



HAL
open science

Learning Multi-modal Representations by Watching Hundreds of Surgical Video Lectures

Kun Yuan, Vinkle Srivastav, Tong Yu, Joel Lavanchy, Pietro Mascagni, Nassir
Navab, Nicolas Padoy

► **To cite this version:**

Kun Yuan, Vinkle Srivastav, Tong Yu, Joel Lavanchy, Pietro Mascagni, et al.. Learning Multi-modal Representations by Watching Hundreds of Surgical Video Lectures. 2023. hal-04214451

HAL Id: hal-04214451

<https://hal.science/hal-04214451>

Preprint submitted on 22 Sep 2023

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

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



Distributed under a Creative Commons Attribution 4.0 International License

Learning multi-modal representations by watching hundreds of surgical video lectures

Kun Yuan^{a,c}, Vinkle Srivastav^{a,b}, Tong Yu^a, Joël Lavanchy^b, Pietro Mascagni^b, Nassir Navab^c, Nicolas Padoy^{a,b}

^aICube, University of Strasbourg, CNRS, Strasbourg, France

^bIHU Strasbourg, Strasbourg, France

^cCAMP, Technische Universität München, Munich, Germany

Recent advancements in surgical computer vision applications have been driven by fully-supervised methods, primarily using only visual data. These methods rely on manually annotated surgical videos to predict a fixed set of object categories, limiting their generalizability to unseen surgical procedures and downstream tasks. In this work, we put forward the idea that the surgical video lectures available through open surgical e-learning platforms can provide effective supervisory signals for multi-modal representation learning without relying on manual annotations. We address the surgery-specific linguistic challenges present in surgical video lectures by employing multiple complementary automatic speech recognition systems to generate text transcriptions. We then present a novel method, *SurgVLP* - Surgical Vision Language Pre-training, for multi-modal representation learning. *SurgVLP* constructs a new contrastive learning objective to align video clip embeddings with the corresponding multiple text embeddings by bringing them together within a joint latent space. To effectively show the representation capability of the learned joint latent space, we introduce several vision-and-language tasks for surgery, such as text-based video retrieval, temporal activity grounding, and video captioning, as benchmarks for evaluation. We further demonstrate that without using any labeled ground truth, our approach can be employed for traditional vision-only surgical downstream tasks, such as surgical tool, phase, and triplet recognition. The code will be made available at <https://github.com/CAMMA-public/SurgVLP>

Keywords: Multi-modal representation learning; Surgical video lectures; Self-supervision; Vision-and-language.

1. Introduction

Recent progress in surgical computer vision has started to pave the way for a new wave of AI-assisted support systems for the operating room (OR) (Maier-Hein et al., 2017, 2022; Ward et al., 2021; Mascagni et al., 2022; Madani et al., 2020; Yuan et al., 2021). The field has seen substantial advancements, evolving from coarse-level surgical workflow recognition (Blum et al., 2008, 2010; Padoy et al., 2012; Twinanda et al., 2016; Dergachyova et al., 2016) to fine-grained surgical scene understanding through surgical action triplet (Nwoye et al., 2022), pixel-level scene segmentation (Allan et al., 2019; Alapatt et al., 2021), and surgical scene reconstruction (Wang et al., 2022; Pfeiffer et al., 2019; Rivoir et al., 2021). Nevertheless, the current advancements present two major limitations. First, these approaches have largely been focused on building task-specific fully-supervised deep-learning models that demand a significant effort from clinical experts to generate labeled ground truth. Second, the efficacy of these approaches has primarily been demonstrated on a limited number of mono-centric, procedure-specific surgical video datasets, which are not representative enough to encompass the complex intricacies of the overall surgical workflow (Eisenmann et al., 2022). Approaches that can extend to multiple downstream tasks using minimal labeled data, while leveraging large-scale multi-

procedural surgical videos, will be instrumental in scaling up the approaches for widespread adoption.

In the general computer vision domain, *multi-modal representation learning* (Radford et al., 2021; Miech et al., 2020) that combines visual and free-form natural text information is emerging as a viable alternative to circumvent the need to collect labeled training data for different downstream tasks (Radford et al., 2021; Jia et al., 2021). These approaches aim to learn a low-dimensional joint latent space by pre-training two parallel encoders - one for vision and one for text - on large-scale paired visual-textual inputs. The shared latent space of the two modalities enables zero-shot transfer learning, i.e., the ability of the pre-trained visual and text encoders to adapt to different downstream tasks without fine-tuning using task-specific labels. This breakthrough has led to promising results in a wide range of general computer vision applications, including zero-shot image classification (Radford et al., 2021), image captioning (Nukrai et al., 2022), semantic image retrieval (Sain et al., 2023), and text-to-shape generation (Sanghi et al., 2022).

Considering this impressive advancement in multi-modal representation learning, a natural question arises: can such high-level joint representations be learned for surgical computer vision? If possible, this could be a significant step forward in the progress of surgical data science (Maier-Hein et al., 2022). By obtaining such representations, it would not only enable us to perform existing surgical video analysis tasks, such as coarse-grained to fine-grained surgical workflow recognition (Twinanda et al., 2016; Nwoye et al., 2022),

*Corresponding author: Tel.: +33-390413530
e-mail: npadoy@unistra.fr (Nicolas Padoy)

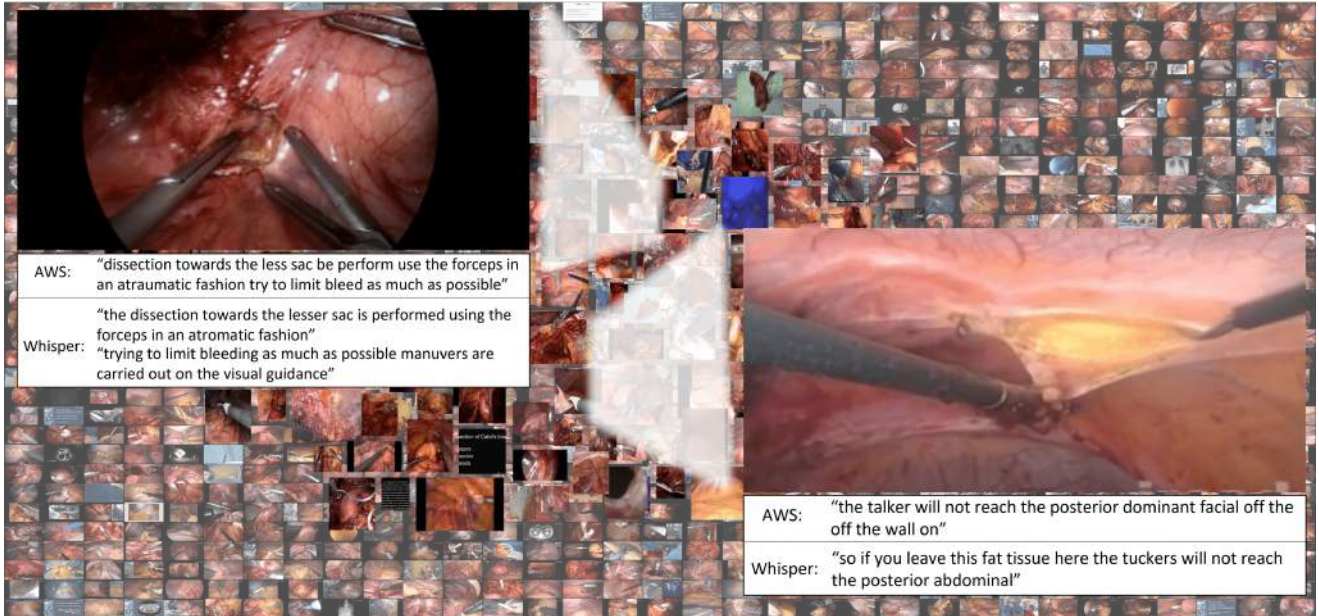


Fig. 1: Examples of video clip-text pairs from SVL dataset. The video clip-text pairs are pairs of video clips and their corresponding transcripts. We generate transcripts for hundreds of surgical video lectures using two ASR systems. The transcripts usually illustrate the essential concept of surgical anatomies, instruments and events. We use large-scale video clip-text pairs to learn joint multi-modal representations.

without using task-specific labels, but it would also open up new avenues for scalable and intelligent cognitive aids in the OR. These include vision-and-language applications, such as surgical visual question answering (Seenivasan et al., 2022), surgical report generation (Xu et al., 2021), and facilitating interactive communication between clinicians and surgical devices.

This work introduces *SurgVLP*, Surgical Vision Language Pre-training, a deep learning approach to perform large-scale multi-modal representation learning for surgical computer vision. Developing such an approach is not without its unique challenges. One of the primary obstacles is the unavailability of large-scale multi-modal multi-procedural surgical datasets compared to the millions of multi-modal visual-textual pairs available in the general computer vision domain (Radford et al., 2021; Grauman et al., 2022; Miech et al., 2019). For instance, a recently developed Ego4D (Grauman et al., 2022) dataset collects 3,000 hours of activity videos and manually narrates them. Such methods are unattainable in the surgical field due to the significant human efforts required in collecting and annotating surgical videos. As our first contribution, we propose to employ surgical video lectures available through open surgical e-learning platforms such as WebSurg (Web-surg, 2023) and EAES (EAES, 2023), and online video sharing platforms such as YouTube (YouTube, 2023), for visual-textual multi-modal learning. Compared to the manually labeled medical imaging reports (Chen et al., 2022a) or surgical instructions (Rojas-Muñoz et al., 2020), we propose to employ unprocessed and possibly noisy audio as the primary source of supervision for multi-modal representation learning. We leverage recent advances in audio speech recognition (ASR) (Mehri et al., 2023) to transcribe the lecture audio into sentences and link them to the corresponding video seg-

ment to construct large amounts of video clip-text pairs, as shown in Fig. 1. The resulting surgical video lecture (SVL) dataset contains diverse descriptions of surgical events, instrument usage, and anatomical status across various surgical procedures, thereby providing sufficient supervision to enable multi-modal representation learning for surgery.

Multi-modal representation learning using the SVL dataset nonetheless poses several linguistic challenges. First, the surgical concepts described in these videos use domain-specific knowledge and scientific terminology not typically encountered in general computer vision. For instance, “*grasping the neck of the gallbladder and retracting it towards the left lower quadrant to open up the hepatocystic triangle*” and “*dissect above an imaginary safety line connecting Rouviere’s sulcus and the base of the fourth liver segment*” are surgery-specific descriptions commonly found in the surgical video lectures for the laparoscopic cholecystectomy procedure. Furthermore, there could be a semantic misalignment between the surgical video clips and the corresponding textual descriptions. In fact, the lecturer describing the surgical procedure might divert from the case at hand and recall a similar case with a bleeding event, even if this is not shown in the associated video. Additionally, these videos have long-range dependencies. For example, the lecturer might comment on the importance of an adequate dissection to obtain tension-free anastomosis, even though that dissection step was shown at the procedure’s beginning or edited out. Finally, while recent ASR models (Chen et al., 2022b; Radford et al., 2023) can effectively transcribe day-to-day speech, their performance is suboptimal in surgical scenarios due to the surgery-specific linguistic challenges, as described before. For example, the state-of-the-art ASR Whisper model (Radford et al., 2023) can understand the sentence structure and common words but

struggles with surgical-specific terms (e.g., transcribing “jejunostomy” as “egenostomy”). Commercial medical-specific solutions, such as AWS (AWS, 2023), is considerably better at transcribing medical terminology but often fail to capture the overall structure and boundaries of the sentences, see Fig. 1.

We propose two key techniques for developing surgery-specific multi-modal representation learning. First, we employ text transcription from two noisy but complementary ASR systems, i.e., Whisper (Radford et al., 2023) and AWS (AWS, 2023), to obtain improved supervisory signals for the learning process, effectively mitigating the limitations and inaccuracies associated with each individual system. Second, we propose a new contrastive learning objective that utilizes these dual text transcriptions from the ASR systems with the corresponding video clip. The proposed contrastive learning objective aims to encourage the embedding vectors of the video clip and the corresponding dual textual transcriptions to be close in the joint latent space. By doing so, the learned multi-modal representation retains common semantics present in the noisy ASR transcripts, enabling a more effective fusion of visual and textual information.

To effectively showcase the representation capability of the learned joint latent space, we introduce various vision-and-language tasks for surgery to serve as multi-modal benchmarks for evaluation. These tasks include *text-based video retrieval*, *temporal activity grounding*, and *video captioning*. The text-based video retrieval task aims to associate a given text query to various video clips, while the temporal activity grounding task involves localizing a given text query to a specific video segment in the entire video. These two tasks examine how well the joint latent space captures the underlying relationship inherent in surgical visual information and its textual descriptions. The video captioning task aims to generate captions for a given surgical video clip. Since this is a generative task, it entails using a text decoder to produce coherent textual output. We propose an approach to construct a *text decoder* and append it to our pre-trained encoders, seamlessly repurposing our pre-trained model to function as a video captioner. The whole process requires only the textual data to train the text decoder model. We demonstrate a notable improvement over the baseline methods for all the vision-and-language tasks.

Next, we assess the robustness and adaptability of our approach when applied to unseen surgical datasets and tasks. Specifically, we examine its performance in traditional vision-only surgical tasks, including surgical tool, phase, and action triplet recognition (Twinanda et al., 2016; Nwoye et al., 2022). We evaluate our approach as a zero-shot transfer learning by processing the category labels (tool, phase, or action triplet) into textual form and classifying the video frames based on the similarity of visual and textual latent vectors. The results show that the general surgical concepts learned by our multi-modal joint representations from various surgical procedures can benefit a specific surgical procedure, such as laparoscopic cholecystectomy. To the best of our knowledge, this is the first work demonstrating the ability of a learned multi-modal representation to recognize surgical tools, phases, and action

triplets without using any ground truth labels. While the results fall short of fully-supervised baselines, they are encouraging and could pave the way for more advanced methodologies to improve the results further. Finally, we conduct extensive ablation studies to shed light on the different components of our approach and their impacts on the results. The contributions of our work can be succinctly summarized in the following four key aspects:

- We propose to harness the knowledge from surgical video lectures accessible via open surgical e-learning platforms for visual-textual multi-modal representation learning. To this end, we introduce a large-scale dataset of surgical video lectures (SVL) comprising 1.4k procedural videos.
- We propose to leverage text transcription from two complementary ASR systems, Whisper and AWS, to enhance the learning process by addressing limitations and inaccuracies inherent in each ASR system.
- We propose a novel contrastive learning objective that leverages dual text transcriptions from ASR systems and the corresponding video clip, aiming to encourage close proximity of embedding vectors in the joint latent space.
- We demonstrate the zero-shot transferability of our learned multi-modal representations by showing superior performance on various vision-and-language and vision-only surgical tasks.

2. Related Works

2.1. Representation learning from unlabeled videos

Traditional fully-supervised approaches learn representations from datasets of images or videos with manually annotated labels. However, manual annotation is laborious and expensive, especially when creating labeled datasets for new problems. Therefore, several works have studied representation learning by deriving supervision from unlabeled visual data, i.e., self-supervised learning (SSL) methods.

Self-supervised learning. Typically, SSL methods rely on heuristics-based pretext tasks by pre-training the model to solve handcrafted tasks having some degree of relevance to downstream tasks. These pretext tasks focus on the consistency of visual entities, including predicting the geometric transformations (Jing et al., 2018), the future representation (Vondrick et al., 2016) as well as colorizing videos (Vondrick et al., 2018). However, these methods heavily rely on the quality of the pretext tasks to show improvement in the downstream tasks. Recently, contrastive learning methods have emerged as an alternative. These methods regulate the distribution of feature vectors within the embedding space. This regulation process is accomplished by generating *positive* and *negative* pairs and then applying a discriminative loss function. This function brings the embeddings of *positive* pairs closer together while simultaneously distancing the embeddings of *negative* pairs. Specifically, MoCo (He et al., 2020) proposes to store embeddings of *negative* samples in a queue-based memory bank which is continuously updated during the

course of contrastive learning. MoCo v2 (Chen et al., 2020b) and SimCLR (Chen et al., 2020a) are subsequent works that proposed to improve representation learning by using more advanced data augmentations.

In the surgical computer vision field, SSL methods have also been studied to support different surgical downstream tasks ranging from estimating remaining surgery duration to surgical activity recognition (Ross et al., 2018; Yengera et al., 2018; Funke et al., 2018; Ramesh et al., 2023; Rivoir et al., 2019). Specifically, these approaches employ SSL methods to pre-train the model to obtain robust and effective representations and then fine-tune the pre-trained model on different surgical downstream tasks to achieve superior performance.

Our approach distinguishes itself from these vision-only SSL methods by performing representation learning using supervisory guidance derived from language. Language supervision is produced by applying automatic speech recognition (ASR) systems to surgical video lectures. We focus on learning the multi-modal representations exclusively from uncensored surgical lecture videos, presenting a more realistic and scalable learning scenario.

Language in representation learning. The potential of language in deep learning has been harnessed through advancements in natural language processing (NLP). Transformer-based models like Llama (Touvron et al., 2023) and T5 (Raffel et al., 2020) have set new standards in tasks such as translation and summarization. Moreover, transformer encoders like Bart (Lewis et al., 2019) and RoBERTa (Liu et al., 2019) have improved bidirectional language understanding, enhancing natural language task performance. In the medical field, models like SciBERT (Beltagy et al., 2019) and BioClinicalBERT (Huang et al., 2019), trained on large-scale biomedical data and electronic health records, respectively, have significantly improved biomedical text mining and healthcare prognostics.

Different from the conventional visual-only SSL methods as discussed before, recent representation learning methods have started to explore language as an alternative semantic supervision. Numerous works have proposed to use the paired image/video and text datasets to learn a joint latent space where the visual and textual data are close if they are semantically similar (Changpinyo et al., 2021; Krishna et al., 2017; Chen et al., 2015). However, these methods still rely on human annotators to describe the visual content. To avoid labeling images, several works (Mahajan et al., 2018) have leveraged image titles, descriptions, and hashtag metadata to provide language supervision. The recently proposed CLIP model utilizes a dataset with 400 million image-text pairs for multi-modal representation learning, achieving superior zero-shot performance on various downstream tasks (Radford et al., 2021). Also, De-CLIP leverages various supervisions, including self-supervision within each modality and multi-view supervision across modalities to efficiently optimize the model using large-scale image-text paired datasets (Li et al., 2021).

Audio in representation learning. While exciting works have been proposed using natural language supervision, the application of natural language supervision for video repre-

sentation learning has received comparatively less attention. A significant hurdle is the high cost of annotating textual descriptions for video content. Some recent works adopt audio to learn a self-supervised multisensory representation (Owens and Efros, 2018; Alayrac et al., 2020). These approaches train a neural network to predict whether video frames and audio are temporally aligned. However, these works operate on the event audio signals, such as instrument or environmental audio, rather than spoken audio. As a result, these approaches are limited to specific audio contexts, and they may not perform optimally in scenarios where spoken language is the primary focus.

Spoken audio has gained more research interest with the rapid development of audio speech recognition (ASR) methods (Seide et al., 2011; Amodei et al., 2016; Narayanan et al., 2018). These ASR methods achieve a much lower word error rate when trained on the 2,000 hours Switchboard dataset (Seide et al., 2011). By leveraging the semi-supervised strategy, authors in (Narayanan et al., 2018) have expanded the training dataset size to 162,000 hours of labeled audio, thereby further decreasing the word error. The recently introduced Whisper model has set the new state-of-the-art performance in the general ASR domain by training a transformer decoder using a large-scale dataset in a weakly supervised manner (Radford et al., 2023).

Powered by the emergence of ASR, HowTo100M (Miech et al., 2019, 2020) constructs a large-scale paired video and text dataset by transcribing audio into textual annotations, thereby providing large-scale language supervision for representation learning. YT-Temporal (Zellers et al., 2021) creates a dataset for learning multimodal knowledge from textual graphs derived from 6 million public YouTube videos. It applies multiple filtering and denoising strategies to obtain clean and better-aligned audio transcripts. By utilizing audio from a wide range of instructional videos, these datasets facilitate the creation of scalable and robust models capable of understanding natural language from diverse, real-world scenarios.

While Whisper has shown promising transcription results in the general ASR domain, it faces challenges in accurately transcribing domain-specific content, such as surgical video lectures. Whisper is primarily trained on extensive multilingual general domain datasets, which are less likely to include a substantial number of surgery-specific words. To address this issue, we propose to use two ASR systems to generate video clip-text pairs from surgical video lectures. Specifically, we use AWS Transcribe Medical (AWS, 2023) as a complementary ASR system to Whisper. AWS Transcribe Medical is a service provided by Amazon Web Services and is specifically designed to transcribe medical keywords accurately. Combining these two ASR systems addresses the limitations associated with the individual ASR system, thus helping to create effective transcriptions for surgical video lectures.

2.2. Downstream tasks

The quality of representation learning is typically assessed by applying the learned representations to various downstream tasks. In our context, we categorize these downstream tasks

into two main groups: vision-and-language and vision-only tasks. Vision-and-language tasks are at the intersection of vision and language modalities and can be classified into two categories: visual-textual understanding and generation tasks. Visual-textual understanding tasks aim to correlate the information across different modalities, such as text-based video retrieval. Visual-textual generation tasks aim to generate one modality output from another, such as video captioning. The objective of vision-only tasks is to understand and interpret the visual data from a single vision modality. A few examples of these tasks include image recognition, object detection, semantic segmentation, and pose estimation. In the following, we describe different vision-and-language and vision-only tasks.

Vision-and-language tasks. With the substantial progress in NLP (Devlin et al., 2018; Brown et al., 2020) and multi-modal self-supervised learning (Radford et al., 2021), research has been shifting towards multi-modal vision-and-language tasks. Text-based video retrieval (Gabeur et al., 2020) enables efficient video content searching and understanding using natural language queries. Temporal activity grounding (Gao et al., 2017) seeks to identify and localize specific action steps within a video sequence given a text description. In the context of surgical vision-and-language tasks, SurgVQA (Seeni-vasan et al., 2022) proposes a visual question-answering system designed to answer the questions related to surgical scenes, with the aim to offer real-time support to junior residents. Surgical report generation is also explored to describe current surgical events and objects (Xu et al., 2022). Similarly, DAISI (Rojas-Muñoz et al., 2020) has introduced a method for generating surgical instructions for different surgical procedures.

Vision-only tasks. Vision-only tasks have made substantial advancements using fully-supervised deep learning methods. These advancements have shown state-of-the-art performance in object detection (Girshick, 2015), action segmentation (Tang et al., 2019; Farha and Gall, 2019), image classification (Deng et al., 2009; He et al., 2016), and action recognition (Soomro et al., 2012). These successes are primarily attributed to the power of deep representations in capturing complex visual patterns.

The surgical computer vision community has embraced these advancements, exploring numerous studies on surgical tool and phase recognition (Twinanda et al., 2016; Dergachyova et al., 2016; Al Hajj et al., 2018; Shah et al., 2023). Following this, researchers proposed more refined spatiotemporal architectures to capture the dynamic and complex nature of surgical tasks (Jin et al., 2017; Czempiel et al., 2020, 2021; Jin et al., 2021; Rivoir et al., 2022). For a more comprehensive understanding, a recent survey provides an overview of surgical phase recognition approaches (Garrow et al., 2021). For a more fine-grained surgical activity recognition, authors in (Nwoye et al., 2022) have introduced the challenging action triplet recognition task for laparoscopic cholecystectomy procedures. They propose a dataset to classify instrument, verb, and target components from surgical scenes. However, all previous methods are vision-only, fully supervised methods that

rely heavily on human-annotated curated data.

In this work, we investigate both vision-and-language and vision-only surgical downstream tasks. We apply our approach to vision-and-language tasks, including text-based video retrieval, temporal activity grounding, and video captioning, to assess how well the model integrates and correlates information across modalities. We also apply the model to vision-only tasks, including surgical tool, phase, and action triplet recognition, to investigate whether our representation learning from uncurated data benefits a specific procedure, for example, laparoscopic cholecystectomy. We take a first step towards generic representation learning for surgery by performing zero-shot recognition on these surgical computer vision tasks. By evaluating our models on both vision-and-language and vision-only tasks, we aim to provide a comprehensive view of the capabilities and limitations of our approach.

3. Approach

This section outlines our *SurgVLP* approach for learning joint multi-modal representations from surgical video lectures and its adaptability to multiple downstream tasks. We first introduce the process of constructing a large-scale video clip-text pair dataset from surgical video lectures. Then, we present our dual-branch architecture consisting of one branch for the visual encoder and another for the text encoder to extract the latent vectors from the video clip and the corresponding multiple text views, respectively. We then describe our contrastive learning objective to optimize the dual-branch architecture using the joint latent vectors. Finally, we describe a variety of vision-and-language and vision-only surgical downstream tasks to evaluate learned multi-modal representations.

3.1. Video clip-text pairs

We utilize the recent advancements in audio speech recognition (ASR) to transcribe lecture audio into textual sentences. In particular, we use two different ASR systems to generate complementary text views, namely the AWS Medical Transcribe ASR system (AWS, 2023) and the Whisper ASR system (Radford et al., 2023), as we observe that ASR systems perform differently on the specific surgical domain. In particular, AWS Medical Transcribe is used to understand medical terminology and surgery-specific terms, and Whisper is used to understand the overall sentence structure and common words.

To obtain the first text view, we start with the AWS ASR system and use it to convert lecture audio into text transcriptions. We then apply text preprocessing to truncate the transcription into sentences based on the pre-defined stop symbols. Then, we apply filtering strategies to keep the sentences containing the surgical terminologies and remove meaningless and extremely short sentences. Thus, for a given surgical lecture video, these operations extract multiple AWS sentences, where each sentence is accompanied by its start and end timestamps. Due to the filtering, the AWS sentences contain surgery-specific terms but are sparsely distributed along

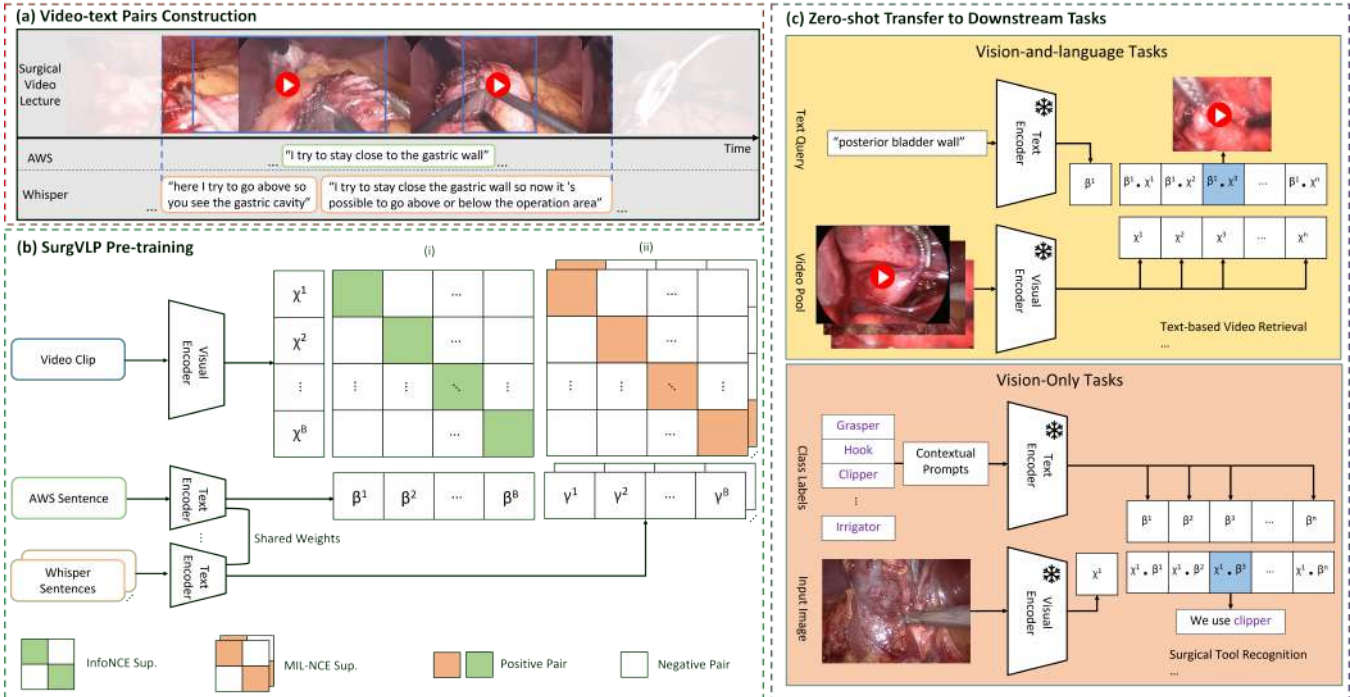


Fig. 2: (a) shows the examples of video clip-text pairs and their construction process. We have two text views and we pair them to random lengths of video clips. (b) presents the contrastive learning objective with AWS sentences and Whisper sentences. *SurgVLP* utilizes the Info-NCE and MIL-NCE losses for AWS and Whisper sentences, respectively. (c) illustrates how to perform downstream tasks in the zero-shot setting. We show the vision-and-language tasks, e.g., text-based video retrieval and temporal activity grounding, at the top and the vision-only tasks at the bottom.

the video time axis. More details on text pre-processing are given in section 4.2.

To obtain the second text view, we use the Whisper ASR system to generate the text transcription. We apply the same text preprocessing strategy as above to truncate the transcription into sentences. We do not apply heavy filtering because the Whisper ASR system generates meaningful sentences with clear stop symbols. Therefore, the processed Whisper sentences possess different punctuation patterns than AWS sentences and are densely distributed along the video time axis. This variation in punctuation and distribution leads to potential overlapping between AWS and Whisper sentences, as shown in Fig. 2(a). Thus, one AWS sentence usually corresponds to $M \geq 1$ overlapped Whisper sentences. To extract the corresponding video clip from given multiple text views, we propose the following steps:

- sample an AWS sentence to ensure that the sentence contains surgery-specific term;
- find M overlapped Whisper sentences;
- merge Whisper sentences and sample a timestamp, called *center timestamp*, within the boundary of merged Whisper sentences;
- grow a video clip of random duration (up to 10 seconds) from this *center timestamp* as the center for the video clip.

We give below our intuitions behind the choices made for the video clip sampling.

Better clip sampling with key surgical terms. Randomly sampling a video clip from a surgical lecture video increases the chances of having a video clip that either does not align with the relevant text or has no correspondence with any text at all. Therefore, sampling the AWS sentences to locate video clips ensures that the sampled video clip contains surgical terms of interest.

Better video clip boundary. The start and end timestamps of AWS sentences provide a rough range for video clip duration. However, sampling video clips based on the exact timestamps of AWS transcribed sentences could produce less relevant video clips. This is because AWS transcribed sentences miss explicit sentence boundaries, thereby producing weak alignment to the visual content. To address this, we propose to sample video clips within the boundary of merged Whisper sentences because they provide better and more explicit activity boundaries. We also sample video clips of different lengths to encourage fine-grained association. Thus, a video clip may have a better chance of being aligned or supervised by nearby text.

These steps produce a set of video clips, each associated with one AWS sentence and M Whisper sentences. This constitutes our *video clip-text pair* dataset, as illustrated in Fig. 1 and Fig. 2(a).

3.2. Dual-branch model

Given the *video clip-text pair* dataset, we denote a video clip as v and its corresponding AWS sentence and Whisper sentences as a and w^m , respectively. Here, m indexes over each Whisper sentence. Given all paired video clips and transcribed

sentences $\{v_i, a_i, w_i^m\}_{i=1}^K$, *SurgVLP* aims to learn a joint latent space that correlates semantically similar video clips with the corresponding multiple view texts. Here, K is the total number of video clip-text pairs in the dataset. Formally, we aim to learn two parameterized mappings: $\mathcal{F} : v \rightarrow \mathbb{R}^d$ maps a video clip into a d dimensional latent vector, and $\mathcal{G} : a/w^m \rightarrow \mathbb{R}^d$ maps an AWS or a Whisper sentence into a d dimensional latent vector. To learn these two mappings, we employ a dual-branch model with a vision branch \mathcal{F} using a CNN-based visual encoder and a text branch \mathcal{G} using a transformer-based text encoder, described below.

3.2.1. Visual encoder

We employ the ResNet-50 model (He et al., 2016) pre-trained on ImageNet as the base architecture for the visual encoder \mathcal{F} . We replace the last layer of ResNet-50 with a linear layer to project the incoming global-averaged pooled vector into the d dimensional vector. For a given video clip v , we first sample a fixed number of frames $\{z_0, \dots, z_T\}$ and then feed them through \mathcal{F} to obtain the latent vectors for each frame, i.e., $\mathcal{F}(z_i) \in \mathbb{R}^d$. Subsequently, we perform average pooling across frames to obtain the final latent vector $\chi \in \mathbb{R}^d$, as shown below:

$$\chi = \frac{1}{T} \sum_i^T \mathcal{F}(z_i). \quad (1)$$

3.2.2. Text encoder

As mentioned above, we use AWS and Whisper ASR systems to transcribe surgical audio into surgical texts. The surgical texts generated using these ASR systems are inherently noisy, missing either the overall sentence structure or domain-specific scientific terminologies. Thus, a robust and domain-specific text encoder is required to handle the complexity of the surgical texts. We propose to use a transformer-based text encoder to encode the transcribed surgical sentences into representative latent vectors. In particular, we employ BioClinicalBert (Huang et al., 2019) as our text encoder \mathcal{G} . It is a base-size Bert model (Devlin et al., 2018), containing 12 encoders with 12 bidirectional self-attention heads totaling 110 million parameters. It is pre-trained on medical texts and generates more representative features than other text encoders, such as SciBert (Beltagy et al., 2019) and Bert (Devlin et al., 2018), see section 4.4 for a comparison.

Given an AWS or Whisper sentence, we first split each word into n_i subwords based on the word-piece tokenizer. Therefore, we generate a total of $N = \sum_{i=1}^L n_i$ subwords for a sentence with L words. For the sentences of different lengths, we truncate or pad them to the fixed length of N subwords. The word-piece tokenizer addresses out-of-vocabulary words and typo errors by splitting them into subword units. It maintains linguistic meaning and minimizes the adverse effects of unknown surgical terminologies. We consider each subword as a token t_i and pass its token ID to the tokenizer’s lookup embedding table to get the corresponding input embedding vectors of size $\mathbb{R}^{N \times 768}$. Then, input embedding vectors are passed through multiple transformer layers of the text encoder \mathcal{G} , resulting in a feature vector of size $\mathbb{R}^{N \times 768}$. Then, \mathcal{G} performs

the global average pooling to aggregate N feature vectors into one vector. Finally, a linear layer (MLP) projects this vector into the d dimensional output latent vector as follows:

$$\beta/\gamma^m = \text{MLP}\left(\frac{1}{N} \sum_i^N \mathcal{G}(t_i)\right). \quad (2)$$

Here, β and γ^m represent a d dimensional latent vector for each sentence transcribed from AWS and Whisper ASR systems, respectively.

3.3. Multiple text-views contrastive supervision

In this section, we describe our multiple text-views contrastive learning objective. Given the latent vectors $\{\chi_i, \beta_i, \gamma_i^m\}_{i=1}^K$ from video clip-text pairs $\{v_i, a_i, w_i^m\}_{i=1}^K$, we first consider the case when we have text supervision from only one ASR, i.e., supervision from AWS text latent vectors (β_i). In this case, we employ the InfoNCE loss, which is a contrastive loss function utilized in self-supervised learning (Oord et al., 2018; He et al., 2020; Chen et al., 2020a) and multi-modal representation learning (Radford et al., 2021; Li et al., 2021); here NCE stands for Noise Contrastive Estimation. By utilizing the InfoNCE loss, denoted as $\mathcal{L}_{\text{InfoNCE}}$, we align the visual latent space with the corresponding AWS textual latent space, as follows:

$$\mathcal{L}_{\text{InfoNCE}} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(\text{sim}(\chi_i, \beta_i)/\tau)}{\sum_{j=1}^B \exp(\text{sim}(\chi_i, \beta_j)/\tau)}. \quad (3)$$

Here, the numerator $\text{sim}(\chi_i, \beta_i)$ is the cosine similarity score between visual and corresponding AWS text latent vectors, i.e., *positive* pairs. The denominator $\text{sim}(\chi_i, \beta_j)$ is the cosine similarity between the visual latent vector and all other AWS text latent vectors in the batch, i.e., *negative* pairs, as illustrated in Fig. 2(b). Also, B is the batch size, and τ is the temperature hyper-parameter to control the probability distribution over the positive and negative pairs in the embedding space (Wu et al., 2018). A high-temperature value smooths out the distribution, allowing the model to consider a large number of negative pairs. A low-temperature value makes the distribution sharper, focusing more on the positive pairs and reducing the influence of negative examples.

While AWS sentences contain medical and surgery-specific terms, they however generate incomplete sentence fractions, eventually providing insufficient supervision when utilizing the InfoNCE learning objective. To address this issue, as our key contribution, we propose to exploit text supervision from multiple complementary ASR sources. In addition to the text supervision from the AWS text latent vectors (β_i), we utilize text supervision from Whisper latent vectors (γ_i^m) that can recognize overall sentence structure, thus enhancing the alignment between text and video clips.

As mentioned in section 3.1, we link $M \geq 1$ Whisper sentences to a video clip. Therefore, we propose to extend the InfoNCE learning objective with the MIL-NCE (Miech et al., 2020) learning objective, where MIL stands for Multiple Instance Learning. The MIL-NCE learning objective, denoted

as $\mathcal{L}_{MIL-NCE}$, aims to align the visual latent vector χ with multiple Whisper latent vectors γ^m , as follows:

$$\mathcal{L}_{MIL-NCE} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\sum_{m=1}^M \exp(\text{sim}(\chi_i, \gamma_i^m)/\tau)}{\sum_{j=1}^B \sum_{m=1}^M \exp(\text{sim}(\chi_i, \gamma_j^m)/\tau)}. \quad (4)$$

Our final multiple text-views contrastive loss \mathcal{L} combines these two loss functions, $\mathcal{L}_{InfoNCE}$ and $\mathcal{L}_{MIL-NCE}$, scaled by the weighting coefficients ϵ , as shown below:

$$\mathcal{L} = \epsilon \mathcal{L}_{InfoNCE} + (1 - \epsilon) \mathcal{L}_{MIL-NCE}. \quad (5)$$

In summary, the $\mathcal{L}_{InfoNCE}$ loss is responsible for aligning surgery-specific terms from AWS sentences with video clips, and the $\mathcal{L}_{MIL-NCE}$ loss is responsible for addressing the misalignment issue that the lecturers might talk about something before or after they actually demonstrate it. Both loss functions aim to maximize the alignment between *positive* pairs of video clips and texts compared to random *negative* pairs in the batch B .

3.4. Zero-shot transfer to downstream tasks

SurgVLP learns generic multi-modal representations from different surgical procedures to correlate the semantically similar video clips and texts. Therefore, in order to assess its learned joint latent representations in handling the multi-modal data, we introduce various *vision-and-language* surgical tasks. These tasks include text-based video retrieval, temporal activity grounding, and video captioning to serve as comprehensive multi-modal benchmarks.

Furthermore, we investigate the zero-shot transferability of *SurgVLP* to the traditional *vision-only* surgical tasks. This evaluation is crucial to assess whether the learned joint latent representations can effectively generalize to an unseen surgical dataset without using task-specific labels. We evaluate *SurgVLP* on Cholec80 (Twinanda et al., 2016), CholecT45 (Nwoye et al., 2022) datasets for surgical phase, tool, and action triplet recognition. Below, we provide a detailed explanation of both vision-and-language and vision-only surgical tasks.

3.4.1. Vision-and-language surgical tasks

Text-based video retrieval. Given a textual query using natural language, text-based video retrieval aims to retrieve matched video clips from a large video repository. This task allows users to search for videos using free-form language queries, which is a flexible and intuitive way of querying. An ideal retrieving system needs to handle novel queries that were not seen during the training process. Also, text-based video retrieval implies the degree of understanding across different data modalities, i.e., video and text, which is a challenging and important problem in surgical computer vision. In this work, we perform this task in a zero-shot manner as evaluation methods relying on model fine-tuning or linear probing require task-specific labels and complex hyper-parameter tuning, which may not directly reflect the true performance of representation learning. We freeze the visual and text encoders after the pre-training. We use the visual encoder to

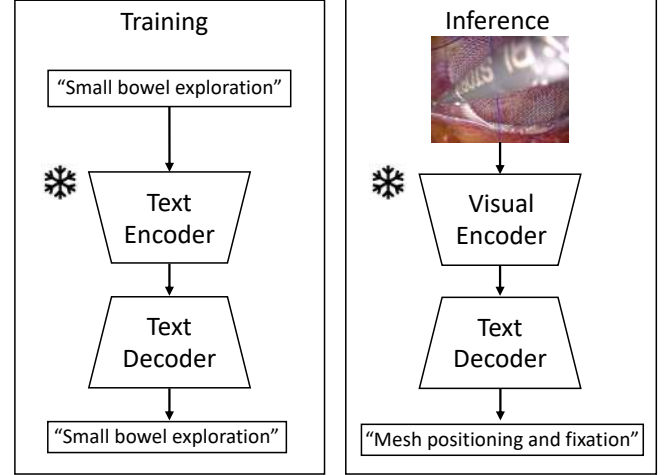


Fig. 3: Text-only-training for video captioning: We use the learned joint embedding space where text is encoded in a representation close to the ones of its corresponding video clips. During training, we train the text decoder to generate captions from text embeddings. During inference, the visual embeddings are fed to the visual encoder and then to the text decoder to generate the text captions.

generate the visual latent vectors of the video clips from the test database. Then, given a language query, we use the text encoder to generate a text latent vector that is matched against the visual latent vectors using cosine similarity. We retrieve top-K highest-matched video clips.

Unlike the pre-training phase, where video clips of random durations are generated, we follow a different process to generate video clips for evaluation. We manually correct Whisper transcriptions on a subset of the dataset with the help of a clinical expert. Subsequently, each individual sentence is aligned with its corresponding video clip, ensuring a precise correspondence between the textual description and the visual content. This process results in a video clip-text pair test dataset that enables a proper multi-modal evaluation of *SurgVLP*.

Temporal activity grounding. Similar to the text-based video retrieval task, temporal activity grounding is a retrieval-based task. However, rather than retrieving from a large pool of video clips, it aims to localize a given textual query to a relevant segment within a video. It therefore requires more fine-grained video text understanding and could be useful in many applications, such as identifying critical events in a given surgical video using textual queries and surgical video editing. We perform the temporal activity grounding in a zero-shot way by dividing a long lecture video into multiple video clips and retrieving the relevant video clip based on the cosine similarity with the given textual query. We utilize the same evaluation dataset as described in the previous section.

Text-only training for video captioning. The above retrieval-based tasks essentially aim to measure *SurgVLP*'s multi-modal representations by its ability to understand and correlate the latent vectors of two different modalities, i.e., video and text. However, it is unknown if the learned representations can be adapted to a textual generation task, such as image or video captioning.

The traditional computer vision approaches for image or

Table 1: Manually designed contextual prompts for the class names of the surgical phase and tool recognition tasks. The main action of scissors is cutting, but this action can be performed by many other instruments, such as hook. Therefore, we use “I use scissors” as the context prompt for the “Scissors” class.

Phase Labels	Prompts	Tool Labels	Prompts
<i>Preparation</i>	In preparation phase I insert trocars to patient abdomen cavity	<i>Grasper</i>	I use grasper or cautery forcep to grasp it
<i>CalotTriangleDissection</i>	In calot triangle dissection phase I use grasper to hold gallbladder and use hook to expose the hepatic triangle area and cystic duct and cystic artery	<i>Bipolar</i>	I use bipolar to coagulate and clean the bleeding
<i>ClippingCutting</i>	In clip and cut phase I use clipper to clip the cystic duct and artery then use scissor to cut them	<i>Hook</i>	I use hook to dissect it
<i>GallbladderDissection</i>	In dissection phase I use the hook to dissect the connective tissue between gallbladder and liver	<i>Scissors</i>	I use scissor
<i>GallbladderPacking</i>	In packaging phase I put the gallbladder into the specimen bag	<i>SpecimenBag</i>	I use specimenbag to wrap it
<i>CleaningCoagulation</i>	In clean and coagulation phase I use suction and irrigation to clear the surgical field and coagulate bleeding vessels	<i>Irrigator</i>	I use irrigator to suck it
<i>GallbladderRetraction</i>	In retraction phase I grasp the specimen bag and remove it from trocar	<i>Clipper</i>	I use clipper to clip it

video captioning require paired visual and textual input to train the model (Chen et al., 2020c). In this work, we however propose to exploit the multi-modal knowledge of *SurgVLP* for the task of video captioning by performing a text-only training process. This approach requires the minimal modifications to our pre-trained model and does not require any visual input for training.

Our approach solely relies on textual data for training, achieved through the construction of a trainable *text decoder* (Nukrai et al., 2022). During the training phase, this text decoder is appended to the frozen pre-trained text encoder and is trained as an encoder-decoder architecture to reconstruct the textual sentences from the textual latent vectors, which are generated from the *SurgVLP*’s pre-trained *text encoder*. During the inference phase, the visual latent vector from a given video clip is generated using the frozen pre-trained visual encoder and fed into the trained text decoder to generate the corresponding text captions. Thanks to the multi-modal knowledge learned through *SurgVLP*, it allows us to effectively pass the latent vectors of video clips to the text decoder. The training and inference phase is illustrated in Fig. 3.

To generate meaningful captions for video clips, instead of using inherently noisy transcribed textual sentences from pre-training, we opt for using *metadata sentences* from the surgical video lectures. These metadata sentences are uploaded by the content creator and typically describe different phases of the surgical procedure in a given video. These sentences are shorter and more abstract such as “*creation of the gastric pouch*” or “*trocar and port placement*” instead of a transcribed sentence such as “*I think the important thing is that we have a good view at the beginning I perform the dissection with electrocautery*”.

These metadata sentences show a significantly different textual distribution than the pre-training sentences. Therefore,

using this task, we aim to assess the effective alignability of our learned multi-modal representations into a different textual domain. We use the metadata sentences from the videos in the pre-training dataset to train the *text decoder*; for the evaluation, we use a separate set of 100 videos containing video clips and corresponding metadata sentences.

3.4.2. Vision-only surgical tasks

We evaluate our approach on traditional vision-only surgical tasks using the Cholec80 (Twinanda et al., 2016) and CholecT45 datasets (Nwoye et al., 2022). Through these evaluations, we aim to assess the domain generalization and robustness to distribution shift of *SurgVLP*, as these datasets contain only laparoscopic cholecystectomy videos. In the following, we describe these tasks and outline our approach to solving them without relying on any ground-truth labels, i.e., in a zero-shot manner.

Surgical tool, phase and triplet recognition. We evaluate three recognition downstream tasks of different granularities: surgical tool and phase recognition as coarse-grained tasks and surgical action triplet recognition as fine-grained task.

SurgVLP is pre-trained to predict if a video clip and a text snippet are paired together in the joint latent space. We reuse this capability by processing the name of all classes as potential text snippets and predicting whether a video clip and class name are paired. Instead of feeding class names directly to *SurgVLP*, we manually create contextual prompts for the names of all classes in Cholec80 and CholecT45. Given the contextual prompts, we compute their text latent vectors using the frozen text encoder. We feed the input testing image to the visual encoder to compute the visual latent vector and compare it to the text latent vectors using cosine similarity. We assign the most probable class to the image using the one with the highest cosine similarity score, as shown in Fig. 2(c).

The zero-shot evaluation on these datasets presents numerous challenges. First, manually defining the class names for a coarse-grained surgical phase introduces ambiguity. For example, the Cholec80 dataset manually defines the class name “*clipping and cutting*”, which does not accurately describe the surgical definition. In the clipping and cutting phase, the surgeon uses the clipper to clip the cystic duct and artery, then use scissors to cut them. Therefore, directly feeding class names into our pre-trained model increases confusion and degrades performance. Second, there exists a significant textual gap between the downstream and the pre-training datasets. Specifically, annotators usually define class as a single word in the downstream dataset. However, pre-training text transcriptions are usually complete sentences. A video clip rarely accompanies a single word in our pre-training dataset.

To address the above issues, we perform prompt engineering. We create contextual prompts from category names. For surgical phase recognition, we create contextual prompts for each phase label using a sentence describing the current phase, which mainly includes instruments, anatomies, and actions, as shown in Table 1. For surgical tool recognition, we create a phrase for each label. The phrase describes the main action performed by the instrument. For example, as shown in Table 1, the category label “*hook*” in tool recognition is transformed to “*I use hook to dissect it*”, the category label “*grasper*” is transformed to “*I use grasper or cautery forceps to grasp it*”, and so on. For surgical action triplet recognition, we use a prompt template “*I use {tool} to {action} {target}*”, where we replace the value in the bracket with corresponding labels. For example, $\langle \textit{grasper}, \textit{grasp}, \textit{gut} \rangle$ is transformed into “*I use grasper to grasp the gut*”. These generated prompts help bridge the textual gap between the pre-training dataset and the downstream dataset.

4. Experiments

We first describe the implementation details of our method in section 4.1 and the pre-training dataset to train our model in section 4.2. We present the downstream tasks and datasets used in our zero-shot evaluation in section 4.3. We then show an ablation study emphasizing key ingredients of our approach in section 4.4. We compare our *SurgVLP* approach to previous self-supervised methods in section 4.5. Finally, we show the effect of contextual prompts and different text encoders in section 4.5.3.

4.1. Implementation details

Feature extraction. We use the ResNet-50 network pre-trained on ImageNet as the backbone of the visual encoder. We uniformly sample $T = 4$ frames for each video clip. We first extract the embedding vector for each frame and then aggregate these vectors by temporal average pooling. Finally, we obtain a $d = 768$ dimensional vector for each video clip. For each sentence transcribed from ASR systems, we tokenize it and then pad or truncate it to obtain the length $N = 77$. We feed the tokenized sentence into the text encoder and generate a $d = 768$ dimensional vector as the text embedding.

Hyper-parameters. We implement our method using the Pytorch framework. We set the batch size as $B = 80$ and the learning rate as 0.0001. The temperature parameter τ is set to 0.3. We use 1 sentence from AWS and $M = 2$ sentences from Whisper as corresponding texts for each video clip. The coefficient weight for two loss functions ϵ is set to 0.5. The pre-training takes approximately 3 days on four A100 GPUs.

4.2. Pre-training dataset

We propose the **Surgical Video Lecture Pre-training dataset (SVL-Pretrain)** for multi-modal representation learning. The SVL-Pretrain dataset contains large amounts of laparoscopic surgical video lectures. Specifically, we query the videos based on keywords from three online platforms, i.e., Websurg (Websurg, 2023), EAES (EAES, 2023), and YouTube (YouTube, 2023). For Websurg, we use “intervention” and “laparoscopic” as keywords to crawl 1,124 surgical lectures. They cover the five main categories, “Hepatobiliary and pancreatic surgery”, “General and digestive surgery”, “Bariatric surgery”, “Hernia surgery” and “Colorectal, transanal and proctological surgery”. For EAES and YouTube, we manually collect 202 surgical video lectures using manually designed keywords. Specifically, we compile a list of surgical keywords for specific surgical procedures, such as “laparoscopic cholecystectomy”, “hernia repairing” and “gastric bypass”. Then we form text queries by adding prefixes and suffixes to the keywords, such as “how to...” and “:101”. We use these text queries to retrieve surgical video lectures from EAES and YouTube. We keep the top 50 retrieved results and exclude videos without any audio. We also remove the video lectures that only contain slides and textbook images. In total, we collect 1,326 lectures for the pre-training set.

Text pre-processing. *SurgVLP* requires video clip-text pairs for the training. Therefore, we use the AWS Medical Transcribe (AWS, 2023) system and the Whisper (Radford et al., 2023) model to transcribe audio into textual sentences from the surgical video lectures. We apply pre-processing strategies to remove noisy and unaligned texts. Specifically, we first truncate the sentences based on the stop symbols, such as “;”, “;”, and so on. We then filter out sentences that do not contain any noun or verb. Since each word transcribed from AWS Medical Transcribe has a confidence score ranging from 0 to 1, we apply the confidence-based filter to remove sentences with a low average confidence score. We set up the threshold to 0.4. Finally, we apply a keyword filtering strategy to remove sentences that do not contain any useful surgical word for the downstream tasks. For the sentences transcribed from Whisper, we follow the above strategy except for the confidence-based and keyword-based filtering. This difference makes the AWS sentences distribute more sparsely than Whisper sentences, providing the rationale for our hyper-parameter setup. For example, we select one AWS sentence and $M \geq 1$ Whisper sentences as correspondences. After pre-processing, each surgical video lecture has multiple transcribed sentences along the time axis. We then generate video clip-text pairs as mentioned in Sec. 3.1. This results in 25,578 video clip-text pairs for the pre-training set.

4.3. Downstream datasets and tasks

To show the generalizability of our learned representations, we perform the evaluation on five diverse downstream tasks using four datasets described below.

4.3.1. Vision-and-language datasets and tasks

Text-based video retrieval. We construct a dataset, SVL-Retrieval, to evaluate *SurgVLP*'s performance on text-based video retrieval tasks. Specifically, SVL-Retrieval contains 20 surgical video lectures. Each video lecture is split into multiple video clip-text pairs, resulting in 537 pairs. The text queries in SVL-Retrieval are initially transcribed from the Whisper ASR system. Therefore, the transcribed sentences are inherently noisy and miss surgery-specific terms. To make a representative evaluation, our clinical collaborator has validated and cleaned the text queries by correcting the typos and punctuation errors.

We follow the same evaluation protocol as described in (Zhukov et al., 2019) to report the retrieval performance using the Recall@K metric (with $K = 1, 5, 10$), which measures the percentage of correct video clips present within the K top-ranked clips that were retrieved (the higher, the better). We also report the median rank of the video clips in the video pool. During the retrieval, each video clip in the SVL-Retrieval dataset receives a ranking score that indicates how close it is to the corresponding text query, compared to the other video clips in the SVL-Retrieval dataset (the lower, the better). The median rank score is the median of the ranks of the ground truth videos across all text queries. A lower median rank score indicates a stronger overall correlation between the video clips and corresponding text queries, thus demonstrating the effectiveness of the retrieval system.

Temporal activity grounding. We use the same SVL-Retrieval dataset to evaluate the performance of temporal activity grounding. We use the Recall@K metric (with $K = 1, 5, 10$) to report the performance. Here, we do not apply the median rank in the temporal activity grounding task. This is because temporal grounding is a video-specific task and does not yield a large enough ranking list to represent the performance of the method accurately.

Text-only training for video captioning. We create a video captioning dataset, called SVL-Caption, to perform text-only training for video captioning. The SVL-Caption dataset is comprised of a training set with 9,074 *metadata sentences* obtained from the videos in the SVL-Pretrain dataset. The testing set of SVL-Caption contains 734 video clip-caption pairs from a separate set of 100 videos obtained using the same collection process as described in section 4.2 but not used during the pre-training. As explained in section 3.4.1, the metadata sentences, which are uploaded by the video creator, describe different phases in the surgical video lecture.

We utilize BLEU (Papineni et al., 2002), METEOR (Banerjee and Lavie, 2005), and ROUGE (Lin, 2004) metrics to evaluate the quality of generated captions. The BLEU metric measures n-gram overlap, with higher scores for more matches. The METEOR metric enhances this by considering synonymy,

stemming, and exact word matches. Finally, the ROUGE metric evaluates n-gram overlap, longest common subsequence, and skip-bigrams. All three metrics effectively assess the capability of the video captioning model to deliver concise, informative summaries encapsulating key information.

4.3.2. Vision-only tasks

Tool recognition. We use the publicly available Cholec80 dataset (Twinanda et al., 2016) to evaluate the surgical tool recognition. We use the test split of this dataset. We apply the average precision (AP) to measure the recognition performance for surgical tools as the area under the precision-recall curve. We report the tool-wise AP and the average AP across classes.

Phase recognition. We use the same Cholec80 dataset test split as above to evaluate the surgical phase recognition. As precision and recall are commonly used in prior works (Twinanda et al., 2016; Czempiel et al., 2020), we choose F1-Score as the balanced metric for precision and recall for the phase recognition task. It ranges from 0 to 1, where 1 indicates perfect precision and recall.

Triplet recognition. We use CholecT45 dataset (Nwoye et al., 2022) to evaluate the surgical action triplet recognition. We use the test split of this dataset to make the comparison to the prior fully-supervised and self-supervised works. In the triplet recognition, we utilize the average precision (AP) metric to recognize individual components, i.e., instrument (AP_i), verb (AP_v), and target (AP_t). We also evaluate the performance of instrument-tissue interactions by looking at the metrics for different sets of triplet components: the instrument-verb (AP_{iv}), instrument-target (AP_{it}), and instrument-verb-target (AP_{ivt}).

4.4. Ablation studies

We perform our ablation studies on the following downstream tasks: phase recognition (Mean AP), triplet recognition (Mean AP), and text-based video retrieval (Recall@10). This subset of downstream tasks has been chosen for their simplicity of evaluation and because they cover a wide range of tasks. As the hyper-parameter decisions, such as number of text views, are made based on the results of the ablation studies, we conduct the evaluation on the validation subsets for all these three tasks. In the following, we give the key takeaways from our study.

Multiple text views introduce complementary knowledge. Here, we remove one of the textual inputs from the video clip-text pairs $\{v_i, a_i, w_i^m\}_{i=1}^K$. When we train the model with sentences transcribed only from the AWS system, we obtain video clip-text pairs $\{v_i, a_i\}_{i=1}^K$. When we train the model with sentences transcribed only from the Whisper system, we obtain video clip-text pairs $\{v_i, w_i^m\}_{i=1}^K$, as shown in Table 2.

As mentioned in section 3.1, we hypothesize that combining AWS Medical Transcribe and Whisper ASR systems will introduce complementary knowledge. In Table 2, we prove these hypotheses by showing that the model trained from *SurgVLP* outperforms its counterparts in phase recognition and text-based video retrieval tasks where the language

Table 2: Ablation studies. We conduct three sets of experiments to demonstrate the effect of key designs in our approach, multiple text views, clips of random lengths, and frame sampling from video clip. $\{v_i, a_i\}_{i=1}^K$: model trained with one AWS text view; $\{v_i, w_i^m\}_{i=1}^K$: model trained with one Whisper text view; $\{v_i, a_i, w_i^m\}_{i=1}^K$: model trained with both text views.

		Phase Recognition Mean AP	Triplet Recognition Mean AP	Text-based Video Retrieval Recall@10
Text Views	$\{v_i, a_i\}_{i=1}^K$	25.5	11.0	21.6
	$\{v_i, w_i^m\}_{i=1}^K$	23.5	9.1	14.4
	$\{v_i, a_i, w_i^m\}_{i=1}^K$	26.7	10.3	25.1
Clip Length	2 seconds	20.0	9.4	19.1
	4 seconds	24.1	9.2	21.6
	10 seconds	17.3	9.0	18.8
	Random	26.7	10.3	25.1
Frame Sampling	1 frame	19.7	9.5	18.1
	2 frames	19.8	9.7	20.6
	4 frames	26.7	10.3	25.1

Table 3: Comparison of different methods in text-based video retrieval and temporal activity grounding tasks.

Method	Text-based Video Retrieval				Temporal Grounding		
	R@1 (%)	R@5 (%)	R@10 (%)	Median Rank	R@1 (%)	R@5 (%)	R@10 (%)
Random	0.0	0.2	0.8	322	4.2	8.7	13.0
CLIP (Radford et al., 2021)	0.0	0.9	2.1	256	4.2	7.8	13.9
CLIP*	0.6	1.2	2.3	244	4.3	9.3	15.3
SurgVLP	2.8	11.8	16.1	70	8.6	19.9	29.7

and boundary information play the most important role. The improvement can be attributed to the fact that sentences transcribed by two ASR systems serve as two augmented versions of the “ground truth text”, helping the model to learn transform-invariant textual features. We also find that a model trained with sentences transcribed only by AWS leads to better performance in action triplet recognition. We believe this downstream task does not require sophisticated language understanding. Instead, detecting and matching the main keywords is sufficient to solve this task.

Different clip lengths affect different granular-level tasks. Here, we use both AWS and Whisper text views to pre-train the model and demonstrate the effect of the different lengths for video clip sampling. We conclude that different lengths of video clips affect the performance of different downstream tasks. Also, randomly growing video clips into different lengths provide better robustness to different granular-level downstream tasks.

Among all three tasks, action triplet recognition is a fine-grained task. Phase recognition and text-based video retrieval are coarse-grained. As shown in Table 2, longer clip lengths improve performance on coarse-grained tasks. For example, the phase recognition and text-based video retrieval tasks are significantly improved when the clip length is increased from 2s to 4s. However, extremely long clip lengths bring no benefit and hinder performance. This is because longer clips provide more redundant frames that are not semantically correlated with the texts, resulting in weakly aligned video clip-text pairs. To cover the tasks with different granularities, we sam-

ple the video clips with different lengths (up to 10s), which gives the best results for all three tasks. We fix the video clip sampling strategy for the rest of the paper by growing the video clips with random lengths.

More frames, the better. We continue to use two text views and the video clips of random lengths to evaluate the quality of our multi-modal representations trained with different numbers of frames. As shown in Table 2, performance is improved when the model includes more frames from the clip. Specifically, the performance of triplet recognition increases from 9.5 to 10.3 when three additional frames are used in the pre-training. This suggests that including more visual information from video clips could further improve the performance of the joint representation. Given the limited computational resources, we sample up to 4 frames for each video clip in the current work.

4.5. Main results and comparisons

We build three baselines for comparisons with our *SurgVLP* approach from different perspectives. One trivial system is a randomly initialized *SurgVLP*, denoted as **Random**, to present it as a lower baseline. We initialize it randomly 3 times and present the average results. We also take the publicly available **CLIP** (Radford et al., 2021) model as a competitive baseline because it is trained on 400 million of image-text pairs and shows promising zero-shot image classification results on the ImageNet benchmark. We further load the CLIP model’s weight and fine-tune it with our SVL-Pretrain dataset. We denote the resulting model as **CLIP***.



Fig. 4: Qualitative results of text-based video retrieval on SVL-Retrieval dataset using *SurgVLP*'s learned joint multi-modal representations. For each language query, we retrieve 3 video clips from the repository. The ground truth video clip is framed in green. It is here always mentioned in the top-3 results.




Fig. 5: Textual-visual activation maps from different sentence queries. The first row shows the ground truth. The second row shows the predicted activation map along the time axis for the raw sentence. The third row shows the newly generated activation maps conditioned by modified sentences. When the whole sentence is decomposed into sub-sentences, the *SurgVLP* approach generates a focused textual-visual activation map for the sentence with clear and less ambiguous words, such as “bloodless” compared to “helpful” on the left. This shows that *SurgVLP* responds to specific surgical terms rather than general terminology.

4.5.1. Vision-and-language tasks

This section analyzes *SurgVLP*'s zero-shot transfer to multiple vision-and-language downstream tasks.

Text-based video retrieval. Quantitatively, as shown in Table 3, *SurgVLP* surpasses the baselines by a large margin. Compared to the publicly available CLIP (Radford et al., 2021) model, multi-modal representations from our *SurgVLP* approach correlate the surgical narrative description to video

clips. It also links the instrument, action, and activity descriptions of different granularities to the corresponding video clips. Qualitatively, Fig. 4 illustrates the top three retrieved video clips from our proposed SVL-Retrieval dataset. The ground truth is framed in green. It shows that *SurgVLP* can capture specific medical terms like bag and liver retractor.

Temporal activity grounding. We then evaluate our learned multi-modal representations on the temporal activity

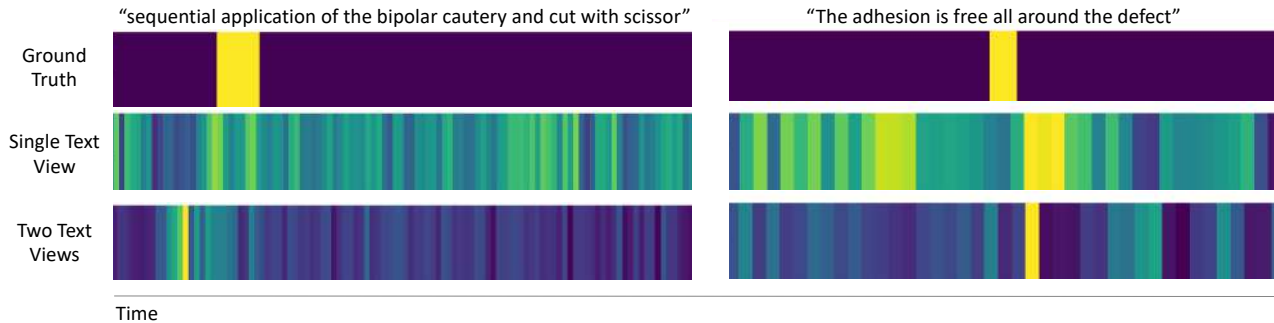


Fig. 6: Textual-visual activation maps of our *SurgVLP*, computed on two language queries from SVL-Retrieval testing set. The language queries are shown at the top of the figure, and the first row shows the ground truth activation map. The second row shows *SurgVLP*’s trained with one text view, i.e., AWS texts. The last row shows that *SurgVLP*’s generates more concrete activation maps with less noise when trained on both AWS and Whisper texts.



Fig. 7: Qualitative results of temporal activity grounding. We show the grounding results of two videos with three language queries. Each set of images represents a video clip. We show top-2 grounded clips for given text queries. Video clips framed in green are the ground truth to the given text. #1: top-1 grounded result. #2: top-2 grounded result.

Table 4: Quantitative results of text-only training for video captioning. We report 6 conventional metrics to measure the similarity between generated text and ground text. Our proposed *SurgVLP* significantly outperforms previous work, especially for ROUGE, which requires an accurate representation of not only individual words but also their correct order.

Methods	BLEU-1	BLEU-2	BLEU-3	BLEU-4	METEOR	ROUGE
Random	0.19	0.02	0.01	0	0.04	0.22
CLIP (Radford et al., 2021)	0.21	0.04	0.02	0.01	0.04	0.24
<i>SurgVLP</i>	0.36	0.22	0.17	0.12	0.18	0.33

grounding task in Table 3. It shows that *SurgVLP* outperforms the public CLIP model despite its use of 400 million images for pre-training. We believe that this improved localization ability is due to the fact that our model is trained on surgical narrative descriptions with temporally grounded video clips as opposed to image-level annotations in CLIP (Radford et al., 2021).

Fig. 5 shows the impact of novel text queries on the temporal activity grounding task. Specifically, we simulate the open-vocabulary scenario by decomposing the sentence into sub-sentence. For example, we remove the sub-sentence “and bloodless field could be obtained” from the complete sentence

“bipolar is very helpful in this task and bloodless field could be obtained”, resulting in a novel text query “bipolar is very helpful in this task”. We plot the visual-textual activation map by comparing the text embedding to all visual frame embeddings of the video. It shows that the model is triggered by precise sentences, such as “bipolar is very helpful in this task” rather than abstract or noisy ones, such as “bloodless field could be obtained”. Therefore, removing a noisy sub-sentence creates a more precise activation map. Fig. 6 also demonstrates that *SurgVLP* with two text views generates a focused attention map, highlighting relevant frames and ignoring adjacent ones. This demonstrates that multiple text view







			
Random	"clinical case presentation and stent catheterization"	"trocar placement and exposure of surgical field after lymph"	"specimen extraction and anastomosis preparation"
CLIP	"trocar placement and exploration of abdominal cavity"	"rectal division and anastomosis"	"rectal division and anastomosis"
SurgVLP	"trocar position and exposure of surgical field"	"cholecystectomy and common bile duct"	"crural repair (with pledgets) and vic"
Ground Truth	"trocar and port placement"	"dissection of the anatomical structure of the hepatic pedicle"	"crural repair"
			
Random	"trocar placement and exposure of surgical field after lymph"	"trocar placement and exposure of surgical field after lymph"	"clinical case presentation and stent catheterization"
CLIP	"intracorporeal anastomosis and mes"	"specimen extraction via a Pfannenstiel"	"trocar placement and exploration of abdominal cavity and operative"
SurgVLP	"gastric pouch creation: 2nd firing with end"	"jejunojunostomy with endo"	"end of procedure with extraction of specimen"
Ground Truth	"creation of the gastric pouch"	"intracorporeal esophagojejunostomy"	"mini-laparotomy for extraction of specimen"

Fig. 8: Caption results from text-only training for video captioning. Random: randomly initialized *SurgVLP*. CLIP (Radford et al., 2021): publicly available joint embedding space from OpenAI pre-trained CLIP model. *SurgVLP* shows more reliable captioning results with more overlap to the ground truth sentence. Also, the *SurgVLP* approach can generate detailed captions with the surgical instrument mentioned, e.g. "pledgets" in the top row last column.

supervision effectively connects the visual and the textual content.

We also show the qualitative results for the temporal activity grounding task. As shown in the first column of the top row of Fig. 7, the system clearly understands the concept of "lesion" to locate the correct video clip. As shown in the second column of the top row, it grounds the "seal" action correctly but fails to ground the "coagulate" action.

Text-only training for video captioning. We conduct text-only training for video captioning, demonstrating that the model pre-trained from the *SurgVLP* approach can be applied for text generation tasks. We compare *SurgVLP* to two baselines, the randomly initialized *SurgVLP* and CLIP. Here, all methods are trained and evaluated using the SVL-Caption dataset. As shown in Table 4, *SurgVLP* performs better than the other baselines, as it is trained with surgical video clip-text pairs. Random initialization yields the lowest results across all metrics because the visual and textual embeddings are not correlated. When a semantically similar video clip is given to this model, the model does not align it to the corresponding textual embedding, thus generating incorrect text output. Our proposed *SurgVLP* approach scores highest in all the metrics, suggesting the effect of its learned multi-modal embedding

space. We also show qualitative results in Fig. 8 to demonstrate that *SurgVLP* can accurately generate abstract metadata sentences from a given surgical video clip.

4.5.2. Vision-only tasks

This section analyzes *SurgVLP*'s zero-shot transfer to different vision-only surgical downstream tasks such as surgical tool, phase, and triplet recognition on laparoscopic cholecystectomy.

Tool recognition. As depicted in Table 5, *SurgVLP* exhibits promising tool recognition performance, especially in identifying tools with unique shapes like the T7 (specimen bag) with a 64.9% average precision (AP). This zero-shot performance is close to a fully supervised ResNet-50 model, as shown in the first row. Nevertheless, there remains ample opportunity for improvement in recognizing more challenging tools, such as scissors and irrigators. These tools are rarely mentioned in the SVL-Pretrain dataset. Also, we find that CLIP* gives a slight improvement but still lags behind *SurgVLP*. This demonstrates that multi-modal pre-training on natural images and texts hampers surgical tool recognition tasks due to the large domain gap between natural and surgical objects.

Phase recognition. As shown in Table 6, *SurgVLP* out-

Table 5: Zero-shot tool recognition on Cholec80. T1: grasper; T2: bipolar; T3: hook; T4: scissor; T5: clipper; T6: irrigator; T7: specimen bag. Fully-supervised: ResNet50 model with full supervision.

Methods	T1	T2	T3	T4	T5	T6	T7	Mean
Fully-supervised	86.9	79.6	99.2	81.2	94.9	61.3	93.8	85.3
Random	48.3	6.1	57.7	2.4	5.0	4.6	11.5	19.3
CLIP (Radford et al., 2021)	60.7	4.3	56.0	1.9	2.5	4.6	2.8	22.6
CLIP*	61.7	5.1	57.2	1.6	2.7	3.5	7.5	23.2
SurgVLP	55.2	21.3	61.4	7.8	16.2	7.9	64.9	33.3

Table 6: Zero-shot phase recognition on Cholec80. P1: preparation; P2: calot triangle dissection; P3: clipping and cutting; P4: gallbladder dissection; P5: gallbladder packing; P6: cleaning and coagulation; P7: gallbladder extraction. F1-Score is used as the evaluation metric. Fully-supervised: ResNet50 model with full supervision.

Methods	P1	P2	P3	P4	P5	P6	P7	Mean
Fully-supervised	62.1	84.5	74.4	82.2	62.7	52.2	52.2	67.3
Random	5.6	0.3	4.7	3.4	6.5	1.0	4.4	1.8
CLIP (Radford et al., 2021)	10.3	38.0	0	20.1	0.3	3.7	0.4	10.4
CLIP*	9.0	38.7	0	20.3	2.3	3.6	14.7	11.5
SurgVLP	30.9	58.2	11.6	11.9	34.5	6.0	15.5	24.0

Table 7: Zero-shot triplet recognition results. We report the average precision for each component and the combination of the components. i: instrument, v: verb, t: target, iv: instrument-verb, it: instrument-target, ivt: instrument-verb-target triplet.

mAP (%)	i	v	t	iv	it	ivt
Fully-supervised (Nwoye et al., 2022)	89.5	63.0	45.2	38.7	37.5	30.8
Random	22.7	14.9	10.9	5.0	4.3	3.2
CLIP (Radford et al., 2021)	24.4	15.3	10.3	7.2	4.0	3.1
CLIP*	25.7	15.8	11.0	7.7	4.9	3.7
SurgVLP	32.6	22.8	17.1	10.8	8.6	7.0




Fig. 9: Effect of our designed contextual prompts to the zero-shot transfer of vision-only downstream tasks. Our contextual prompts outperform their counterparts by encoding more specific action and anatomy information, thus boosting phase recognition and instrument-verb recognition.

performs the baseline methods by a large margin. This indicates that our surgery-specific pre-training method helps to recognize surgical phases without explicitly training for them. Also, *SurgVLP* and CLIP* largely outperform the vanilla CLIP (Radford et al., 2021) at P7 (gallbladder extraction) and P5 (gallbladder packing), which involves specimen-bag during the operation. This is consistent with the improvement

in recognizing T7 (specimen bag), indicating that T7 contains generic knowledge that can be transferred across different procedures.

Triplet recognition. As mentioned in (Cole et al., 2022), contrastive learning lags far behind supervised learning on fine-grained visual classification tasks. Therefore, we conduct the action triplet recognition task to study a similar question

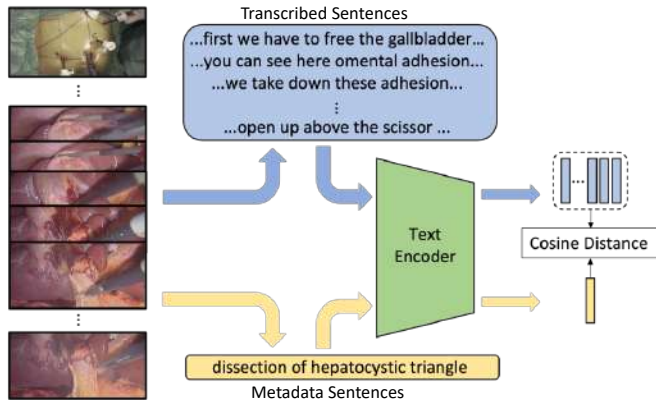


Fig. 10: Text architecture selection. We calculate the cosine similarity score between the transcript texts from ASR and pre-segment texts from metadata to measure which text encoder retains the semantic information between these two texts.

Table 8: Cosine similarity scores comparison among different text encoders.

Models	Similarity Score
Word2Vec Meanpooling	0.11
Word2Vec Maxpooling	0.56
Bert (Devlin et al., 2018)	0.63
CLIP’s Text Encoder (Radford et al., 2021)	0.69
SciBert (Beltagy et al., 2019)	0.80
BioClinicalBert (Huang et al., 2019)	0.86

for surgical computer vision. We aim to assess whether the multi-modal features of *SurgVLP* are only effective for “easy” classification tasks or are also useful for more challenging “fine-grained” visual concepts. As shown in Table 7, *SurgVLP* outperforms the baseline methods in zero-shot triplet recognition by enhancing the recognition of combinations of instruments, verbs, and targets in cholecystectomy. Specifically, we find that *SurgVLP* significantly improves verb and target recognition performance. This suggests that the model effectively learns better anatomical and action knowledge from our SVL-Pretrain dataset and can transfer the learned knowledge to procedure-specific downstream tasks. The zero-shot performance is still lagging behind the fully-supervised method by a large margin, and we expect this could be addressed by scaling the dataset.

4.5.3. Text encoder and prompt analysis

Finally, we present the results of two experiments. First, we conduct experiments to identify an appropriate text encoder by exploring various text encoders for their ability to understand surgery-specific texts. Second, we study the impact of different types of prompts on the vision-only downstream tasks for zero-shot evaluations. In the following, we give two key takeaways from these two experiments.

BioClinicalBert text encoder excels at encoding surgery-specific texts. Conventional text encoders like Word2Vec (Mikolov et al., 2013) show difficulties in handling the out-of-vocabulary surgical terms. Also, a large

textual domain gap exists between the surgical datasets and conventional computer vision datasets. Therefore, we provide a study demonstrating which text encoders can generate representative text embeddings for multi-modal representation learning in the surgical field. Specifically, we consider Bert (Lu et al., 2019) pre-trained on a large-scale web corpus, SciBert (Beltagy et al., 2019) pre-trained on the bio-medical scientific papers, CLIP’s text encoder (Radford et al., 2021) pre-trained on large-scale natural image-caption pairs, and Word2Vec (Mikolov et al., 2013) with different pooling strategies, to compare with *SurgVLP*’s text encoder, i.e., BioClinicalBert (Huang et al., 2019).

We select the video clips from the SVL-Caption dataset. Each clip is associated with a metadata sentence. Given the start and end timestamps of this sentence, we extract the transcribed sentences within the video clip. We assume that the metadata sentences and their corresponding transcribed sentences are paired and have high similarity. Therefore, we feed the sentences into the text encoders mentioned above to generate text embeddings. We then calculate the cosine similarity score between the corresponding metadata and transcription text embeddings, as reported in Table 8. Overall, BioClinicalBert achieves the best results compared to the other text encoders. We show that the bag-of-words approach, e.g., Word2Vec, achieves the lowest score because transcribed sentences from the SVL-Pretrain dataset require complex textual understanding, which Word2Vec fails to capture. Also, as CLIP’s text encoder is trained on a web corpus that contains large amounts of natural texts, it shows the difficulty in handling surgical-specific texts. We use BioClinicalBert in our *SurgVLP* approach as it outperforms all other text encoders in this experiment.

Contextual prompts aligned to downstream tasks result in better performance. To bridge the textual gap between the pre-training and downstream datasets, we manually design the contextual prompts for class names. We use these prompts as text inputs to perform the zero-shot transfer experiments on vision-only downstream tasks. We compare our contextual prompts, explained in section 3.4.2, to two variants: the textual form of the class names and CLIP’s default prompts. The CLIP’s default prompt follows a template “the photo of {class name}” where the class name is replaced with the phase, tool, or triplet class name. As shown in Fig. 9, our designed contextual prompts improve the performance on all three tasks. For example, the *PO* phase, which is “preparation”, is significantly improved. Our contextual prompts also outperform CLIP’s prompts in terms of instrument-verb, instrument-target, and instrument-verb-target recognition. This shows that our proposed contextual prompts encode more action and anatomical linguistic information to benefit fine-grained vision-only downstream tasks.

5. Conclusion

The expensive and laborious process of creating manually annotated datasets has been a main hindrance in developing scalable surgical computer vision AI systems. In this work,

we argue that surgical video lectures available through open surgical e-learning platforms can provide a wealth of multi-modal knowledge to train a scalable system for multi-modal representation learning. We have harnessed this knowledge by creating a multi-modal and multi-procedural dataset comprising 1.4k surgical video lectures. In order to derive effective supervisory signals without manual annotations, we utilize the recent advancements in automatic speech recognition (ASR) systems to transcribe the audio from these videos into textual descriptions. This automated process has resulted in a visual-textual multi-modal surgical dataset consisting of descriptions of surgical events, instrument usage, and anatomical status across various surgical procedures.

In order to tackle the surgery-specific linguistic challenges inherently present in these videos, we utilize text transcriptions from two complementary ASR models, namely Whisper and AWS. The AWS model understands specific surgical terms, whereas the Whisper model understands the overall sentence structure. By combining the complementary knowledge of these two systems, we overcome inherent limitations and inaccuracies present in each ASR system. We then propose a novel contrastive learning objective for multi-modal representation learning. Our approach, called *SurgVLP*, learns effective multi-modal representations by bringing embeddings of multiple text transcriptions and video clip to close proximity in the joint latent space.

To demonstrate the efficacy of the learned joint latent space, we present a range of vision-and-language tasks tailored for surgical computer vision. These tasks include text-based video retrieval, temporal activity grounding, and video captioning, serving as benchmarks for evaluating the multi-modal representation capability of *SurgVLP*. We demonstrate that the learned multi-modal representations are not only useful for these vision-and-language tasks but can also be seamlessly applied to traditional vision-only surgical downstream tasks. We show promising results on these vision-only surgical tasks, namely surgical tool, phase, and triplet recognition, without using any manual annotations.

Acknowledgment

This work was supported by French state funds managed by the ANR within the National AI Chair program under Grant ANR-20-CHIA-0029-01 (ChairAI4ORSafety) and within the Investments for the future program under Grant ANR-10-IAHU-02 (IHU Strasbourg). This work was granted access to the HPC resources of IDRIS under the allocation AD011013704R1 and AD011011631R2 made by GENCI.

We also appreciate the tremendous effort put forth by the open education platforms Websurg (IRCAD), EAES, and YouTube to provide high-quality educational content and make it freely accessible to learners around the world. We are incredibly grateful to the clinicians who have devoted their time and expertise to create and share content on these platforms, making this work possible.

References

- Al Hajj, H., Lamard, M., Conze, P.H., Cochener, B., Quellec, G., 2018. Monitoring tool usage in surgery videos using boosted convolutional and recurrent neural networks. *Medical image analysis* 47, 203–218.
- Alapatt, D., Mascagni, P., Vardazaryan, A., Garcia, A., Okamoto, N., Mutter, D., Marescaux, J., Costamagna, G., Dallemagne, B., Padoy, N., 2021. Temporally constrained neural networks (tcnn): A framework for semi-supervised video semantic segmentation. *arXiv preprint arXiv:2112.13815*.
- Alayrac, J.B., Recasens, A., Schneider, R., Arandjelović, R., Ramapuram, J., De Fauw, J., Smaira, L., Dieleman, S., Zisserman, A., 2020. Self-supervised multimodal versatile networks. *Advances in Neural Information Processing Systems* 33, 25–37.
- 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*.
- Amodei, D., Ananthanarayanan, S., Anubhai, R., Bai, J., Battenberg, E., Case, C., Casper, J., Catanzaro, B., Cheng, Q., Chen, G., et al., 2016. Deep speech 2: End-to-end speech recognition in english and mandarin, in: *International conference on machine learning*, PMLR. pp. 173–182.
- AWS, 2023. Amazon transcribe medical. URL: <https://aws.amazon.com/transcribe/medical/>.
- Banerjee, S., Lavie, A., 2005. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments, in: *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pp. 65–72.
- Beltagy, I., Lo, K., Cohan, A., 2019. Scibert: A pretrained language model for scientific text. *arXiv preprint arXiv:1903.10676*.
- Blum, T., Feußner, H., Navab, N., 2010. Modeling and segmentation of surgical workflow from laparoscopic video, in: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2010: 13th International Conference, Beijing, China, September 20–24, 2010, Proceedings, Part III* 13, Springer. pp. 400–407.
- Blum, T., Padoy, N., Feußner, H., Navab, N., 2008. Modeling and online recognition of surgical phases using hidden markov models, in: *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2008: 11th International Conference, New York, NY, USA, September 6–10, 2008, Proceedings, Part II* 11, Springer. pp. 627–635.
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al., 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33, 1877–1901.
- Changpinyo, S., Sharma, P., Ding, N., Soricut, R., 2021. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3558–3568.
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G., 2020a. A simple framework for contrastive learning of visual representations, in: *International conference on machine learning*, PMLR. pp. 1597–1607.
- Chen, X., Fan, H., Girshick, R., He, K., 2020b. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*.
- Chen, X., Fang, H., Lin, T.Y., Vedantam, R., Gupta, S., Dollár, P., Zitnick, C.L., 2015. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*.
- Chen, Z., Du, Y., Hu, J., Liu, Y., Li, G., Wan, X., Chang, T.H., 2022a. Multi-modal masked autoencoders for medical vision-and-language pre-training, in: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part V*, Springer. pp. 679–689.
- Chen, Z., Song, Y., Chang, T.H., Wan, X., 2020c. Generating radiology reports via memory-driven transformer, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 1439–1449.
- Chen, Z., Zhang, Y., Rosenberg, A., Ramabhadran, B., Moreno, P., Bapna, A., Zen, H., 2022b. Maestro: Matched speech text representations through modality matching. *arXiv preprint arXiv:2204.03409*.
- Cole, E., Yang, X., Wilber, K., Mac Aodha, O., Belongie, S., 2022. When does contrastive visual representation learning work?, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14755–14764.
- Czempiel, T., Paschali, M., Keicher, M., Simson, W., Feussner, H., Kim, S.T., Navab, N., 2020. Tecno: Surgical phase recognition with multi-

- stage temporal convolutional networks, in: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part III 23*, Springer. pp. 343–352.
- Czempiel, T., Paschali, M., Ostler, D., Kim, S.T., Busam, B., Navab, N., 2021. Opera: Attention-regularized transformers for surgical phase recognition, in: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part IV 24*, Springer. pp. 604–614.
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L., 2009. Imagenet: A large-scale hierarchical image database, in: *2009 IEEE conference on computer vision and pattern recognition*, Ieee. pp. 248–255.
- Dergachyova, O., Bouget, D., Huauilmé, A., Morandi, X., Jannin, P., 2016. Automatic data-driven real-time segmentation and recognition of surgical workflow. *International journal of computer assisted radiology and surgery* 11, 1081–1089.
- Devlin, J., Chang, M.W., Lee, K., Toutanova, K., 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- EAES, 2023. The european association of endoscopic surgery. URL: <https://eaes.eu/>.
- Eisenmann, M., Reinke, A., Weru, V., Tizabi, M.D., Isensee, F., Adler, T.J., Godau, P., Cheplygina, V., Kozubek, M., Ali, S., et al., 2022. Biomedical image analysis competitions: The state of current participation practice. *arXiv preprint arXiv:2212.08568*.
- Farha, Y.A., Gall, J., 2019. Ms-tcn: Multi-stage temporal convolutional network for action segmentation, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 3575–3584.
- Funke, I., Jenke, A., Mees, S.T., Weitz, J., Speidel, S., Bodenstedt, S., 2018. Temporal coherence-based self-supervised learning for laparoscopic workflow analysis, in: *OR 2.0 Context-Aware Operating Theaters, Computer Assisted Robotic Endoscopy*, Springer. pp. 85–93.
- Gabeur, V., Sun, C., Alahari, K., Schmid, C., 2020. Multi-modal transformer for video retrieval, in: *European Conference on Computer Vision*, Springer. pp. 214–229.
- Gao, J., Sun, C., Yang, Z., Nevatia, R., 2017. Tall: Temporal activity localization via language query, in: *Proceedings of the IEEE international conference on computer vision*, pp. 5267–5275.
- Garrow, C.R., Kowalewski, K.F., Li, L., Wagner, M., Schmidt, M.W., Engelhardt, S., Hashimoto, D.A., Kenngott, H.G., Bodenstedt, S., Speidel, S., et al., 2021. Machine learning for surgical phase recognition: a systematic review. *Annals of surgery* 273, 684–693.
- Girshick, R., 2015. Fast r-cnn, in: *Proceedings of the IEEE international conference on computer vision*, pp. 1440–1448.
- Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., et al., 2022. Ego4d: Around the world in 3,000 hours of egocentric video, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18995–19012.
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R., 2020. Momentum contrast for unsupervised visual representation learning, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738.
- He, K., Zhang, X., Ren, S., Sun, J., 2016. Deep residual learning for image recognition, in: *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778.
- Huang, K., Altosaar, J., Ranganath, R., 2019. Clinicalbert: Modeling clinical notes and predicting hospital readmission. *arXiv preprint arXiv:1904.05342*.
- Jia, C., Yang, Y., Xia, Y., Chen, Y.T., Parekh, Z., Pham, H., Le, Q., Sung, Y.H., Li, Z., Duerig, T., 2021. Scaling up visual and vision-language representation learning with noisy text supervision, in: *International Conference on Machine Learning*, PMLR. pp. 4904–4916.
- Jin, Y., Dou, Q., Chen, H., Yu, L., Qin, J., Fu, C.W., Heng, P.A., 2017. Svrnet: workflow recognition from surgical videos using recurrent convolutional network. *IEEE transactions on medical imaging* 37, 1114–1126.
- Jin, Y., Long, Y., Chen, C., Zhao, Z., Dou, Q., Heng, P.A., 2021. Temporal memory relation network for workflow recognition from surgical video. *IEEE Transactions on Medical Imaging* 40, 1911–1923.
- Jing, L., Yang, X., Liu, J., Tian, Y., 2018. Self-supervised spatiotemporal feature learning via video rotation prediction. *arXiv preprint arXiv:1811.11387*.
- Krishna, R., Zhu, Y., Groth, O., Johnson, J., Hata, K., Kravitz, J., Chen, S., Kalantidis, Y., Li, L.J., Shamma, D.A., et al., 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision* 123, 32–73.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., Zettlemoyer, L., 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Li, Y., Liang, F., Zhao, L., Cui, Y., Ouyang, W., Shao, J., Yu, F., Yan, J., 2021. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv preprint arXiv:2110.05208*.
- Lin, C.Y., 2004. Rouge: A package for automatic evaluation of summaries, in: *Text summarization branches out*, pp. 74–81.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., Stoyanov, V., 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Lu, J., Batra, D., Parikh, D., Lee, S., 2019. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *Advances in neural information processing systems* 32.
- Madani, A., Namazi, B., Altieri, M.S., Hashimoto, D.A., Rivera, A.M., Pucher, P.H., Navarrete-Welton, A., Sankaranarayanan, G., Brunt, L.M., Okrainec, A., et al., 2020. Artificial intelligence for intraoperative guidance: using semantic segmentation to identify surgical anatomy during laparoscopic cholecystectomy. *Annals of surgery*.
- Mahajan, D., Girshick, R., Ramanathan, V., He, K., Paluri, M., Li, Y., Bharambe, A., Van Der Maaten, L., 2018. Exploring the limits of weakly supervised pretraining, in: *Proceedings of the European conference on computer vision (ECCV)*, pp. 181–196.
- Maier-Hein, L., Eisenmann, M., Sarikaya, D., März, K., Collins, T., Malpani, A., Fallert, J., Feussner, H., Giannarou, S., Mascagni, P., et al., 2022. Surgical data science—from concepts toward clinical translation. *Medical image analysis* 76, 102306.
- Maier-Hein, L., Vedula, S.S., Speidel, S., Navab, N., Kikinis, R., Park, A., Eisenmann, M., Feussner, H., Forestier, G., Giannarou, S., et al., 2017. Surgical data science for next-generation interventions. *Nature Biomedical Engineering* 1, 691–696.
- Mascagni, P., Alapatt, D., Sestini, L., Altieri, M.S., Madani, A., Watanabe, Y., Alseidi, A., Redan, J.A., Alfieri, S., Costamagna, G., et al., 2022. Computer vision in surgery: from potential to clinical value. *npj Digital Medicine* 5, 163.
- Mehri, A., Majumder, N., Bhardwaj, R., Poria, S., 2023. A review of deep learning techniques for speech processing. *arXiv preprint arXiv:2305.00359*.
- Miech, A., Alayrac, J.B., Smaira, L., Laptev, I., Sivic, J., Zisserman, A., 2020. End-to-end learning of visual representations from uncurated instructional videos, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9879–9889.
- Miech, A., Zhukov, D., Alayrac, J.B., Tapaswi, M., Laptev, I., Sivic, J., 2019. Howto100m: Learning a text-video embedding by watching hundred million narrated video clips, in: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 2630–2640.
- Mikolov, T., Chen, K., Corrado, G., Dean, J., 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Narayanan, A., Misra, A., Sim, K.C., Pundak, G., Tripathi, A., Elfeky, M., Haghani, P., Strohmaier, T., Bacchiani, M., 2018. Toward domain-invariant speech recognition via large scale training, in: *2018 IEEE Spoken Language Technology Workshop (SLT)*, IEEE. pp. 441–447.
- Nukrai, D., Mokady, R., Globerson, A., 2022. Text-only training for image captioning using noise-injected clip. *arXiv preprint arXiv:2211.00575*.
- Nwoye, C.I., Yu, T., Gonzalez, C., Seeliger, B., Mascagni, P., Mutter, D., Marescaux, J., Padoy, N., 2022. Rendezvous: Attention mechanisms for the recognition of surgical action triplets in endoscopic videos. *Medical Image Analysis* 78, 102433.
- Oord, A.v.d., Li, Y., Vinyals, O., 2018. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*.
- Owens, A., Efros, A.A., 2018. Audio-visual scene analysis with self-supervised multisensory features, in: *Proceedings of the European conference on computer vision (ECCV)*, pp. 631–648.
- Padoy, N., Blum, T., Ahmadi, S.A., Feussner, H., Berger, M.O., Navab, N., 2012. Statistical modeling and recognition of surgical workflow. *Medical image analysis* 16, 632–641.

- Papineni, K., Roukos, S., Ward, T., Zhu, W.J., 2002. Bleu: a method for automatic evaluation of machine translation, in: Proceedings of the 40th annual meeting of the Association for Computational Linguistics, pp. 311–318.
- Pfeiffer, M., Funke, I., Robu, M.R., Bodenstedt, S., Strenger, L., Engelhardt, S., Roß, T., Clarkson, M.J., Gurusamy, K., Davidson, B.R., et al., 2019. Generating large labeled data sets for laparoscopic image processing tasks using unpaired image-to-image translation, in: Medical Image Computing and Computer Assisted Intervention–MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part V 22, Springer. pp. 119–127.
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sasstry, G., Askell, A., Mishkin, P., Clark, J., et al., 2021. Learning transferable visual models from natural language supervision, in: International Conference on Machine Learning, PMLR. pp. 8748–8763.
- Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I., 2023. Robust speech recognition via large-scale weak supervision, in: International Conference on Machine Learning, PMLR. pp. 28492–28518.
- Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J., 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research* 21, 5485–5551.
- Ramesh, S., Srivastav, V., Alapatt, D., Yu, T., Murali, A., Sestini, L., Nwoye, C.I., Hamoud, I., Sharma, S., Fleurentin, A., et al., 2023. Dissecting self-supervised learning methods for surgical computer vision. *Medical Image Analysis* 88, 102844.
- Rivoir, D., Bodenstedt, S., von Bechtolsheim, F., Distler, M., Weitz, J., Speidel, S., 2019. Unsupervised temporal video segmentation as an auxiliary task for predicting the remaining surgery duration, in: International Workshop on OR 2.0 Context-Aware Operating Theaters, Springer. pp. 29–37.
- Rivoir, D., Funke, I., Speidel, S., 2022. On the pitfalls of batch normalization for end-to-end video learning: A study on surgical workflow analysis. *arXiv preprint arXiv:2203.07976*.
- Rivoir, D., Pfeiffer, M., Docea, R., Kolbinger, F., Riediger, C., Weitz, J., Speidel, S., 2021. Long-term temporally consistent unpaired video translation from simulated surgical 3d data, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 3343–3353.
- Rojas-Muñoz, E., Couperus, K., Wachs, J., 2020. Daisi: Database for ai surgical instruction. *arXiv preprint arXiv:2004.02809*.
- Ross, T., Zimmerer, D., Vemuri, A., Isensee, F., Wiesenfarth, M., Bodenstedt, S., Both, F., Kessler, P., Wagner, M., Müller, B., et al., 2018. Exploiting the potential of unlabeled endoscopic video data with self-supervised learning. *International journal of computer assisted radiology and surgery* 13, 925–933.
- Sain, A., Bhunia, A.K., Chowdhury, P.N., Koley, S., Xiang, T., Song, Y.Z., 2023. Clip for all things zero-shot sketch-based image retrieval, fine-grained or not, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 2765–2775.
- Sanghi, A., Chu, H., Lambourne, J.G., Wang, Y., Cheng, C.Y., Fumero, M., Malekshan, K.R., 2022. Clip-forge: Towards zero-shot text-to-shape generation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 18603–18613.
- Seenivasan, L., Islam, M., Krishna, A.K., Ren, H., 2022. Surgical-vqa: Visual question answering in surgical scenes using transformer, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 33–43.
- Seide, F., Li, G., Chen, X., Yu, D., 2011. Feature engineering in context-dependent deep neural networks for conversational speech transcription, in: 2011 IEEE Workshop on Automatic Speech Recognition & Understanding, IEEE. pp. 24–29.
- Shah, N.A., Sikder, S., Vedula, S.S., Patel, V.M., 2023. Glsformer: Gated-long, short sequence transformer for step recognition in surgical videos. *arXiv preprint arXiv:2307.11081*.
- Soomro, K., Zamir, A.R., Shah, M., 2012. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*.
- Tang, Y., Ding, D., Rao, Y., Zheng, Y., Zhang, D., Zhao, L., Lu, J., Zhou, J., 2019. Coin: A large-scale dataset for comprehensive instructional video analysis, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 1207–1216.
- Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al., 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Twinanda, A.P., Shehata, S., Mutter, D., Marescaux, J., De Mathelin, M., Padoy, N., 2016. Endonet: a deep architecture for recognition tasks on laparoscopic videos. *IEEE transactions on medical imaging* 36, 86–97.
- Vondrick, C., Pirsivavash, H., Torralba, A., 2016. Anticipating visual representations from unlabeled video, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 98–106.
- Vondrick, C., Shrivastava, A., Fathi, A., Guadarrama, S., Murphy, K., 2018. Tracking emerges by colorizing videos, in: Proceedings of the European conference on computer vision (ECCV), pp. 391–408.
- Wang, Y., Long, Y., Fan, S.H., Dou, Q., 2022. Neural rendering for stereo 3d reconstruction of deformable tissues in robotic surgery, in: Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VII, Springer. pp. 431–441.
- Ward, T.M., Mascagni, P., Ban, Y., Rosman, G., Padoy, N., Meireles, O., Hashimoto, D.A., 2021. Computer vision in surgery. *Surgery* 169, 1253–1256.
- Websurg, 2023. Websurg, the online university of ircad. URL: <https://websurg.com/en/>.
- Wu, Z., Xiong, Y., Yu, S.X., Lin, D., 2018. Unsupervised feature learning via non-parametric instance discrimination, in: Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 3733–3742.
- Xu, M., Islam, M., Lim, C.M., Ren, H., 2021. Class-incremental domain adaptation with smoothing and calibration for surgical report generation, in: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part IV 24, Springer. pp. 269–278.
- Xu, M., Islam, M., Ren, H., 2022. Rethinking surgical captioning: End-to-end window-based mlp transformer using patches, in: International Conference on Medical Image Computing and Computer-Assisted Intervention, Springer. pp. 376–386.
- Yengera, G., Mutter, D., Marescaux, J., Padoy, N., 2018. Less is more: Surgical phase recognition with less annotations through self-supervised pre-training of cnn-lstm networks. *arXiv preprint arXiv:1805.08569*.
- YouTube, 2023. Youtube. URL: <https://www.youtube.com/>.
- Yuan, K., Holden, M., Gao, S., Lee, W.S., 2021. Surgical workflow anticipation using instrument interaction, in: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part IV 24, Springer. pp. 615–625.
- Zellers, R., Lu, X., Hessel, J., Yu, Y., Park, J.S., Cao, J., Farhadi, A., Choi, Y., 2021. Merlot: Multimodal neural script knowledge models. *Advances in Neural Information Processing Systems* 34, 23634–23651.
- Zhukov, D., Alayrac, J.B., Cinbis, R.G., Fouhey, D., Laptev, I., Sivic, J., 2019. Cross-task weakly supervised learning from instructional videos, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 3537–3545.