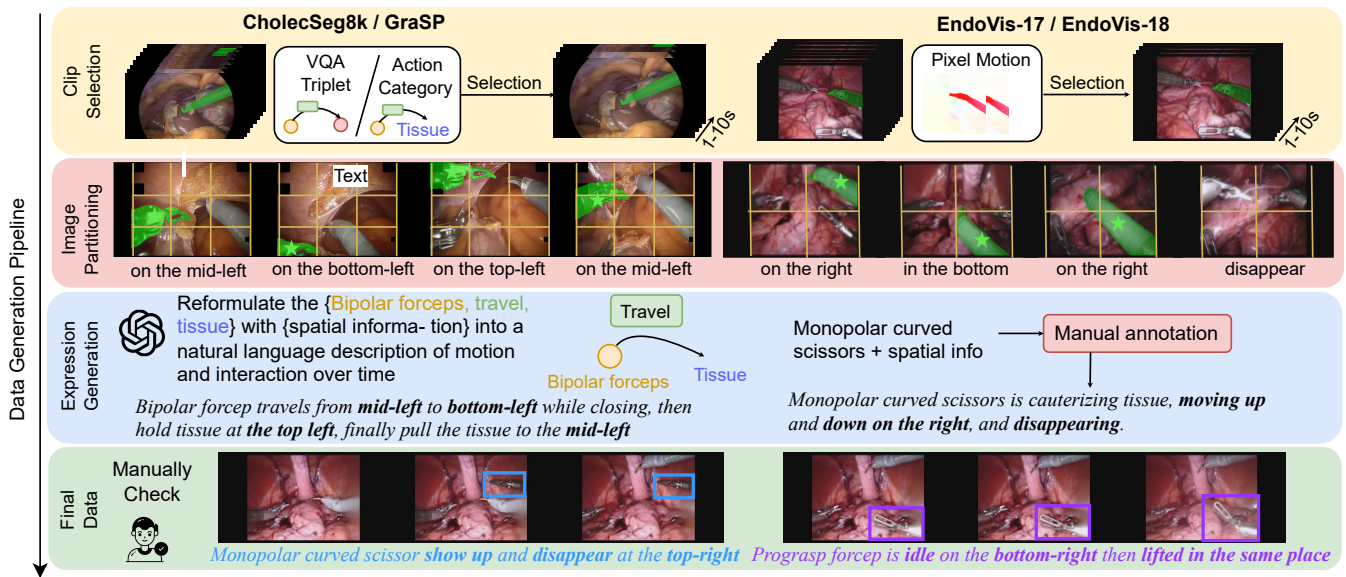


Figure 1: Data construction process of the EndoVis-IM17, EndoVis-IM18, GraSP-IM, and CholecSeg8k-IM.

architecture that employs depth-wise separable convolutions to enable real-time inference with reduced computational cost. Ayobi et al. (Ayobi et al. 2023) proposed MATIS, a transformer-based framework that combines pixel-wise attention, masked attention modules focused on instrument regions, and video transformers for handling short-range temporal dependencies.

Surgical Video Referring Segmentation. Beyond conventional vision-based models, vision-language approaches is the new trend for surgical instrument segmentation because of their human-understandable interactions. Wang et al. (Wang et al. 2024b) introduced VIS-Net, the first method for referring instrument segmentation in the surgical domain. Their approach incorporates a video-instrument synergistic network to jointly learn video-level and instrument-level representations. Zhou et al. (Zhou et al. 2023) proposed TP-SIS, a text-promptable segmentation framework that integrates pretrained vision-language models and utilizes a mixture-of-prompts mechanism to improve robustness.

Despite recent progress, existing methods fall short in two key aspects: they lack long-term temporal modeling and overly rely on static appearance cues. This overlooks motion, the most consistent and discriminative signal in surgical scenes, especially when tools are visually similar or occluded. Without motion-aware modeling, current approaches struggle to robustly localize and segment instruments in complex, real-world surgical scenarios.

Surgical Instrument Segmentation Dataset. Several benchmarks have been developed for surgical instrument segmentation, including EndoVis-17, EndoVis-18, Cholec80 (Hong et al. 2020) and GraSP (Ayobi et al. 2024). The EndoVis-17 dataset consists of stereo video sequences acquired from the da Vinci robotic system, annotated with binary, part-based, and type-level segmentation masks for

multiple surgical instruments. EndoVis-18 extends this to full-scene segmentation, providing labels for both instruments and anatomical structures across longer and more varied procedures. CholecSeg8k is a semantic segmentation dataset derived from the Cholec80 surgical video benchmark (Twinanda et al. 2016), offering dense annotations for surgical instruments. GraSP is a benchmark dataset designed to support surgical scene understanding through a hierarchy of complementary tasks at varying levels of granularity. It includes pixel-level instance annotations, as well as instrument and action category labels.

These datasets lack fine-grained motion annotations, only providing object identity, spatial location, and short-term action descriptions without describing the trajectory, direction, and motion of instruments. This limits their applicability for training or evaluating models that aim to understand and leverage motion cues in referring segmentation tasks.

Methodology

Dataset Construction

We tackle the task of segmenting surgical objects in a video guided by motion-centric language expressions. To this end, we construct the motion-guided referring dataset, i.e., Ref-IMotion, by annotating publicly available surgical videos, EndoVis-17, EndoVis-18, CholecSeg8k, and GraSP, with rich motion-based textual descriptions.

For each dataset, we follow its own preprocessing pipelines to extract high-quality, mask-annotated frames. From these, we select clips that exhibit high motion complexity, including instrument appearances, disappearances, and interactions. Specifically, we extract frames from EndoVis-17, EndoVis-18, and GraSP where tool interactions are detected via abrupt pixel intensity changes or rapid motion. For CholecSeg8k, we use scene graph triplets from SSG-VQA (Yuan et al. 2024) to locate frames with distinct

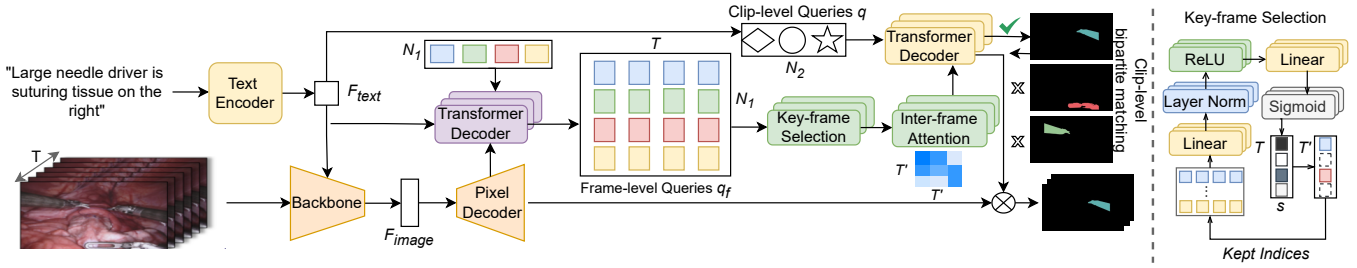


Figure 2: Baseline model for SurgRef, which integrates a Swin Transformer backbone for spatial feature extraction, a frozen RoBERTa encoder for language grounding, a language-guided transformer decoder for cross-modal reasoning, and a key-frame attention module for selecting semantically aligned frames.

tool-object interactions. Frames are grouped into short clips based on visual continuity and action diversity, resulting in a curated set of 319 video clips totaling 21,350 frames, the largest video referring segmentation dataset in surgical computer vision to date.

Our core contribution lies in introducing motion-centric expressions that capture not only the presence and action of surgical instruments, but also their spatiotemporal dynamics. To reflect this focus on motion, we prefix our newly annotated datasets with “IM.”

- For EndoVis-IM17 (EV-IM17), we annotate expressions aligned with motion trajectories across frame sequences (e.g., “Grasping retractor appears from the left, moves to the top-right, and disappears”), supplemented with coarse spatial references using a 2×2 grid.
- EndoVis-IM18 (EV-IM18) features attribute-based motion expressions (e.g., “Suction instrument is suctioning blood”) that focus on tool action and spatial movement.
- For CholecSeg8k-IM (C8K-IM), we augment existing segmentation masks with semantic triplets and 3×3 grid-based spatial tags to construct expressions such as “Grasper is grasping gallbladder in the mid-left.” While some triplets are adapted from SSG-VQA (Yuan et al. 2024), we perform manual verification and correction to ensure consistency and domain accuracy.
- To construct GraSP-IM (GSP-IM), we first selected surgical videos based on specific procedure phases and steps, assigning each video a unique identifier in the format *case-phase-step*. Using the previously defined action categories, we generated a corresponding motion expression for each object in each video segment. Object instances were identified from the annotations, and their bounding boxes were extracted accordingly. The key point of each object was defined as the center of its bounding box. To estimate motion direction, we calculated the displacement of this point across consecutive frames. Each frame was divided into a 3×3 spatial grid, allowing us to localize movements within spatial regions and combine them with action labels. We then used a prompt to feed the structured inputs, comprising instrument, action, object, and spatial information, into a large language model (OpenAI 2025) to generate natural language descriptions of motion and interaction over time.

Note that we merged video segments into their neighboring steps if the number of images was fewer than 10, resulting in 258 unique videos and 10,944 frames in total.

Together, these curated annotations form the existing publicly available datasets that compose our Ref-IMotion, supporting fine-grained spatio-temporal language grounding and advancing motion-aware surgical scene understanding.

Each surgical video consists of T RGB frames, denoted as $\{I_t\}_{t=1}^T$, where each frame $I_t \in \mathbb{R}^{H \times W \times 3}$ with spatial resolution $H \times W$. To extract spatial features from each frame, we employ a vision transformer backbone based on the Swin Transformer (Liu et al. 2021), yielding a multi-scale feature representation for each frame. Let F_{image} denote the extracted visual features from frame I_t .

Given a natural language expression that refers to surgical activity (e.g., “Grasper is retracting tissue on the right”), our goal is to perform referring video object segmentation: producing a segmentation mask for the object(s) referred to by the expression in each frame. To interpret the referring expression, we use a frozen RoBERTa-base (Liu et al. 2019) model as our text encoder. This encoder transforms the unstructured language input into a dense semantic embedding. Let $F_{text} \in \mathbb{R}^d$ denote the global sentence-level embedding extracted from RoBERTa.

Unlike the object queries in mask2former (Cheng et al. 2022), SurgRef initializes N_2 language-driven queries $\{q_i^{(0)}\}_{i=1}^{N_2}$ by injecting the language embedding F_{text} into the initial query vectors:

$$q_i^{(0)} = W_{init} \cdot F_{text} + b_i, \quad (1)$$

where $W_{init} \in \mathbb{R}^{d_q \times d}$ is a learnable projection matrix and $b_i \in \mathbb{R}^{d_q}$ is a learnable positional bias for the i -th query. Our transformer decoder is based on the Mask2Former (Cheng et al. 2022) architecture, adapted to incorporate language guidance. Specifically, we first replace standard object queries with language-driven initializations, then inject textual semantics into the cross-attention layers for spatial reasoning. After L transformer decoder layers, the final query representations $\{q_i^{(L)}\}_{i=1}^{N_2}$ are passed through two parallel heads:

Classification Head. Each query produces a classification score indicating whether it corresponds to a relevant object:

Dataset	Static Attr.	Motion	Total	Motion %
EndoVis-IM17	119	45	164	27.43
EndoVis-IM18	137	21	158	13.29
CholecSeg8k-IM	104	34	138	24.63
GraSP-IM	N/A	258	258	100
Ref-IMotion (total)	360	358	718	49.86

Dataset	Video-based	TC	RE	Inst.M	SE
EndoVis-RS17/18 (Wang et al. 2024a)	✓	✗	✓	✓	RGB color ^a
CholecSeg8k (Hong et al. 2020)	✗	✗	✗	✓	✗
SSG-VQA (Yuan et al. 2024)	✗	✗	Triplet	✗	✗
Ref-IMotion	✓	✓	✓	✓	2×2/3×3 ^b

Table 1: Left: Ref-IMotion expression statistics (Attr.: Attribute). Right: Comparison of multimodal-temporal properties with other surgical datasets (TC: Temporal Continuity; RE: Referring Expression; Inst.M: Instance Mask; SE: Spatial encoding). ^aColor-coded segmentation masks for instrument instances. ^bText-based spatial grid descriptions for instrument positioning.

$$\hat{y}_i^{\text{cls}} = \text{Linear} \left(\left[q_i^{(L)} \parallel F_{\text{text}} \right] \right) \in \mathbb{R}^{C+1}, \quad (2)$$

where \parallel denotes vector concatenation, and C is the number of object classes (with one additional background class).

Mask Embedding Head. Each query also generates a mask embedding:

$$\hat{y}_i^{\text{mask}} = \text{MLP}(q_i^{(L)}) \in \mathbb{R}^{d_m}. \quad (3)$$

Given the pixel-level mask features $F_{\text{mask}} \in \mathbb{R}^{d_m \times H \times W}$, the binary segmentation mask \hat{M}_i for query i is computed by a dot product and sigmoid activation:

$$\hat{M}_i = \sigma(e_i^\top \cdot F_{\text{mask}}), \quad (4)$$

where $e_i = \hat{y}_i^{\text{mask}}$ is the i -th mask embedding, and σ denotes the sigmoid function. To determine valid object-mask pairs, we apply a threshold on classification confidence. All queries with $\max_{c=1}^C \text{softmax}(\hat{y}_i^{\text{cls}})_c > \tau$ (where $\tau = 0.8$) are retained. This thresholding strategy allows the model to flexibly handle both single-object and multi-object expressions, which is essential for parsing complex surgical scenes.

Key-frame Selection. Surgical videos contain significant temporal redundancy, with many frames showing repetitive activity. To improve temporal reasoning and reduce computational cost, we introduce a key-frame selection module that selects expression-relevant frames using language-guided queries.

For each video clip, the transformer decoder generates language-guided frame-level queries $\mathbf{q}_f \in \mathbb{R}^{C_Q \times N_1}$ of dimension C_Q from the transformer decoder, which incorporate both spatial and linguistic context, yielding a tensor of size $[T, N_1, C_Q]$. For each frame I_t , we extract its frame-level representation $\mathbf{e}_t \in \mathbb{R}^{C_Q}$ by aggregating the corresponding object queries, resulting in embeddings of size $[T, C_Q]$. We then compute a scalar relevance score $s_t \in [0, 1]$ for each frame using a lightweight MLP applied to \mathbf{e}_t :

$$s_t = \sigma(\mathbf{W}_2 \cdot \text{ReLU}(\mathbf{W}_1 \cdot \mathbf{e}_t) + \mathbf{b}), \quad (5)$$

where $\mathbf{W}_1 \in \mathbb{R}^{d \times C_Q}$, $\mathbf{W}_2 \in \mathbb{R}^{1 \times d}$ are learnable projection matrices, \mathbf{b} is a bias term, and $\sigma(\cdot)$ denotes the sigmoid function. These scores measure the visual-text alignment between each frame and the referring expression.

After computing relevance scores $\{s_t\}_{t=1}^T$ for all frames, we select the top- T' frames with the highest scores, maintaining temporal order. Only these key frames are passed to

the segmentation decoder for final mask prediction. For example, for the expression ‘‘scissors traveling,’’ frames showing active motion score > 0.8 , while idle frames score < 0.3 . This strategy allows the model to focus computation on frames that are semantically aligned with the referring expression, improving both efficiency and expression-specific segmentation performance in long surgical sequences.

Inter-frame Attention. After key-frame selection, we apply the inter-frame attention (Ding et al. 2023) to aggregate temporal motion information across the selected key frames. The language-guided frame-level queries \mathbf{q}_f now yield a tensor of size $[T', N_1, C_Q]$, where T' denotes the number of selected key frames and N_1 is the number of queries per frame. We flattend \mathbf{q}_f into $\mathbf{q}'_f \in \mathbb{R}^{(T' \cdot N_1) \times C_Q}$ then apply multi-head self-attention followed by feed-forward networks with residual connections, which enables queries to exchange information across the temporal dimension, capturing object motion. Then \mathbf{q}'_f is reshaped back and passed to the transformer decoder for mask generation.

Training Loss. The model is trained with a composite loss that combines frame-level loss, which includes classification loss (cross-entropy) and mask prediction losses (binary cross-entropy and Dice), and video-level losses, including temporal similarity loss that enforces consistency of query embeddings across adjacent frames and video-level mask prediction loss (see supplementary).

Experiments

Evaluation Metrics

We adopt J, F, J&F, Dice, and IoU as evaluation metrics. Our surgical referring video segmentation model requires tracking referred objects across the entire video, even under occlusion. Therefore, to capture both spatial and boundary accuracy, we use J (region similarity) and F (contour accuracy), with J&F denoting their mean. All metrics are computed frame-wise between predicted and ground-truth masks, averaged temporally per object, and then aggregated across all referred objects. We also report Dice and IoU for pixel-wise overlap quality. Details are in the supplementary.

Implementation Details

We use a Swin Transformer (Liu et al. 2021) as the visual backbone and a frozen RoBERTa-base (Liu et al. 2019) model as the textual encoder. We set T' to 8 for the key-frame selection module for optimal efficiency-accuracy

Dataset	Model	J	F	J&F	Dice	IoU
EV-IM17	VIS-Net	79.21	82.52	80.86	82.90	78.24
	VISA	86.85	84.93	85.89	83.39	81.57
	MPG-SAM 2	88.20	90.61	89.41	84.37	82.94
	Ours	87.11	88.29	87.70	83.14	80.16
	Ours + KFS	89.91	88.93	89.42	87.94	82.12
EV-IM18	VIS-Net	58.51	59.56	59.03	54.24	49.97
	VISA	83.46	81.96	82.71	63.77	58.72
	MPG-SAM 2	83.11	82.94	83.03	62.14	59.67
	Ours	80.08	80.45	80.27	57.86	55.23
	Ours + KFS	85.91	83.05	84.48	66.89	61.72
GSP-IM	Ours	81.36	81.67	81.52	70.63	66.92
	Ours + KFS	84.92	85.01	84.97	73.56	70.07
C8K-IM*	Ours	64.58	71.23	67.91	67.89	63.54
	Ours + KFS	66.92	74.31	70.62	68.13	66.23

Table 2: Comparison of state-of-the-art methods, our baseline model adapted from LMPM (Ding et al. 2023), and our full approach with the key-frame selection (KFS) module. Models are trained and tested separately on EndoVis-IM17, EndoVis-IM18, and GraSP-IM datasets. CholecSeg8k-IM is an unseen test set, where SurgRef is trained on GraSP-IM and evaluated zero-shot on CholecSeg8k-IM.

trade-off. The model is trained for 100,000 iterations using the AdamW optimizer with an initial learning rate of 5×10^{-5} , decayed by a cosine scheduler. The batch size is set to 8 and the weight decay to 0.05. Training takes approximately 23 hours on two NVIDIA A100 GPUs with 40 GB of memory. Implementation details are in the supplementary.

Dataset Details

Table 1 presents a comprehensive comparison of Ref-IMotion with existing referring video segmentation datasets in the general computer vision and surgical computer vision fields. It analyzes these datasets across multiple dimensions, including input modality, temporal reasoning, referring expression types, and spatial annotations. Overall, Ref-IMotion is the only dataset that integrates all these features, which enable fine-grained spatiotemporal grounding. Prior referring datasets like Ref-COCO and Ref-Youtube-VIS support temporal language grounding but are built on natural scenes and lack the visual complexity of surgical domains. Surgical datasets such as EndoVis-IM17/18 and CholecSeg8k-IM provide valuable surgical segmentation masks but contain either limited or no referring language expression, restricting their utility for motion-aware language-driven tasks. Our proposed Ref-IMotion dataset with carefully annotated motion-centric expressions that describe not just what an instrument is, but how it moves through space and time, captures spatiotemporal dynamics such as entry trajectories, retraction paths, and tool-tissue interactions. For instance, our EndoVis-IM17 subset includes trajectory-aligned motion expressions mapped to frame sequences, while CholecSeg8k-IM subset includes expressions enriched with semantic action triplets and grid-based spatial tags. Unlike attribute-only phrases (e.g., “suctioning blood”), our expressions support dynamic interac-

tion grounding (e.g., “grasper enters from the left and retracts the gallbladder medially”), which is crucial for real-world surgical understanding.

State-of-the-art Comparison

In this work, we compare to prior state-of-the-art methods from both the surgical domain (Wang et al. 2024a) and the general computer vision domain (Yan et al. 2024; Rong et al. 2025). VIS-Net (Wang et al. 2024a) is a surgery-specific model that combines spatial-temporal graph features with language-guided attention. VISA (Yan et al. 2024) and MPG-SAM 2 (Rong et al. 2025) are general-purpose frameworks that use large language models and SAM-based propagation for referring video object segmentation. As shown in Table 2, even though prior works emphasize vision-language alignment and temporal modeling, our proposed approach consistently outperforms them across multiple datasets, e.g., achieving 89.91 J on EndoVis-IM17 compared to 88.2 by MPG-SAM 2, and 84.48 J&F on EndoVis-IM18 compared to 83.03. This demonstrates our model’s superior ability to fuse motion-guided expressions with language and video cues, enabling more precise spatial-temporal localization. As shown in Figure 3, our model produces accurate spatial-temporal segmentations in challenging scenarios.

Generalizability Analysis

While prior methods demonstrate the benefits of vision-language alignment and temporal modeling, Table 2 shows that our approach not only achieves state-of-the-art performance on standard surgical benchmarks (EndoVis-IM17/18) but also exhibits strong zero-shot generalization across different surgical procedures and modalities. Specifically, our model, trained under a unified setting without any dataset-specific tuning, generalizes effectively to CholecSeg8k-IM (laparoscopic cholecystectomy) using the model trained on GraSP-IM (robot-assisted prostatectomy), both of which differ significantly from the training data in terms of anatomy, tool types, imaging modalities, and institutional sources. The consistent performance gains on the unseen dataset highlight the robustness of our motion-guided language grounding framework, which captures instrument behavior beyond static appearance or naming cues, making it broadly applicable to diverse clinical environments.

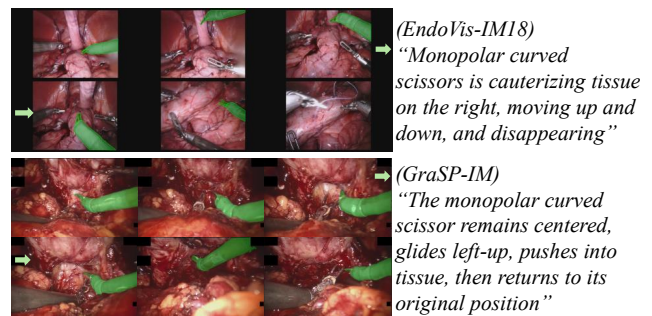


Figure 3: Visualization examples of our SurgRef on EndoVis-IM18 and GraSP-IM dataset.

Dataset	Expr.	w/o Motion			w/ Motion		
		J	F	J&F	J	F	J&F
EV-IM17	Appr.	73.92	74.58	74.25	85.42	84.91	85.17
	Spatial	76.33	76.95	76.64	87.52	87.13	87.33
	Motion	79.92	78.14	79.03	89.91	88.93	89.42
EV-IM18	Appr.	67.78	68.53	68.16	82.77	79.26	81.02
	Spatial	69.24	69.05	69.15	84.13	81.97	83.05
	Motion	71.79	70.32	71.06	85.91	83.05	84.48
GSP-IM	Appr.	68.01	68.75	68.38	81.02	81.93	81.48
	Spatial	68.92	69.88	69.40	81.53	82.74	82.14
	Motion	70.18	71.34	70.76	84.92	85.01	84.97

Table 3: Evaluation on different testing expressions (Expr.): appearance-based instrument descriptions (Appr.), instrument spatial information-based expressions (Spatial), and motion-aware temporal dynamics expressions (Motion). Performance is compared between models trained with (w/) and without (w/o) motion-aware expressions.

Motion-guided Expression Analysis

To evaluate the impact of motion-centric supervision, we compare models trained with and without motion-based expressions across different test-time referring styles: (1) appearance-based instrument descriptions (i.e., describing the name, appearance, and function of a given instrument), (2) instrument spatial information-based expressions (e.g., “on the left”), and (3) full motion-based descriptions (e.g., “the tool entering from the right and grasping the gallbladder”). As shown in Table 3, models trained with motion-aware expressions consistently outperform those trained without, across all datasets and expression styles. For example, on EndoVis-IM17, SurgRef trained with motion expressions achieves a J&F of 89.42, compared to 79.03 for non-motion-trained counterparts. Even when tested on expression with static features including appearance-based and spatial-based expressions, motion-trained models show substantial gains (e.g., 85.17 vs. 74.25 on appearance-based expressions). This consistent improvement confirms that motion-based supervision from our Ref-IMotion dataset enhances not only motion expression understanding but also generalizes to static feature descriptions, demonstrating stronger robustness in real-world surgical contexts.

Ablation Study

Expression Style. Table 4 presents a detailed analysis of model’s performance under varying referring expression styles on EndoVis-IM17. In the same-style setting, where the model is both trained and tested on the same expression type, performance is highest for original expressions (J&F = 87.70), while removing location or name cues leads to substantial drops, especially without spatial cues (J&F = 64.13), confirming the utility of such information. In the cross-style setting, performance degrades further when there’s a mismatch between training and testing styles. For example, the origin-trained model tested on expressions with no location yields a J&F of 59.66, compared to 87.70 when tested on origin expressions. Notably, the model trained without instrument names still generalizes well to origin expressions

(a) Expression style ablation					
Evaluation	Training → Testing	J	F	J&F	
Same-style	origin → origin	87.11	88.29	87.70	
	no location → no location	59.24	69.02	64.13	
	no name → no name	78.05	79.98	79.02	
Cross-style	origin → no location	57.78	61.54	59.66	
	origin → no name	62.24	64.23	63.23	
	no location → origin	59.24	69.07	64.15	
	no location → no name	34.79	42.50	38.64	
	no name → origin	78.62	80.69	79.65	
	no name → no location	41.42	43.84	42.63	
(b) Key-frame selection strategy and T' ablation					
Selection Strategy	$T'=4$	$T'=8$	$T'=16$	$T'=24$	$T'=32$
Uniform sampling	79.28	82.16	84.67	85.92	86.77
Cosine similarity	78.93	85.71	86.82	86.95	87.09
Ours	80.33	89.42	89.56	88.87	88.35

Table 4: Ablation study on EndoVis-IM17 (%). (a) Expression style generalization with the same style and cross-style for training expression and testing expression. (b) Key-frame selection strategies across different T' values (J&F scores).

(J&F = 79.65), highlighting that motion and spatial context can compensate for missing lexical identifiers. Conversely, models trained without location cues generalize poorly to other styles, particularly when both name and location cues are absent (J&F = 38.64 or 42.63). Overall, these results emphasize that while spatial and naming cues are helpful, motion-centric training enhances generalization, especially under challenging linguistic conditions.

Key-Frame Selection. As shown in Table 4, our object-level key-frame selection achieves optimal J&F of 89.42 at $T' = 8$, demonstrating superior sample efficiency. The performance saturation while T' increases indicates that our method effectively identifies the most expression-relevant frames, while additional frames introduce significant training cost without improving performance. In contrast, uniform sampling and Cosine similarity-based selections exhibit monotonic improvement, requiring more frames to capture task-relevant temporal content.

Conclusion

In this work, we introduced SurgRef, a motion-guided referring video segmentation framework that grounds natural language expressions in surgical videos by explicitly modeling instrument motion. Paired with our curated Ref-IMotion dataset, comprising motion-centric expressions and dense spatial-temporal annotations, SurgRef advances beyond static appearance-based methods by capturing dynamic tool behaviors and semantic trajectories. Our model achieves state-of-the-art performance across multiple surgical benchmarks and demonstrates strong zero-shot generalization to unseen procedures, tools, and modalities. Together, SurgRef and Ref-IMotion form a strong foundation for motion-aware, language-driven surgical understanding.

Acknowledgments

Meng Wei is supported by the UKRI EPSRC CDT in Smart Medical Imaging [EP/S022104/1]. Tom Vercauteren is a co-founder and shareholder of Hypervision Surgical. This work is supported by core funding from the Wellcome Trust / EPSRC [WT203148/Z/16/Z; NS/A000049/1] and has received funding from the European Union (ERC, COMP-SURG, 101088553). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. This work was also partially supported by French state funds managed within the Plan Investissements d’Avenir by the ANR under reference ANR-10-IAHU-02 (IHU Strasbourg).

References

- Allan, M.; Kondo, S.; Bodenstedt, S.; Leger, S.; Kadkhodamohammadi, R.; Luengo, I.; Fuentes, F.; Flouty, E.; Mohammed, A.; Pedersen, M.; et al. 2020. 2018 robotic scene segmentation challenge. *arXiv preprint arXiv:2001.11190*.
- Allan, M.; Shvets, A.; Kurmann, T.; Zhang, Z.; Duggal, R.; Su, Y.-H.; Rieke, N.; Laina, I.; Kalavakonda, N.; Bodenstedt, S.; et al. 2019. 2017 robotic instrument segmentation challenge. *arXiv preprint arXiv:1902.06426*.
- Ayobi, N.; Pérez-Rondón, A.; Rodríguez, S.; and Arbeláez, P. 2023. Matis: Masked-attention transformers for surgical instrument segmentation. In *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI)*, 1–5. IEEE.
- Ayobi, N.; Rodríguez, S.; Pérez, A.; Hernández, I.; Aparicio, N.; Dessevres, E.; Peña, S.; Santander, J.; Caicedo, J. I.; and Fernández, N. 2024. Pixel-wise recognition for holistic surgical scene understanding. *arXiv preprint arXiv:2401.11174*.
- Carstens, M.; Rinner, F. M.; Bodenstedt, S.; Jenke, A. C.; Weitz, J.; Distler, M.; Speidel, S.; and Kolbinger, F. R. 2023. The dresden surgical anatomy dataset for abdominal organ segmentation in surgical data science. *Scientific Data*, 10(1): 1–8.
- Cheng, B.; Misra, I.; Schwing, A. G.; Kirillov, A.; and Girshick, R. 2022. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1290–1299.
- Cheng, H. K.; and Schwing, A. G. 2022. Xmem: Long-term video object segmentation with an atkinson-shiffrin memory model. In *European conference on computer vision*, 640–658. Springer.
- Ding, H.; Liu, C.; He, S.; Jiang, X.; and Loy, C. C. 2023. MeViS: A large-scale benchmark for video segmentation with motion expressions. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2694–2703.
- Grammatikopoulou, M.; Sanchez-Matilla, R.; Bragman, F.; Owen, D.; Culshaw, L.; Kerr, K.; Stoyanov, D.; and Luengo, I. 2024. A spatio-temporal network for video semantic segmentation in surgical videos. *International Journal of Computer Assisted Radiology and Surgery*, 19(2): 375–382.
- Gupta, A.; Kocielnik, R.; Wang, J.; Nasriddinov, F.; Yang, C.; Wong, E.; Anandkumar, A.; and Hung, A. 2024. Multi-Modal Self-Supervised Learning for Surgical Feedback Effectiveness Assessment. *arXiv preprint arXiv:2411.10919*.
- Hong, W.-Y.; Kao, C.-L.; Kuo, Y.-H.; Wang, J.-R.; Chang, W.-L.; and Shih, C.-S. 2020. Cholecseg8k: a semantic segmentation dataset for laparoscopic cholecystectomy based on cholec80. *arXiv preprint arXiv:2012.12453*.
- Huang, N.; Zheng, W.; Xu, C.; Keutzer, K.; Zhang, S.; Kanazawa, A.; and Wang, Q. 2025a. Segment Any Motion in Videos. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 3406–3416.
- Huang, Y.; Bai, L.; Cui, B.; Yuan, K.; Wang, G.; Islam, M.; Padoy, N.; Navab, N.; and Ren, H. 2025b. SurgTPGS: Semantic 3D Surgical Scene Understanding with Text Promptable Gaussian Splatting. *arXiv preprint arXiv:2506.23309*.
- Killeen, B. D.; Chaudhary, S.; Osgood, G.; and Unberath, M. 2024. Take a shot! Natural language control of intelligent robotic X-ray systems in surgery. *International journal of computer assisted radiology and surgery*, 19(6): 1165–1173.
- Kim, B.; Yim, J.; and Kim, J. 2020. Highway driving dataset for semantic video segmentation. *arXiv preprint arXiv:2011.00674*.
- Lavanchy, J. L.; Ramesh, S.; Dall’Alba, D.; Gonzalez, C.; Fiorini, P.; Müller-Stich, B. P.; Nett, P. C.; Marescaux, J.; Mutter, D.; and Padoy, N. 2024. Challenges in multi-centric generalization: phase and step recognition in Roux-en-Y gastric bypass surgery. *International journal of computer assisted radiology and surgery*, 19(11): 2249–2257.
- Liao, G.; Jogan, M.; Koushik, S.; Eaton, E.; and Hashimoto, D. A. 2025. Disentangling spatio-temporal knowledge for weakly supervised object detection and segmentation in surgical video. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 8013–8023. IEEE.
- Liu, H.; Gao, M.; Luo, X.; Wang, Z.; Qin, G.; Wu, J.; and Jin, Y. 2025. ReSurgSAM2: Referring Segment Anything in Surgical Video via Credible Long-term Tracking. *arXiv preprint arXiv:2505.08581*.
- Liu, Y.; Ott, M.; Goyal, N.; Du, J.; Joshi, M.; Chen, D.; Levy, O.; Lewis, M.; Zettlemoyer, L.; and Stoyanov, V. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10012–10022.
- Long, Y.; Lin, A.; Kwok, D. H. C.; Zhang, L.; Yang, Z.; Shi, K.; Song, L.; Fu, J.; Lin, H.; Wei, W.; et al. 2025. Surgical embodied intelligence for generalized task autonomy in laparoscopic robot-assisted surgery. *Science Robotics*, 10(104): eadt3093.
- Ma, R.; Ramaswamy, A.; Xu, J.; Trinh, L.; Kiyasseh, D.; Chu, T.; Wong, E.; Lee, R.; Rodriguez, I.; DeMeo, G.; et al. 2022. Surgical gestures as a method to quantify surgical performance and predict patient outcomes. *NPJ Digit Med* 5 (1): 187.

- Muhammad, K.; Hussain, T.; Ullah, H.; Del Ser, J.; Rezaei, M.; Kumar, N.; Hijji, M.; Bellavista, P.; and De Albuquerque, V. H. C. 2022. Vision-based semantic segmentation in scene understanding for autonomous driving: Recent achievements, challenges, and outlooks. *IEEE Transactions on Intelligent Transportation Systems*, 23(12): 22694–22715.
- Murali, A.; Alapatt, D.; Mascagni, P.; Vardazaryan, A.; Garcia, A.; Okamoto, N.; Mutter, D.; and Padoy, N. 2022. Latent Graph Representations for Critical View of Safety Assessment. *arXiv preprint arXiv:2212.04155*.
- Ni, Z.-L.; Bian, G.-B.; Hou, Z.-G.; Zhou, X.-H.; Xie, X.-L.; and Li, Z. 2020. Attention-guided lightweight network for real-time segmentation of robotic surgical instruments. In *2020 IEEE international conference on robotics and automation (ICRA)*, 9939–9945. IEEE.
- Nwoye, C. I.; Elgohary, K.; Srinivas, A.; Zaid, F.; Lavanchy, J. L.; and Padoy, N. 2025. Cholectrack20: A multi-perspective tracking dataset for surgical tools. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 8942–8952.
- Nwoye, C. I.; Yu, T.; Gonzalez, C.; Seeliger, B.; Mascagni, P.; Mutter, D.; Marescaux, J.; and Padoy, N. 2022. Rendezvous: Attention Mechanisms for the Recognition of Surgical Action Triplets in Endoscopic Videos. *Medical Image Analysis*, 78: 102433.
- Oh, S. W.; Lee, J.-Y.; Xu, N.; and Kim, S. J. 2019. Video object segmentation using space-time memory networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9226–9235.
- OpenAI. 2025. ChatGPT (August 2025 Version). Accessed: 2025-08-01.
- Rong, F.; Lan, M.; Zhang, Q.; and Zhang, L. 2025. MPG-SAM 2: Adapting SAM 2 with Mask Priors and Global Context for Referring Video Object Segmentation. *International Conference on Computer Vision*.
- Seenivasan, L.; Islam, M.; Krishna, A. K.; and Ren, H. 2022. Surgical-vqa: Visual question answering in surgical scenes using transformer. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 33–43. Springer.
- Siam, M.; Gamal, M.; Abdel-Razek, M.; Yogamani, S.; Jagersand, M.; and Zhang, H. 2018. A comparative study of real-time semantic segmentation for autonomous driving. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 587–597.
- Sun, X. 2024. *Assessing Model Robustness in Complex Visual Environments*. Ph.D. thesis, The Australian National University (Australia).
- Sun, X.; Leng, X.; Wang, Z.; Yang, Y.; Huang, Z.; and Zheng, L. 2023a. Cifar-10-warehouse: Broad and more realistic testbeds in model generalization analysis. *arXiv preprint arXiv:2310.04414*.
- Sun, X.; Yao, Y.; Wang, S.; Li, H.; and Zheng, L. 2023b. Alice Benchmarks: Connecting Real World Re-Identification with the Synthetic. *arXiv preprint arXiv:2310.04416*.
- Twinanda, A. P.; Shehata, S.; Mutter, D.; Marescaux, J.; De Mathelin, M.; and Padoy, N. 2016. EndoNet: a deep architecture for recognition tasks on laparoscopic videos. *IEEE transactions on medical imaging*, 36(1): 86–97.
- Wang, H.; Yang, G.; Zhang, S.; Qin, J.; Guo, Y.; Xu, B.; Jin, Y.; and Zhu, L. 2024a. Video-instrument synergistic network for referring video instrument segmentation in robotic surgery. *IEEE Transactions on Medical Imaging*.
- Wang, H.; Yang, G.; Zhang, S.; Qin, J.; Guo, Y.; Xu, B.; Jin, Y.; and Zhu, L. 2024b. Video-instrument synergistic network for referring video instrument segmentation in robotic surgery. *IEEE Transactions on Medical Imaging*.
- Wei, M.; Budd, C.; Garcia-Peraza-Herrera, L. C.; Dorent, R.; Shi, M.; and Vercauteren, T. 2025. SegMatch: semi-supervised surgical instrument segmentation. *Scientific Reports*, 15(1): 14042.
- Yan, C.; Wang, H.; Yan, S.; Jiang, X.; Hu, Y.; Kang, G.; Xie, W.; and Gavves, E. 2024. Visa: Reasoning video object segmentation via large language models. In *European Conference on Computer Vision*, 98–115. Springer.
- Yang, Z.; Wei, Y.; and Yang, Y. 2021. Associating objects with transformers for video object segmentation. *Advances in Neural Information Processing Systems*, 34: 2491–2502.
- Yang, Z.; and Yang, Y. 2022. Decoupling features in hierarchical propagation for video object segmentation. *Advances in Neural Information Processing Systems*, 35: 36324–36336.
- Yuan, K.; Chen, T.; Li, S.; Lavanchy, J. L.; Heiliger, C.; Özsoy, E.; Huang, Y.; Bai, L.; Navab, N.; Srivastav, V.; et al. 2025a. Recognizing Surgical Phases Anywhere: Few-Shot Test-time Adaptation and Task-graph Guided Refinement. *arXiv preprint arXiv:2506.20254*.
- Yuan, K.; Kattel, M.; Lavanchy, J. L.; Navab, N.; Srivastav, V.; and Padoy, N. 2024. Advancing surgical vqa with scene graph knowledge. *International journal of computer assisted radiology and surgery*, 19(7): 1409–1417.
- Yuan, K.; Srivastav, V.; Yu, T.; Lavanchy, J. L.; Marescaux, J.; Mascagni, P.; Navab, N.; and Padoy, N. 2025b. Learning multi-modal representations by watching hundreds of surgical video lectures. *Medical Image Analysis*, 103644.
- Zhou, Z.; Alabi, O.; Wei, M.; Vercauteren, T.; and Shi, M. 2023. Text promptable surgical instrument segmentation with vision-language models. *Advances in Neural Information Processing Systems*, 36: 28611–28623.