

HecVL: Hierarchical Video-Language Pretraining for Zero-shot Surgical Phase Recognition

Kun Yuan^{1,3}, Vinkle Srivastav^{1,2}, Nassir Navab³, and Nicolas Padoy^{1,2}

¹ University of Strasbourg, CNRS, INSERM, ICube, UMR7357, Strasbourg, France

² IHU Strasbourg, Strasbourg, France

³ CAMP, Technische Universität München, Munich, Germany

Abstract. Natural language could play an important role in developing generalist surgical models by providing a broad source of supervision from raw texts. This flexible form of supervision can enable the model’s transferability across datasets and tasks as natural language can be used to reference learned visual concepts or describe new ones. In this work, we present *HecVL*, a novel hierarchical video-language pretraining approach for building a generalist surgical model. Specifically, we construct a hierarchical video-text paired dataset by pairing the surgical lecture video with three hierarchical levels of texts: at clip-level, atomic actions using transcribed audio texts; at phase-level, conceptual text summaries; and at video-level, overall abstract text of the surgical procedure. Then, we propose a novel fine-to-coarse contrastive learning framework that learns separate embedding spaces for the three video-text hierarchies using a single model. By disentangling embedding spaces of different hierarchical levels, the learned multi-modal representations encode short-term and long-term surgical concepts in the same model. Thanks to the injected textual semantics, we demonstrate that the *HecVL* approach can enable zero-shot surgical phase recognition without any human annotation. Furthermore, we show that the same *HecVL* model for surgical phase recognition can be transferred across different surgical procedures and medical centers. The source code will be made available at <https://github.com/CAMMA-public/HecVL>.

1 Introduction

Developing a single neural network model capable of adapting to different datasets and tasks is a key objective for computer vision. Recent breakthroughs in computer vision methods have begun to fulfill this goal by transitioning from task-specific models [2] to generalist models [27, 18]. These generalist models have shown potential in solving a wide range of downstream tasks and datasets, including various types of object segmentation [28] and zero-shot image and video classification [11]. An essential feature of these models is their ability to be supervised through natural language texts. The generality of natural language allows it to express a broader set of visual concepts, thereby effectively guiding and supervising these models [3, 21].




Fig. 1. Hierarchical video-text pairs in surgical lecture videos. Conventional methods [26] utilize only clip-level video-text pairs, while our HecVL utilizes different hierarchical levels of pairs to perform video-language pretraining.

Yet, within the domain of surgical video analysis, predominant methods still lean toward task-specific models [14, 15, 25, 22]. This is mainly due to the inherent complexity present in the surgical videos, i.e., surgical videos can last several hours while capturing intricate hierarchical surgical activities. Therefore, those methods manually define different levels of categories and annotate large amounts of frames to provide extensive supervision. However, the procedure- and center-specific annotations lead to degraded transferability across procedures and medical centers [10]. While a surgical foundation model [23, 19] is proposed to address the above issue, it focuses only on pure images and ignores the complementary information from other modalities, i.e., language. Also, it still requires finetuning on the downstream dataset to enable transferability. Considering that natural language texts have become a unifying element for generalist models, this work explores whether they can be used to both understand the hierarchical intricacies of surgical videos and enable the generalized zero-shot transfer by processing category labels into texts, without the need for manual annotation. As the task of surgical phase recognition is essential for computer-assisted surgery [17, 1, 20, 9, 4, 7], we use it as a suitable test bench to evaluate our joint visual and textual hierarchical representations.

This work introduces *HecVL*, a Hierarchical Encoded Contrastive Video-language pretraining framework, which learns rich multi-modal representations at different hierarchies of surgical video. Developing such an approach presents a significant challenge due to the lack of surgical video datasets with hierarchical textual supervision. SurgVLP [26] has introduced the first large-scale video-text paired dataset, i.e., *SVL*, by transcribing hundreds of surgical lecture videos into narration texts. We obtain the *SVL* dataset, courtesy of the authors of SurgVLP, and further extend it by incorporating hierarchical-level texts using the metadata of each lecture video. We construct three levels of the hierarchical video-text pairs for each surgical lecture video: *clip-level*, *phase-level*, and *video-*

level. The clip-level video-text pairs contain short video clips of few seconds duration along with narration texts transcribed from lecture audio for capturing the short-term activity. The phase-level video-text pairs contain longer video segments with conceptual text summaries for capturing longer surgical video activity. Finally, the video-level video-text pairs are the entire surgical lecture videos paired with abstract paragraphs encapsulating the goal and the main key points of the surgery. These three levels of hierarchical video-text pairs allow for a more detailed understanding of surgical procedures, capturing both the atomic details and broader contexts, as illustrated in Fig. 1.

Given the hierarchical video-text pair dataset, we propose a *fine-to-coarse contrastive* learning strategy to effectively exploit the hierarchical textual information encoded in the dataset. We construct three separate embedding spaces for each type of hierarchical video-text pair. We first build up a fine-grained embedding space using clip-level video-text pairs, followed by aggregating the fine-grained features to construct the coarse-grained embedding spaces, which embed phase-level and video-level text pairs. We learn these three different embedding spaces through multi-modal contrastive learning using the InfoNCE loss [16]. We show in the experiments that our fine-to-coarse contrastive learning strategy outperforms the approach of projecting all hierarchical texts into a single embedding space and learning only one such space.

We demonstrate the zero-shot transferability and the generalization of our approach by performing surgical phase recognition on three different surgical procedures, cholecystectomy [20], hysterectomy [24], and gastric bypass [10], without using any ground truth labels. The learned multi-modal representations demonstrate their transferability in not only identifying surgical concepts across various surgical procedures but also in extending to different medical centers. We hope that the *HecVL* approach could pave the path for developing more generalist models in the domain of surgical computer vision.

2 Method

We propose *HecVL*, a novel hierarchical video-language pretraining method that learns multi-modal embeddings by capturing clip-, phase-, and video-level video-text pairs from surgical lecture videos. Fig. 2 gives an overview of our method. Sec. 2.1 describes the construction of hierarchical video-text pairs. Sec. 2.2 formalizes the fine-to-coarse contrastive learning strategy. Sec. 2.3 and 2.4 describe the loss function and the training pipeline, respectively.

2.1 Hierarchical video-text pairs

The *HecVL* approach is designed to leverage a hierarchically annotated video-text pair dataset, $D = \{(V_i, N_i, C_i, A_i)\}_{i=1}^{|D|}$, where V_i is a long surgical lecture video composed of a sequence of short-term video clips (each lasting tens of seconds). Each lecture video V_i is paired with three levels of textual annotations from different levels of granularities ranging from fine-grained to coarse-grained,




Fig. 2. Pipeline of the *HecVL* approach. (a) Conventional video-language methods embed video clips and texts of different granularities into the same embedding space. (b) The *HecVL* approach considers the granularity differences and constructs three embedding spaces for clip-, phase-, and video-level representation learning. (c) The fine-grained embedding space ($S_{narration}$) is learned first, followed by learning of coarse-space embedding spaces ($S_{abstract}$ and $S_{concept}$) using a temporal aggregation function to aggregate the visual and the textual embeddings.

i.e., **clip-level** narration texts (N_i), **phase-level** concept texts (C_i), and **video-level** abstract texts (A_i).

The clip-level narration texts (N_i) are sequences of narrations describing the atomic actions for short-term video clips, the phase-level concept texts (C_i) are sequences of conceptual text summaries describing the high-level surgical activities for long-term video phases, and the video-level abstract paragraph texts (A_i) are the abstract paragraph texts summarizing the entire surgical lecture video including patient’s history and surgical technique. These three levels of video-text pairs provide complementary textual supervision at multiple hierarchies for representation learning, as illustrated in Fig. 1.

2.2 Fine-to-coarse contrastive learning

Given the hierarchically annotated video-text pair dataset as described above, we aim to optimize a visual encoder $\mathcal{F}_v$ and a textual encoder $\mathcal{F}_t$ for the multi-modal hierarchical representation learning. This is achieved by constructing different embedding spaces for hierarchical video-text pairs, as described below.

Embedding spaces at different hierarchical levels: given the visual encoder $\mathcal{F}_v$ and textual encoder $\mathcal{F}_t$, we first extract **clip-level** visual embeddings

$\mathcal{F}_v(v_{ij})$ and textual embeddings $\mathcal{F}_t(n_{ij})$ from the short-term video clips $v_{ij} \in V_i$ and their corresponding clip-level narration texts $n_{ij} \in N_i$. These multi-modal embeddings are represented in the fine-grained embedding space $S_{narration}$.

Then, we construct another embedding space $S_{concept}$ by exploiting **phase-level** textual supervision. We define V^c and N^c as the sets of short-term video clips v_{ij} and narration texts n_{ij} temporally corresponding to phase-level concept texts $c_{ij} \in C_i$. We extract the textual embeddings $\mathcal{F}_t(c_{ij})$ using the textual encoder $\mathcal{F}_t$. Subsequently, we define an aggregator function $Agg()$, which takes clip-level visual and textual embeddings, V^c and N^c , as input and performs average pooling on them. The aggregated visual embeddings $Agg(\mathcal{F}_v(V^c))$ and textual embeddings $Agg(\mathcal{F}_t(N^c))$ are represented in embedding space $S_{concept}$.

Finally, we construct a **video-level** embedding space, $S_{abstract}$, using video-level abstract texts (A_i). Similar to constructing the phase-level embedding space, we define V^a and N^a as the evenly sampled sets of short-term video clips v_{ij} and narration texts n_{ij} . As there is only one abstract text for the entire video, V^a and N^a correspond to the video-level abstract texts A_i . We extract the textual embeddings $\mathcal{F}_t(A_i)$ from the abstract text A_i . Then, the aggregator function $Agg()$ is employed to aggregate the clip-level visual and textual embeddings. The resulting aggregated visual embeddings $Agg(\mathcal{F}_v(V^a))$ and textual embeddings $Agg(\mathcal{F}_t(N^a))$ are represented in space $S_{abstract}$. The construction of all three embedding spaces is illustrated in the Fig. 2.

Fig. 2 (a) shows another way to conduct video-language pretraining by projecting all the video clips and the three levels of texts to a single embedding space. We show in experiments that this increases the ambiguity as video clips might be pushed to both narration and concept texts with dissimilar semantics.

2.3 Training objectives

We propose a joint contrastive loss function to enhance the similarity score between matching visual and textual embeddings compared to non-matching pairs in $S_{narration}$, $S_{concept}$, and $S_{abstract}$. At the clip level, we use the loss function from the SurgVLP [26] $\mathcal{L}_{clip}$ (also given in the supplementary) to correlate short-term video clips with the narrations from two different audio transcription systems (ASR). At the phase- and the video-level, we use the InfoNCE [16] loss to correlate the aggregated short-term visual and textual embeddings to phase- and video-level textual embeddings, respectively, as given below:

$$\mathcal{L}_{phase} = -\frac{1}{B} \sum_{i=1}^B \log \left(\frac{\exp(Agg(\mathcal{F}_v(V^c))^T \cdot \mathcal{F}_t(c_i)/\tau)}{\sum_{j=1}^B \exp(Agg(\mathcal{F}_v(V^c))^T \cdot \mathcal{F}_t(c_j)/\tau)} + \frac{\exp(Agg(\mathcal{F}_t(N^c))^T \cdot \mathcal{F}_t(c_i)/\tau)}{\sum_{j=1}^B \exp(Agg(\mathcal{F}_t(N^c))^T \cdot \mathcal{F}_t(c_j)/\tau)} \right) \quad (1)$$

$$\mathcal{L}_{video} = -\frac{1}{B} \sum_{i=1}^B \log \left(\frac{\exp(Agg(\mathcal{F}_v(V^a))^T \cdot \mathcal{F}_t(A_i)/\tau)}{\sum_{j=1}^B \exp(Agg(\mathcal{F}_v(V^a))^T \cdot \mathcal{F}_t(A_j)/\tau)} + \frac{\exp(Agg(\mathcal{F}_t(N^a))^T \cdot \mathcal{F}_t(A_i)/\tau)}{\sum_{j=1}^B \exp(Agg(\mathcal{F}_t(N^a))^T \cdot \mathcal{F}_t(A_j)/\tau)} \right) \quad (2)$$

Here, B is the batch size, and τ is a temperature hyper-parameter, which regulates the probability distribution over positive and negative pairs within the embedding space. The numerator term denotes the cosine similarity between

matched visual and the textual pairs, i.e., *positive pairs*, and the denominator term denotes the cosine similarity between the unmatched visual and the textual pairs, i.e., *negative pairs*. The total loss function $\mathcal{L}$ is the sum of clip-level, phase-level, and video-level loss functions, i.e., $\mathcal{L} = \mathcal{L}_{clip} + \mathcal{L}_{phase} + \mathcal{L}_{video}$.

2.4 Training pipeline

Given the previously described training objectives, the challenge lies in effectively training $\mathcal{F}_v$ and $\mathcal{F}_t$ across all three levels of embedding spaces. We aim to train only one set of visual and textual encoders for all three levels of embedding spaces, ensuring the encoders are optimized for capturing both short-term and long-term semantics. We propose an *alternating training strategy*, i.e., we first optimize L_{clip} for m batches; subsequently, we optimize L_{phase} for n batches and L_{video} for l batches, and then we repeat. We observe that the proposed training strategy not only converges faster but also circumvents the catastrophic forgetting issue [5] that could arise when training the model on clip-level embeddings and then fine-tuning it for phase- and video-level embeddings, or vice-versa.

3 Experiment setup

3.1 Dataset

Pretraining dataset: we use the Surgical Video Lecture dataset (SVL), which was kindly shared by the authors of SurgVLP [26]. We further expand the dataset by including additional phase- and video-level video-text pairs using the metadata of each lecture video. The metadata for each lecture video contains the title of the procedure, the abstract summary, and the key steps. In total, we have 25,578 clip-level, 10,304 phase-level, and 1,076 video-level video-text pairs.

Downstream datasets and evaluation: we perform the evaluation on three public datasets: Cholec80 [20], AutoLaparo [24], StrasBypass70 and BernBypass70 [10]. In our evaluation, we perform surgical phase recognition in the zero-shot setting, which directly evaluates the model on downstream datasets without performing any fine-tuning. Here, class labels are transformed into textual prompts, and their embeddings categorize the image visual embeddings, reflecting the joint embedding space’s effectiveness (details of the constructed textual prompts are given in the Supplementary).

3.2 Implementation details

Network architecture: We use the ResNet-50 model [6] pretrained on ImageNet as visual encoder $\mathcal{F}_v$ and BioClinicalBert [8] as the textual encoder $\mathcal{F}_t$. We sample 4, 8, and 32 frames for each clip-level, phase-level, and video-level video segment. We encode the frames, followed by an average pooling to generate a feature vector for a video segment. The architectures of visual and text encoders, $\mathcal{F}_v$ and $\mathcal{F}_t$, are the same as SurgVLP [26] for a fair comparison.

Table 1. Zero-shot phase recognition results on cholecystectomy and hysterectomy. CLIP**/CLIP*: random/OpenAI initialized CLIP model and pretraining with SVL dataset [26].

Model	Pretraining dataset	Cholec-80		Autolaparo	
		Top-1	Acc. F1 Score	Top-1	Acc. F1 Score
MIL-NCE [13]	Howto100M	7.8	7.3	9.9	7.9
CLIP [18]	CLIP400M	30.8	13.1	17.4	9.1
CLIP** [18]	SVL-Pretrain	29.4	10.4	15.3	10.9
CLIP* [18]	SVL-Pretrain	33.8	19.6	18.9	16.2
SurgVLP [26]	SVL-Pretrain	34.7	24.4	21.3	16.6
HecVL	SVL-Pretrain	41.7	26.3	23.3	18.9

Training parameters: we pretrain the model with one 80 GB NVIDIA A100 GPUs for 200 epochs. We use the AdamW [12] optimizer with a learning rate of $5e - 5$. We alternatively train with $m = 25$ batches of clip-level pairs, followed by $n = 15$ and $l = 115$ batches of phase- and video-level pairs. We use a batch size B of 120/60/10 per GPU for clip-/phase-/video-level video-text pairs.

4 Results and discussions

4.1 Zero-shot phase recognition

Results on zero-shot surgical phase recognition demonstrate if the learned joint visual and textual representations can correlate semantically similar surgical scene images and surgical texts. We compare our method to CLIP [18] to show the benefits of surgical-specific pretraining, and to SurgVLP to emphasize hierarchical pretraining advantages. Tab. 1 and 2 show that our *HecVL* achieves state-of-the-art performance for all the datasets in the zero-shot setting. The consistent boost across cholecystectomy [20], hysterectomy [24], and gastric bypass [10] procedures show the generalizable and transferable features of *HecVL* across different surgical types. Also, we show significant improvement compared to the methods pretrained on the conventional computer vision datasets, i.e., MIL-NCE [13], CLIP [18], which fails in recognizing the surgical concepts.

4.2 Multi-center phase recognition

Here, we examine the ability of the *HecVL* approach to transfer the knowledge learned from hierarchical video-text data to different medical centers, as shown in Tab. 2. Overall, our *HecVL* model achieves the best performance across two medical centers compared to the other methods. Interestingly, the performance of the BernBypass70 is lower than the StrasBypass70. This may be because there are significant differences in the workflow followed at the Bern center, with many phases and steps not routinely performed. Therefore, the textual prompts

Table 2. Zero-shot phase recognition results across medical centers on gastric bypass surgery. We evaluate our model on the test split of StrasBypass70 and BernBypass70 from the University Hospital of Strasbourg and Bern University Hospital.

Model	Pretraining dataset	StrasBypass70		BernBypass70	
		Top-1	Acc. F1 Score	Top-1	Acc. F1 Score
MIL-NCE [13]	Howto100M	3.8	2.7	2.5	2.0
CLIP [18]	CLIP400M	12.0	6.8	9.7	5.1
CLIP** [18]	SVL-Pretrain	6.3	3.5	4.9	2.3
CLIP* [18]	SVL-Pretrain	15.8	8.6	17.8	7.1
SurgVLP [26]	SVL-Pretrain	10.8	6.9	11.4	7.2
HecVL	SVL-Pretrain	26.9	18.3	22.8	13.6

Table 3. Ablation study on different levels. We conduct zero-shot phase recognition and report F1 Score. Single: single embedding space as in Fig. 1(a).

Model	Action	Phase	Abstract	Cholec80	AutoLaparo	StrasBypass70	BernBypass70
SurgVLP	✓	×	×	24.4	16.6	6.9	7.2
<i>HecVL</i>	✓	✓	×	25.2	16.9	15.5	12.9
Single	✓	✓	✓	25.6	15.7	14.6	14.0
<i>HecVL</i>	✓	✓	✓	26.3	18.9	18.3	13.6

designed based on the Strasbourg center’s protocol (see Supplementary for more details) lead to degraded performance when applied to the different centers. To address this, a center-specific textual prompts construction is required.

4.3 Ablation study

Tab. 3 provides the ablation analysis of the contribution of each level of video-text pairs. Specifically, adding the phase-level video-text pairs yields significant improvements in gastric bypass surgical phase recognition compared to the SurgVLP. This trend is pronounced in the zero-shot scenario across all surgical procedures. We also compare to a baseline model: Single in Fig. 2 (a), which embeds action, phase, and abstract texts in a single embedding space, to support our fine-to-coarse strategy. Tab. 3 shows that Single model has inconsistent performance across datasets. This inconsistency implies that the single embedding approach might blur essential distinctions across hierarchical levels compared to maintaining separate embedding spaces for different hierarchical levels.

5 Conclusion

The next generation of scalable and generalizable surgical computer vision systems demands multi-modality models capable of adapting to different surgical procedures with little to no manual annotations. In this work, we design *HecVL*,

a single multi-modality model *HecVL* capable of adapting to the different surgical procedures and centers without using any manual annotations. The core of our contribution lies in developing a hierarchical contrastive learning strategy to exploit textual supervision at multiple granular levels, ranging from short-term surgical actions to long-term high-level surgical concepts. Extensive experimental results demonstrate its efficacy in achieving zero-shot surgical phase recognition across different procedures and medical centers.

Acknowledgments. This work has received funding from the European Union (ERC, CompSURG, 101088553). Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. This work was also partially supported by French state funds managed by the ANR under Grants ANR-20-CHIA-0029-01 and ANR-10-IAHU-02. This work was granted access to the HPC resources of IDRIS under the allocation AD011013704R1 and AD011011631R2 made by GENCI. This work was granted access to the HPC resources of IDRIS under the allocations AD011013704R1, AD011011631R2, and AD011011631R3 made by GENCI. The authors would like to acknowledge the High Performance Computing Center of the University of Strasbourg for supporting this work by providing scientific support and access to computing resources. Part of the computing resources were funded by the Equipex Equip@Meso project (Programme Investissements d’Avenir) and the CPER Alsacalcul/Big Data.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Blum, T., Feußner, H., Navab, N.: 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. pp. 400–407. Springer (2010)
2. Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence* **40**(4), 834–848 (2017)
3. Chen, T., Saxena, S., Li, L., Lin, T.Y., Fleet, D.J., Hinton, G.E.: A unified sequence interface for vision tasks. *Advances in Neural Information Processing Systems* **35**, 31333–31346 (2022)
4. Czempiel, T., Paschali, M., Keicher, M., Simson, W., Feussner, H., Kim, S.T., Navab, N.: 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. pp. 343–352. Springer (2020)
5. Goodfellow, I.J., Mirza, M., Xiao, D., Courville, A., Bengio, Y.: An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211* (2013)

6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
7. Hu, M., Xia, P., Wang, L., Yan, S., Tang, F., Xu, Z., Luo, Y., Song, K., Leitner, J., Cheng, X., et al.: Ophnet: A large-scale video benchmark for ophthalmic surgical workflow understanding. arXiv preprint arXiv:2406.07471 (2024)
8. Huang, K., Altosaar, J., Ranganath, R.: Clinicalbert: Modeling clinical notes and predicting hospital readmission. arXiv preprint arXiv:1904.05342 (2019)
9. Jin, Y., Dou, Q., Chen, H., Yu, L., Qin, J., Fu, C.W., Heng, P.A.: Sv-rcnet: workflow recognition from surgical videos using recurrent convolutional network. IEEE transactions on medical imaging **37**(5), 1114–1126 (2017)
10. Lavanchy, J.L., Ramesh, S., Dall’Alba, D., Gonzalez, C., Fiorini, P., Muller-Stich, B., Nett, P.C., Marescaux, J., Mutter, D., Padoy, N.: Challenges in multi-centric generalization: Phase and step recognition in roux-en-y gastric bypass surgery. arXiv preprint arXiv:2312.11250 (2023)
11. Lin, W., Karlinsky, L., Shvetsova, N., Possegger, H., Kozinski, M., Panda, R., Feris, R., Kuehne, H., Bischof, H.: Match, expand and improve: Unsupervised finetuning for zero-shot action recognition with language knowledge. arXiv preprint arXiv:2303.08914 (2023)
12. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
13. Miech, A., Alayrac, J.B., Smaira, L., Laptev, I., Sivic, J., Zisserman, A.: 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 (2020)
14. Ming, H., Lin, W., Siyuan, Y., Don, M., Qingli, R., Peng, X., Wei, F., Peibo, D., Lie, J., Zongyuan, G.: Nurvid: A large expert-level video database for nursing procedure activity understanding. In: Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track (2023)
15. Nwoye, C.I., Yu, T., Gonzalez, C., Seeliger, B., Mascagni, P., Mutter, D., Marescaux, J., Padoy, N.: Rendezvous: Attention mechanisms for the recognition of surgical action triplets in endoscopic videos. Medical Image Analysis **78**, 102433 (2022)
16. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
17. Padoy, N., Blum, T., Ahmadi, S.A., Feussner, H., Berger, M.O., Navab, N.: Statistical modeling and recognition of surgical workflow. Medical image analysis **16**(3), 632–641 (2012)
18. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
19. Ramesh, S., Srivastav, V., Alapatt, D., Yu, T., Murali, A., Sestini, L., Nwoye, C.I., Hamoud, I., Sharma, S., Fleurentin, A., et al.: Dissecting self-supervised learning methods for surgical computer vision. Medical Image Analysis **88**, 102844 (2023)
20. Twinanda, A.P., Shehata, S., Mutter, D., Marescaux, J., De Mathelin, M., Padoy, N.: Endonet: a deep architecture for recognition tasks on laparoscopic videos. IEEE transactions on medical imaging **36**(1), 86–97 (2016)
21. Wang, H., Jin, Y., Zhu, L.: Dynamic interactive relation capturing via scene graph learning for robotic surgical report generation. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 2702–2709. IEEE (2023)

22. Wang, H., Zhu, L., Yang, G., Guo, Y., Zhang, S., Xu, B., Jin, Y.: Video-instrument synergistic network for referring video instrument segmentation in robotic surgery. arXiv preprint arXiv:2308.09475 (2023)
23. Wang, Z., Liu, C., Zhang, S., Dou, Q.: Foundation model for endoscopy video analysis via large-scale self-supervised pre-train. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 101–111. Springer (2023)
24. Wang, Z., Lu, B., Long, Y., Zhong, F., Cheung, T.H., Dou, Q., Liu, Y.: Autolaparo: A new dataset of integrated multi-tasks for image-guided surgical automation in laparoscopic hysterectomy. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 486–496. Springer (2022)
25. Wu, L., Hu, Z., Ji, Y., Luo, P., Zhang, S.: Multi-frame collaboration for effective endoscopic video polyp detection via spatial-temporal feature transformation. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part V 24. pp. 302–312. Springer (2021)
26. Yuan, K., Srivastav, V., Yu, T., Lavanchy, J., Mascagni, P., Navab, N., Padoy, N.: Learning multi-modal representations by watching hundreds of surgical video lectures. arXiv preprint arXiv:2307.15220 (2023)
27. Zou, X., Dou, Z.Y., Yang, J., Gan, Z., Li, L., Li, C., Dai, X., Behl, H., Wang, J., Yuan, L., et al.: Generalized decoding for pixel, image, and language. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15116–15127 (2023)
28. Zou, X., Yang, J., Zhang, H., Li, F., Li, L., Wang, J., Wang, L., Gao, J., Lee, Y.J.: Segment everything everywhere all at once. *Advances in Neural Information Processing Systems* **36** (2024)