

Operating Room Workflow Analysis via Reasoning Segmentation over Digital Twins

Yiqing Shen^[0000-0001-7866-3339], Chenjia Li, Bohan Liu, Cheng-Yi Li, Tito Porras, and Mathias Unberath^(✉)

Johns Hopkins University, Baltimore, MD, USA
{yshen92, unberath}@jhu.edu

Abstract. Analyzing operating room (OR) workflows to derive quantitative insights into OR efficiency is important for hospitals to maximize patient care and financial sustainability. Prior work on OR-level workflow analysis has relied on end-to-end deep neural networks. While these approaches work well in constrained settings, they are limited to the conditions specified at development time and do not offer the flexibility necessary to accommodate the OR workflow analysis needs of various OR scenarios (*e.g.*, large academic center *vs.* rural provider) without data collection, annotation, and retraining. Reasoning segmentation (RS) based on foundation models offers this flexibility by enabling automated analysis of OR workflows from OR video feeds given only an implicit text query related to the objects of interest. Due to the reliance on large language model (LLM) fine-tuning, current RS approaches struggle with reasoning about semantic/spatial relationships and show limited generalization to OR video due to variations in visual characteristics and domain-specific terminology. To address these limitations, we first propose a novel digital twin (DT) representation that preserves both semantic and spatial relationships between the various OR components. Then, building on this foundation, we propose ORDiRS (Operating Room Digital twin representation for Reasoning Segmentation), an LLM-tuning-free RS framework that reformulates RS into a “reason-retrieval-synthesize” paradigm. Finally, we present ORDiRS-Agent, an LLM-based agent that decomposes OR workflow analysis queries into manageable RS sub-queries and generates responses by combining detailed textual explanations with supporting visual evidence from RS. Experimental results on both an in-house and a public OR dataset demonstrate that our ORDiRS achieves a cIoU improvement of 6.12%-9.74% compared to the existing state-of-the-arts.

Keywords: Operation Room Efficiency · Reasoning Segmentation · Digital Twin Representation · Large Language Model (LLM).

1 Introduction

Operating room (OR) efficiency is important for optimizing surgical throughput, resource utilization, and patient safety, with direct implications for hospital financial sustainability [2]. Video analysis provides objective monitoring of

OR workflow patterns and resource utilization, generating data-driven insights that may enable operational improvements [9]. However, traditional deep learning methods for OR video analysis are limited to narrowly defined closed-set tasks, which are typically formulated in an end-to-end manner [8]. These previous methods can identify surgical instruments or track staff movements, but their closed-set nature 1) limits their flexibility inhibiting their utility for the various clinical settings these algorithms may be used in, and 2) results in them failing to interpret open-set text queries that require reasoning. To overcome these limitations, reasoning segmentation (RS) [5] was introduced. Unlike traditional segmentation methods, RS processes text queries that only implicitly refer to the required segmentation targets, thus requiring both spatial and semantic understanding [5, 18, 14, 17]. Thus, RS may enable automated OR efficiency analysis by identifying workflow bottlenecks, monitoring safety protocol compliance, and evaluating team coordination during procedural steps [12, 3]

However, current RS approaches, including LISA [5, 20], VISA [18], Reason3D [1], rely heavily on fine-tuning multimodal large language models (MLLMs) for both perception and reasoning. These models struggle with implicit text queries that require multi-step reasoning across semantic and spatial relationships due to inherent limitations in MLLM token representations, which arbitrarily discretize continuous physical information in the OR. This limitation further hinders OR efficiency analysis, which depends on processing complex sequences of entity interactions while maintaining spatial awareness. Moreover, these RS models struggle with domain adaptation in OR settings, as they are typically developed and fine-tuned on general-purpose datasets like *ReasonSeg* [5], which differ from medical environments in both visual characteristics and domain-specific terminology [13]. Even with OR-specific fine-tuning, generalization remains challenging due to data distribution variability across hospitals and healthcare institutions. Finally, reliance on MLLM fine-tuning requires frequent updates to maintain compatibility with evolving backbones, potentially disrupting continuous monitoring and increasing implementation costs in OR environments [16].

To overcome these limitations, we introduce a digital twin (DT) representation that models the OR environment from video. Our DT representation integrates specialized vision models, such as SAM for identifying and segmenting objects in the OR [4, 10] and DepthAnything2 for estimating spatial depth [19], to capture and encode semantic and spatial relationships among personnel, equipment, and activities. We then propose a tuning-free RS approach that restructures RS into a “reason-retrieve-synthesize” paradigm, enabling adaptive segmentation without fine-tuning. The large language model (LLM) performs zero-shot reasoning on the DT representation in three steps: (1) interpreting the implicit text query, (2) analyzing relevant objects and their spatial and semantic relationships, and (3) generating segmentation masks. For example, when tasked with segmenting non-essential personnel entering sterile fields, the LLM first interprets the query requirements, then identifies relevant staff members by analyzing semantic relationships in the DT representation, evaluates their spatial positions relative to sterile zones, and generates segmentation masks accordingly.

In contrast to prior approaches to RS that leverage MLLMs to process query and video directly, our approach separates perception from reasoning, allowing the LLM to focus on high-level analysis using structured information from the DT representation.

The major contributions are three-fold. First, we propose a novel DT representation for OR video that captures and preserves semantic and spatial relationships. Second, we introduce a tuning-free RS approach for OR data, called Operating Room Digital twin representation for Reasoning Segmentation (ORDiRS), which leverages the DT representation as an intermediate layer to separate perception from reasoning in the LLM. Third, building on ORDiRS, we propose **ORDiRS-Agent**, an agent framework designed specifically for OR efficiency analysis. It first breaks down user queries into manageable RS sub-queries for ORDiRS to process. It then aggregates the RS results to generate multi-modal responses, combining textual explanations with corresponding RS masks to enhance transparency.

2 Methods

OR Digital Twin Representation Construction Our DT representation creates a representation of the OR environment by integrating outputs from multiple specialized vision foundation models to preserve both semantic relationships between OR components and their spatial configurations, as shown in Fig. 1. We process the OR video sequence $I = [I^{(1)}, I^{(2)}, \dots, I^{(T)}]$ frame by frame, where each frame $I^{(t)}$ at time t is encoded into a structured JSON $\mathcal{J}^{(t)}$. We select JSON as our DT representation format due to its hierarchical structure that naturally accommodates nested relationships between frame-level metadata and instance-level attributes while maintaining compatibility with LLMs and enabling efficient querying of spatial-temporal information.

The DT representation construction process begins with OWLv2 [7], an open vocabulary detection model that generates a set of bounding boxes $B^{(t)} = \{b_i^{(t)}\}_{i=1}^N$ for each frame $I^{(t)}$, where each $b_i^{(t)} = [x_{i,\min}^{(t)}, y_{i,\min}^{(t)}, x_{i,\max}^{(t)}, y_{i,\max}^{(t)}]$ represents the spatial coordinates of a detected object with a confidence score $\alpha_i^{(t)} \in [0, 1]$. Using each bounding box $b_i^{(t)}$ as a prompt, we employ the SAM2 [10] to generate precise instance segmentation masks $M^{(t)} = \{m_i^{(t)}\}_{i=1}^N$, where $m_i^{(t)}$ is a binary mask indicating pixel-wise object boundaries with confidence score $\beta_i^{(t)}$ derived from the IoU prediction head in SAM. For semantic understanding, we use LLaVA-7B to analyze image regions cropped from $B^{(t)}$ [6]. LLaVA-7B generates semantic descriptors $S^{(t)} = \{s_i^{(t)}\}_{i=1}^N$, where each $s_i^{(t)}$ contains natural language descriptions that capture object attributes, roles, and contextual relationships within the OR environment. Analogously, spatial relationships are encoded through depth information using DepthAnything2 [19], which generates a dense depth map $D^{(t)}$ for each frame uniformly. For each instance i , we compute depth statistics within its mask region as $d_i^{(t)} = \{D^{(t)}(p) | p \in m_i^{(t)}\}$, where p represents pixel coordinates. The mean depth $\mu_i^{(t)}$ and standard devi-

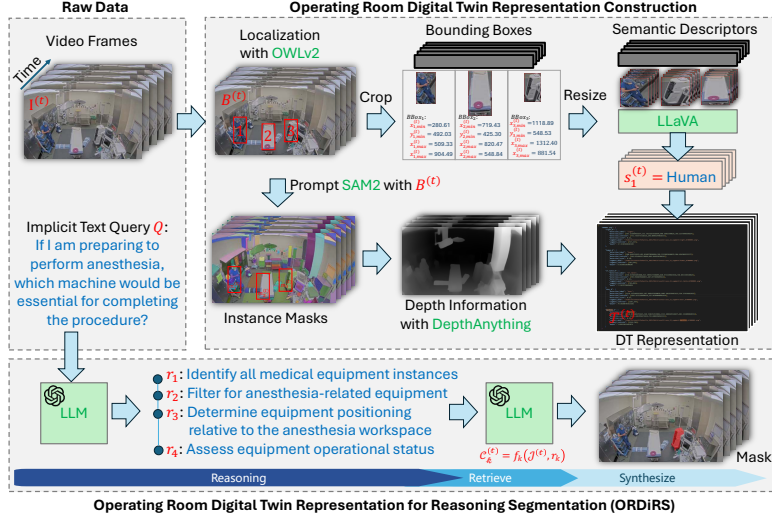


Fig. 1. Overview of the ORDiRS framework. The pipeline consists of two main components. (1) DT representation construction: Processing raw OR video frames through multiple vision foundation models, culminating in a structured JSON; (2) RS stage: Implementing a three-stage “reason-retrieve-synthesize” paradigm where LLM-based reasoning decomposes an implicit text query into atomic reasoning requirements.

ation $\sigma_i^{(t)}$ are calculated from $d_i^{(t)}$ to characterize the instance’s position and depth variation. This depth information enables monitoring of spatial relationships for maintaining proper distances between sterile and non-sterile zones in the OR environment.

Reasoning Segmentation with DT Representation Given the $\mathcal{J}^{(t)}$ and an implicit text query Q , our RS approach (ORDiRS) follows a three-stage paradigm: “reason-retrieve-synthesize” (Fig. 1). This design enables zero-shot reasoning by leveraging LLMs’ inherent capabilities while maintaining interpretability through structured intermediate outputs. The reasoning stage begins by decomposing the implicit query Q into explicit reasoning requirements $R = \{r_k\}_{k=1}^K$, where K represents the number of atomic requirements identified by the LLM through chain-of-thought prompting. Each requirement r_k specifies the semantic or spatial conditions that the target instances must satisfy. The retrieval stage processes these requirements against the DT representation through a series of filtering operations. For each requirement r_k , we define a filtering function f_k that evaluates instances based on their attributes in $\mathcal{J}^{(t)}$: $\mathcal{C}_k^{(t)} = f_k(\mathcal{J}^{(t)}, r_k)$, where $\mathcal{C}_k^{(t)}$ represents the set of candidate instances satisfying requirement r_k at time t . To facilitate semantic filtering, f_k leverages the LLM’s natural language understanding capabilities by constructing prompts that evaluate instance descriptions $s_i^{(t)}$ against semantic requirements. For spatial filtering,

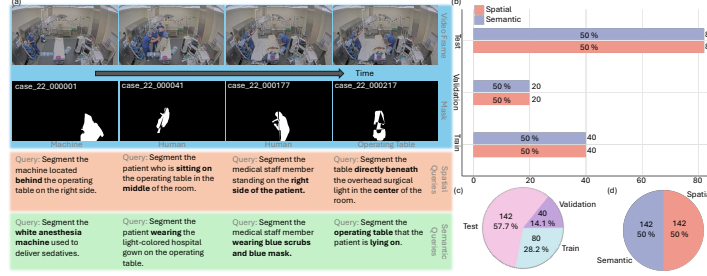


Fig. 2. Visualization of the OR reasoning segmentation in-house benchmark dataset. (a) Sample video frame sequence from case ID 22 demonstrating paired video frames, segmentation masks, and corresponding spatial/semantic implicit text queries. (b) Query type distribution per dataset split showing balanced representation. (c) Dataset split proportions across train, validation, and test datasets. (d) Overall equal distribution between spatial and semantic queries.

f_k prompts LLM to utilize the depth information and instance mask to assess relative positions between instances by generating relative codes. The final candidate set $\mathcal{C}^{(t)}$ is obtained through the intersection of all requirement-specific candidates *i.e.*, $\mathcal{C}^{(t)} = \bigcap_{k=1}^K \mathcal{C}_k^{(t)}$. The synthesis stage generates the final segmentation by combining instance masks from the candidate set. For each frame t , we construct the reasoning segmentation mask $M_Q^{(t)}$ as $M_Q^{(t)} = \bigcup_{i \in \mathcal{C}^{(t)}} m_i^{(t)}$, where $m_i^{(t)}$ represents the instance mask for candidate i .

Agent for OR Workflow Analysis Based on our RS method, we propose ORDiRS-Agent to enable OR workflow analysis through natural language interaction. Given a user query Q_{OR} about OR efficiency, ORDiRS-Agent first decomposes it into a set of RS sub-queries $\{Q_1, Q_2, \dots, Q_L\}$ through LLM-based planning. We implement the LLM-based planning process by a two-step chain-of-thought prompting, where the LLM first analyzes Q_{OR} to identify key efficiency aspects that require investigation, then generates targeted RS sub-queries for each aspect. After LLM-based planning, we construct the DT representation $\{\mathcal{J}^{(t)}\}_{t=1}^T$ across the entire OR video sequence and leverage ORDiRS to process each sub-query Q_l to generate corresponding RS masks $\{M_{Q_l}^{(t)}\}_{t=1}^T$. Then, we define an efficiency analysis function \mathcal{E}_l for each sub-query: $\mathcal{A}_l = \mathcal{E}_l(\{M_{Q_l}^{(t)}\}_{t=1}^T, \{\mathcal{J}^{(t)}\}_{t=1}^T)$, where \mathcal{A}_l represents the analysis results containing temporal patterns, statistical measures, and identified objects. The efficiency analysis function \mathcal{E}_l operates in two modes depending on the analysis requirements of Q_l . For semantic analysis, such as evaluating team coordination patterns or protocol compliance, \mathcal{E}_l directly employs the LLM to reason over the sequence of RS masks and DT representations. For quantitative analysis tasks, such as computing door opening frequencies or phase transition durations, \mathcal{E}_l

Table 1. The comparison of RS across different reasoning categories (semantic, spatial, mixed) on the in-house dataset. The values are reported as mean \pm standard deviation. The efficiency metric is the mean inference time per image in seconds. The upward arrow (\uparrow) signifies that higher values indicate better performance.

Methods	cIoU (\uparrow)			gIoU (\uparrow)			Time(s)
	Semantic	Spatial	Mixed	Semantic	Spatial	Mixed	
LISA-7B-SAM1 [5]	25.17 \pm 6.34	59.93 \pm 5.07	39.85 \pm 5.34	34.94 \pm 7.96	63.77 \pm 4.40	48.97 \pm 4.89	1.51 \pm 0.16
LISA-13B-SAM1 [5]	28.15 \pm 4.36	65.35 \pm 5.66	42.08 \pm 3.87	34.19 \pm 4.68	68.85 \pm 4.77	50.75 \pm 4.19	2.03 \pm 0.21
LISA-7B-SAM1 (ft) [5]	69.68 \pm 3.35	69.00 \pm 5.53	68.30 \pm 3.03	64.62 \pm 3.59	65.38 \pm 5.50	64.56 \pm 3.06	1.52 \pm 0.11
LISA-13B-SAM1 (ft) [5]	53.32 \pm 6.57	64.84 \pm 5.92	59.38 \pm 5.09	54.33 \pm 6.13	65.25 \pm 5.86	58.19 \pm 5.79	1.85 \pm 0.10
LISA-7B-SAM2 (ft) [5]	51.14 \pm 5.00	51.42 \pm 4.13	52.81 \pm 3.87	43.84 \pm 5.94	46.99 \pm 4.16	47.05 \pm 3.91	1.42 \pm 0.13
LISA-13B-SAM2 (ft) [5]	49.06 \pm 5.87	49.45 \pm 3.69	51.21 \pm 4.48	39.22 \pm 6.09	42.00 \pm 4.23	42.69 \pm 4.71	2.09 \pm 0.19
V* [15]	3.09 \pm 1.44	3.35 \pm 1.25	3.26 \pm 1.00	2.93 \pm 1.48	3.27 \pm 1.26	3.15 \pm 1.06	22.62 \pm 0.39
Ours	75.80\pm3.58	78.74\pm3.19	77.25\pm2.34	76.85\pm3.23	82.48\pm2.49	79.55\pm1.97	89.05 \pm 2.51

utilizes the LLM to generate Python code that processes the mask sequences and extracts relevant statistical metrics. Finally, ORDiRS-Agent aggregates all analysis results through LLM to obtain the final response to Q_{OR} by $\mathcal{R} = \text{LLM}(\{\mathcal{A}_i\}_{i=1}^L, Q_{OR})$, where \mathcal{R} represents the final response containing both textual explanations and supporting visual evidence in terms of the most relative RS mask.

3 Experiments

Implementation Details For LLMs used in our method, we utilize the GPT-4o. We leverage the cumulative intersection over the cumulative union (cIoU) and the average of all per-image intersection-over-unions (gIoU) as the metrics for RS [5]. We compare our method against LISA [5] (7B and 13B parameter versions) including both the original and fine-tuned LISA (on our dataset) [5] with their 7B and 13B parameter variants; and V* [15]. Each LISA variant was evaluated using two different SAMs, *i.e.*, SAM1 [4] and SAM2 [10].

Benchmark Dataset To evaluate RS performance on OR videos, we construct a benchmark dataset from in-house OR video recordings, collected under IRB approval with appropriate consent from all participants. The dataset encompasses recordings from four operating rooms (case IDs 22, 25, 26, and 28), specifically focusing on OR workflows centering on patients, anesthesiologists, anesthesia machines, and operating tables. Two experts manually annotated each frame with ground truth segmentation masks and the corresponding video-level implicit text query, as shown in Fig. 2(a). The annotation process incorporated both spatial and semantic queries, where spatial queries address location-based requirements and semantic queries focus on attribute-based or function-based identification. The dataset comprises 142 frames in total, divided into training (1 OR, 40 frames, 28.2%), validation (1 OR, 20 frames, 14.1%), and test sets (2 OR, 82 frames, 57.7%) on OR-level, as illustrated in Fig. 2(b,c). The training




Fig. 3. Qualitative comparison of RS results. Two representative cases are shown: a semantic RS task (top row) requiring identification of a patient with specific clothing attributes, and a spatial RS task (bottom row) involving positional understanding. White regions represent the segmentation masks, with ground truth shown in the left-most column.

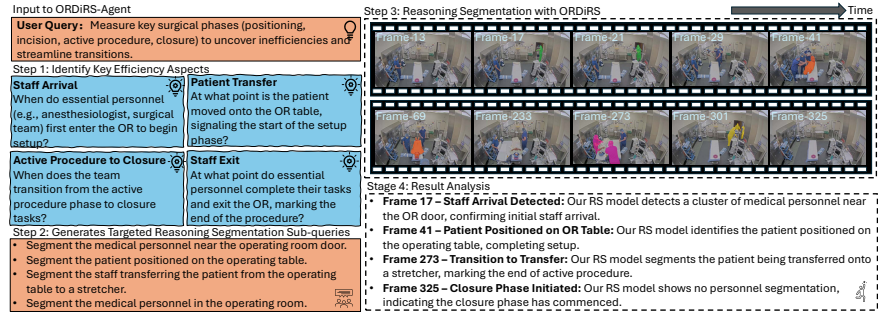
and validation sets are only used for the baseline methods *i.e.*, LISA [5], and we conduct comparisons of all methods on the test set. The query composition maintains an equal balance between spatial and semantic types across all splits (Fig. 2(d)), with each category representing 50% of the total queries. In addition to the in-house OR dataset, we extended a publicly available *MOVR* dataset [11] of 732 frames with the RS annotations and obtained *MOVR-Reason*, following the same annotation workflow.

Reasoning Segmentation Performance Table 1 presents the comparison of RS results and the computational efficiency. Our method demonstrates superior performance across all reasoning categories. Specifically, ORDiRS attains the highest cIoU scores of 75.80%, 78.74%, and 77.25% for semantic, spatial, and mixed reasoning tasks respectively, surpassing the previous best results from LISA-7B-SAM1 (ft) by margins of 6.12%, 9.74%, and 8.95%. Similar improvements are observed in gIoU metrics, where our method achieves 76.85%, 82.48%, and 79.55% across the three categories, representing substantial gains over existing approaches. However, it comes at the cost of increased computational overhead, with our method requiring 89.05 seconds per image compared to 1.51-22.62 seconds for baseline methods. This suggests that our method, without additional improvements to its real-time inference capabilities, is most suitable for offline workflow analysis tasks. Similar results can be observed from the *MOVR-Reason* dataset, in Table 2. Notably, while fine-tuned LISA variants (denoted by ‘ft’) show improved performance over their base counterparts, they still fall short of our tuning-free approach, highlighting the effectiveness of our “reason-retrieval-synthesize paradigm” in leveraging LLM’s inherent reasoning capabilities for OR environment understanding, as shown in Fig. 3.

Case Study on Operating Room Efficiency Analysis We evaluated ORDiRS-Agent through a case study analyzing surgical workflow efficiency in Fig. 4. Given the query “*Measure key surgical phases to uncover inefficiencies and streamline transitions,*” ORDiRS-Agent can detect important transition points. It demonstrates ORDiRS-Agent’s capability to decompose complex efficiency queries into manageable RS tasks, and provide comprehensive visual-temporal insights.

Table 2. The comparison of RS across different reasoning categories (semantic, spatial, mixed) on the *MVOR-Reason* dataset.

Methods	cIoU (\uparrow)			gIoU (\uparrow)			Time(s)
	Semantic	Spatial	Mixed	Semantic	Spatial	Mixed	
LISA-7B-SAM1 [5]	32.45 \pm 3.20	24.18 \pm 6.70	28.12 \pm 3.20	27.51 \pm 3.04	22.45 \pm 2.15	24.24 \pm 1.08	1.72 \pm 0.22
LISA-13B-SAM1 [5]	34.25 \pm 6.20	40.45 \pm 2.13	38.20 \pm 1.99	32.20 \pm 3.87	30.05 \pm 3.43	31.22 \pm 2.98	1.93 \pm 0.54
LISA-7B-SAM1 (ft) [5]	45.34 \pm 2.15	49.45 \pm 4.56	46.67 \pm 5.14	50.20 \pm 4.74	47.30 \pm 3.40	49.49 \pm 5.14	1.77 \pm 0.41
LISA-13B-SAM1 (ft) [5]	52.45 \pm 3.54	48.99 \pm 2.38	51.74 \pm 1.45	56.84 \pm 3.16	55.39 \pm 2.15	55.96 \pm 3.45	1.98 \pm 0.28
LISA-7B-SAM2 (ft) [5]	56.38 \pm 7.21	58.20 \pm 3.98	57.56 \pm 4.45	50.45 \pm 3.24	56.18 \pm 3.06	53.49 \pm 1.81	1.98 \pm 0.22
LISA-13B-SAM2 (ft) [5]	59.47 \pm 2.12	61.31 \pm 2.54	60.26 \pm 3.88	59.32 \pm 2.11	53.22 \pm 2.58	58.90 \pm 3.68	2.39 \pm 0.42
V* [15]	41.22 \pm 3.56	44.52 \pm 3.52	43.30 \pm 4.98	39.54 \pm 3.45	45.76 \pm 5.81	41.20 \pm 4.12	25.30 \pm 2.41
Ours	83.21\pm2.81	86.11\pm3.76	84.35\pm3.40	82.54\pm2.54	80.19\pm3.22	81.54\pm4.50	58.21 \pm 7.31

**Fig. 4.** A case study for the workflow of ORDiRS-Agent for analyzing operating room efficiency. The process begins with a user query about surgical phase transitions, followed by the identification of key efficiency aspects. It then generates targeted reasoning segmentation sub-queries (Step 2), performs reasoning segmentation using ORDiRS (Step 3), and concludes with result analysis (Step 4). The visualization demonstrates how ORDiRS-Agent tracks critical workflow events across frames, including staff arrival (Frame 17), patient positioning (Frame 41), transition to transfer (Frame 273), and closure phase initiation (Frame 325).

4 Conclusion

We present ORDiRS, an RS framework that utilizes DT representations to analyze OR videos. By separating perception from reasoning through the DT representation, ORDiRS eliminates the need for fine-tuning while ensuring robust performance across diverse OR environments. We also introduce ORDiRS-Agent, an analytical framework designed to evaluate OR efficiency using ORDiRS. Our results demonstrate that structured intermediate representations effectively bridge the gap between raw visual data (*e.g.*, video) and high-level reasoning. Future work could extend the DT representation framework to capture temporal dependencies and long-term patterns in surgical workflows, enabling deeper efficiency analyses. Moreover, the ‘reason-retrieve-synthesize’ paradigm could be adapted for other healthcare settings requiring complex visual reasoning, such as emer-

gency departments and intensive care units. Finally, accelerating the workflow with more efficient foundation models is also one promising direction.

Disclosure of Interests

The authors declare no conflicts of interest in relation to this work.

Acknowledgments

This work was supported in part by the JHU Amazon Initiative for Artificial Intelligence (AI2AI) fellowship program and NSF CAREER award (NSF Award No. 2239077).

References

1. Tianrun Chen, Chunan Yu, Jing Li, Jianqi Zhang, Lanyun Zhu, Deyi Ji, Yong Zhang, Ying Zang, Zejian Li, and Lingyun Sun. Reasoning3d-grounding and reasoning in 3d: Fine-grained zero-shot open-vocabulary 3d reasoning part segmentation via large vision-language models. *arXiv preprint arXiv:2405.19326*, 2024.
2. David M Friedman, Suzanne M Sokal, Yuchiao Chang, and David L Berger. Increasing operating room efficiency through parallel processing. *Annals of surgery*, 243(1):10–14, 2006.
3. Rayan Ebnali Harari, Roger D Dias, Lauren R Kennedy-Metz, Giovanna Varni, Matthew Gombolay, Steven Yule, Eduardo Salas, and Marco A Zenati. Deep learning analysis of surgical video recordings to assess nontechnical skills. *JAMA Network Open*, 7(7):e2422520–e2422520, 2024.
4. Alexander Kirillov et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023.
5. Xin Lai, Zhuotao Tian, Yukang Chen, et al. Lisa: Reasoning segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9579–9589, 2024.
6. Haotian Liu, Chunyuan Li, Qingyang Wu, et al. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024.
7. Matthias Minderer, Alexey Gritsenko, and Neil Houlsby. Scaling open-vocabulary object detection. *Advances in Neural Information Processing Systems*, 36, 2024.
8. Niall O’Mahony, Sean Campbell, Anderson Carvalho, Suman Harapanahalli, Gustavo Velasco Hernandez, Lenka Krpalkova, Daniel Riordan, and Joseph Walsh. Deep learning vs. traditional computer vision. In *Advances in Computer Vision: Proceedings of the 2019 Computer Vision Conference (CVC), Volume 1*, pages 128–144. Springer, 2020.
9. Arnaud Pasquer, Simon Ducarroz, Jean Christophe Lifante, Sarah Skinner, Gilles Poncet, and Antoine Duclos. Operating room organization and surgical performance: a systematic review. *Patient Safety in Surgery*, 18(1):5, 2024.
10. Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.
11. Vinkle Srivastav, Thibaut Issenhuth, Abdolrahim Kadkhodamohammadi, Michel de Mathelin, Afshin Gangi, and Nicolas Padoy. Mvor: A multi-view rgb-d operating room dataset for 2d and 3d human pose estimation. *arXiv preprint arXiv:1808.08180*, 2018.

12. Adriana Vladu, Timea Claudia Ghitea, Lucia Georgeta Daina, Dorel Petru Țîrț, and Mădălina Diana Daina. Enhancing operating room efficiency: The impact of computational algorithms on surgical scheduling and team dynamics. In *Healthcare*, volume 12, page 1906. MDPI, 2024.
13. Kushala VM, Harikrishna Warriar, Yogesh Gupta, et al. Fine tuning llm for enterprise: Practical guidelines and recommendations. *arXiv preprint arXiv:2404.10779*, 2024.
14. Junchi Wang and Lei Ke. Llm-seg: Bridging image segmentation and large language model reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1765–1774, 2024.
15. Penghao Wu and Saining Xie. V*: Guided visual search as a core mechanism in multimodal llms. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13084–13094, 2024.
16. Yuchen Xia, Jiho Kim, Yuhan Chen, Haojie Ye, Souvik Kundu, Cong Callie Hao, and Nishil Talati. Understanding the performance and estimating the cost of llm fine-tuning. In *2024 IEEE International Symposium on Workload Characterization (IISWC)*, pages 210–223. IEEE, 2024.
17. Zhuofan Xia, Dongchen Han, Yizeng Han, et al. Gsva: Generalized segmentation via multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3858–3869, 2024.
18. Cilin Yan, Haochen Wang, Shilin Yan, et al. Visa: Reasoning video object segmentation via large language models. *arXiv preprint arXiv:2407.11325*, 2024.
19. Lihe Yang, Bingyi Kang, Zilong Huang, et al. Depth anything v2. *arXiv preprint arXiv:2406.09414*, 2024.
20. Senqiao Yang, Tianyuan Qu, Xin Lai, Zhuotao Tian, Bohao Peng, Shu Liu, and Jiaya Jia. An improved baseline for reasoning segmentation with large language model. *CoRR*, 2023.