

# A Guideline-Aware AI Agent for Zero-Shot Target Volume Auto-Delineation

Yoon Jo Kim<sup>1\*</sup>, Wonyoung Cho<sup>1\*</sup>, Jongmin Lee<sup>1</sup>, Han Joo Chae<sup>1</sup>, Hyunki Park<sup>2</sup>, Sang Hoon Seo<sup>2</sup>, Jae Myung Noh<sup>2</sup>, Kyungmi Yang<sup>2</sup>, Dongryul Oh<sup>2\*\*</sup>, and Jin Sung Kim<sup>1\*\*</sup>

<sup>1</sup> Oncosoft Inc., Seoul, South Korea

<sup>2</sup> Department of Radiation Oncology, Samsung Medical Center, Sungkyunkwan University School of Medicine, Seoul, South Korea  
dongryul.oh@samsung.com, jin@oncosoft.io

**Abstract.** Delineating the clinical target volume (CTV) in radiotherapy involves complex margins constrained by tumor location and anatomical barriers. While deep learning models automate this process, their rigid reliance on expert-annotated data requires costly retraining whenever clinical guidelines update. To overcome this limitation, we introduce OncoAgent, a novel guideline-aware AI agent framework that seamlessly converts textual clinical guidelines into three-dimensional target contours without any target volume annotation. Evaluated on esophageal cancer cases, the agent achieves a Dice similarity coefficient of 0.842 for the CTV and 0.880 for the planning target volume, demonstrating performance highly comparable to a fully supervised nnU-Net baseline. Notably, in a blinded clinical evaluation, physicians strongly preferred OncoAgent over the supervised baseline, rating it higher in guideline compliance, modification effort, and clinical acceptability. Furthermore, without any retraining, the framework generalizes to alternative esophageal guidelines and shows preliminary extensibility to other anatomical sites (e.g., prostate). Beyond mere volumetric overlap, our agent-based paradigm offers near-instantaneous adaptability to alternative guidelines, providing a scalable and transparent pathway toward interpretability in radiotherapy treatment planning.

**Keywords:** Clinical Target Volume · Esophageal Cancer · AI Agent · Radiotherapy Guidelines · Auto-Delineation

## 1 Introduction

Radiation therapy requires precise delineation of the clinical target volume (CTV) to encompass microscopic disease while strictly sparing organs at risk (OARs). Such delineation frequently involves complex, non-isotropic margins determined by tumor location and surrounding anatomical barriers. As a result, this manual

\* Equal contribution

\*\* Co-corresponding authors

process remains labor-intensive and highly prone to substantial inter-observer variability [11, 1]. While deep learning (DL) models like nnU-Net [6] have automated some aspects of medical segmentation, they remain “black boxes” that require extensive, expert-annotated datasets. Moreover, these models are inherently rigid: any update to clinical guidelines requires the costly re-annotation of entire datasets and subsequent model retraining.

Recent advances in large language models (LLMs) have given rise to *AI agents*: systems that combine an LLM’s reasoning capabilities with external tools to perform complex, multi-step tasks [16, 13]. An AI agent can interpret a clinical guideline written in natural language, decompose it into an executable plan, and invoke specialized tools—such as OAR segmentation models and morphological operators—to autonomously construct patient-specific target volumes.

This paradigm offers three key advantages over data-driven DL: (1) *no requirement for expert annotations*, since the delineation logic is directly derived from the guideline text rather than learned from examples; (2) *instant adaptability* to evolving guidelines or institutional variations, as merely updating the textual input is sufficient to reflect new clinical protocols; and (3) *interpretability*, as the agent produces a human-readable plan before execution that clinicians can review and approve.

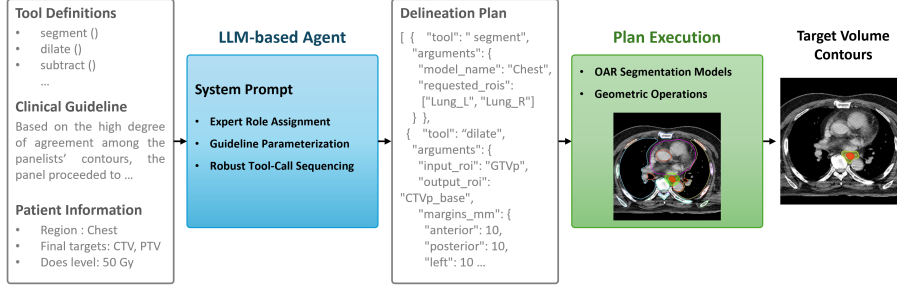
In this work, we present OncoAgent, an AI agent framework for universal, guideline-aware target volume auto-delineation and evaluate it on esophageal cancer patients. Our contributions are threefold:

1. We propose a general, two-phase agentic architecture—*planning* and *execution*—that converts free-text clinical guidelines into three-dimensional target volumes, using pre-trained OAR segmentation models and geometric operation tools without any target volume annotation.
2. Using a mid-thoracic esophageal cancer dataset, we demonstrate that OncoAgent achieves performance highly comparable to a fully supervised nnU-Net baseline, while inherently ensuring superior anatomical safety and strict guideline compliance.
3. We validate the framework’s broad applicability through cross-guideline transferability experiments, demonstrating that OncoAgent adapts to alternative esophageal protocols and shows preliminary extensibility to other anatomical sites (e.g., prostate).

## 2 Method

### 2.1 Overview

OncoAgent automatically converts textual clinical guidelines into 3D target volumes in a zero-shot manner—without any CTV or planning target volume (PTV) annotation (Fig. 1). Given tool definitions, a clinical guideline, and patient information, a large language model (LLM) interprets the raw text guideline to formulate a delineation plan—an ordered sequence of tool calls expressed as a JSON array. The plan is validated against the tool-call schema [3]; if violations



**Fig. 1.** Overview of the proposed OncoAgent framework for zero-shot target volume auto-delineation.

are detected, the LLM iteratively self-refines [9] until conformance. A deterministic execution engine then dispatches each call on the patient’s CT and gross tumor volume (GTV) contours, invoking pre-trained OAR segmentation models and geometric operations to produce the final target volume.

## 2.2 System Prompt Strategies

The system prompt defines the core logic of the agent, structuring raw text guidelines into deterministic, actionable tool-call sequences. It encodes three complementary prompting strategies (Fig. 2) that enable OncoAgent to operate effectively in a zero-shot setting.

*Clinician Role Assignment.* Following the role prompting paradigm [16], the system prompt assigns the LLM the persona of an expert radiation oncology agent. To ensure deterministic, guideline-compliant delineation, the prompt encodes the standard target volume definitions [5, 15] as fixed computational templates following the program-aided reasoning paradigm [2]:

$$\text{CTV} = (\text{GTV} \oplus \mathbf{m}_{\text{ctv}}) \setminus \text{OARs}, \quad \text{PTV} = \text{CTV} \oplus \mathbf{m}_{\text{ptv}}, \quad (1)$$

where  $\mathbf{m}$  denotes a direction-specific margin vector. These domain-specific templates guide the agent to produce delineation plans that are explicitly aligned with clinical guidelines.

*Guideline Parameterization.* Clinical guidelines often exhibit terminological variation and present margin values as ranges. The system prompt provides example OAR aliases and a list of segmentable structures, enabling the LLM to map guideline terms to model-specific OAR names (e.g., “Lung”  $\rightarrow$  Lung\_L, Lung\_R). When guidelines specify ranges (e.g., “5–10 mm”), OncoAgent resolves them using patient-specific context (e.g., dose level, physician preferences).

---

**System Prompt** (simplified)

---

**Clinician Role Assignment**

- “You are an expert radiation oncology guideline-to-implementation agent.”
- Mandatory template:  $CTV = (GTV \oplus \mathbf{m}_{ctv}) \setminus OARs$ ;  $PTV = CTV \oplus \mathbf{m}_{ptv}$

**Guideline Parameterization**

- *OAR alias examples:*  
Lung  $\rightarrow$  Lung\_L, Lung\_R; Kidney  $\rightarrow$  Kidney\_L, Kidney\_R; VBs  $\rightarrow$  VB\_whole
- Multi-OAR aliases: apply operations per component, then union
- Resolve margin ranges using patient context (e.g., dose level, physician preferences), producing patient-specific delineation plans

**Robust Tool-Call Sequencing**

1. Query existing ROIs and segmentation models
  2. **segment**: delineate required OARs
  3. **union**: OARs  $\rightarrow$  unified exclusion mask
  4. **dilate**: GTV  $\rightarrow$  CTV base
  5. **subtract**: CTV base  $-$  exclusion mask  $\rightarrow$  final CTV
  6. **dilate**: CTV  $\rightarrow$  PTV
  7. Validate against schema; self-refine if violations found
- 

**Fig. 2.** Simplified system prompt. The three strategies are detailed in Sec. 2.2.

*Robust Tool-Call Sequencing.* Adapted from algorithmic prompting [19], the system prompt specifies a step-by-step execution order of segmentation and geometric tools, such as **segment**, **dilate**, and **subtract**, to avoid delineation errors. Additionally, consistent naming conventions ensure correct ROI referencing throughout the generated plan. If the generated plan is invalid, the LLM iteratively self-refines the plan until it works.

### 2.3 Plan Execution

The execution engine receives the validated tool-call sequence and dispatches each call sequentially, passing the output of each step as input to subsequent ones. OncoAgent has access to two classes of tools (Fig. 2): a *segmentation tool* that invokes pre-trained deep learning models to delineate anatomical structures on the patient CT, and *geometric operation tools* that apply direction-specific morphological and Boolean operations. After all planned calls have been executed, post-processing—contour smoothing and hole-filling—is applied to produce the final target volume.

### 3 Results

#### 3.1 Experimental Setup

*Dataset.* We retrospectively collected planning CT scans from 40 patients with mid-thoracic esophageal cancer, provided by Samsung Medical Center under the approval of the Data Review Board (DRB No. D2025-0006) and the Institutional Review Board (IRB No. SMC 2024-11-122). For each case, an experienced radiation oncologist delineated ground truth (GT) contours for the GTV, CTV, and PTV in accordance with the IJROBP expert consensus guideline [17]. All CT images were anonymized and exported from the Eclipse treatment planning system.

*Comparative Methods.* We compared OncoAgent against three supervised DL baselines: (1) nnU-Net [6], a robust, self-configuring medical image segmentation framework; (2) DDAU-Net [18], a state-of-the-art dual-attention model for esophageal target volume segmentation; and (3) nnU-Net (GTV Prior), a strong baseline utilizing both CT images and GT GTV masks as a 2-channel input for GTV-guided segmentation [8]. The DL models were trained on 32 cases, and both the DL models and OncoAgent were evaluated on the remaining 8 test cases.

*Evaluation Metrics and Clinical Assessment.* Quantitative performance was evaluated using Dice Similarity Coefficient (DSC), Mean Surface Distance (MSD), Sensitivity, and Precision, with statistical significance determined by a paired t-test ( $p < 0.05$ ). To assess planning accuracy, we additionally used *Tool Call F1*, the set-level F1 score between the agent-generated and a manually authored reference tool-call sequence.

For clinical assessment, two senior radiation oncologists (17 and 20 years’ experience) blindly rated the 8 test cases generated by nnU-Net (GTV Prior) and OncoAgent. Using a 5-point Likert scale, they evaluated Guideline Compliance, Modification Effort, and Clinical Acceptability, respectively. Qualitative analysis assessed anatomical violations via visual comparison.

*Implementation Details.* The system integrates GPT-5.2 [12] (`temperature=1.0`, `reasoning_effort=none`) as its reasoning engine, utilizing a pre-trained 3D U-Net (trained on 2,870 CT/MR volumes, Dice 0.80–0.98) for OAR segmentation and the VTK library [14] for geometric operations. This end-to-end workflow achieved an average execution time of 1.74 minutes per patient, requiring a mean of 1.13 LLM inference calls per case ( $\sim 13\%$  triggering self-refinement) to generate a complete set of target contours. The source code is publicly available at <https://github.com/Oncosoft-Research/OncoAgent>.

#### 3.2 Quantitative Evaluation

The quantitative performance of OncoAgent is presented in Table 1. While baselines relying solely on CT images (DDAU-Net, nnU-Net) struggled with invisible,

**Table 1.** Quantitative evaluation of target volume segmentation (mean  $\pm$  SD). MSD is in mm. Bold: best or no significant difference ( $p \geq 0.05$ ) from best.

| Class | Method              | DSC $\uparrow$                    | MSD $\downarrow$                | Sens. $\uparrow$                  | Prec. $\uparrow$                  |
|-------|---------------------|-----------------------------------|---------------------------------|-----------------------------------|-----------------------------------|
| CTV   | DDAU-Net            | 0.463 $\pm$ 0.107                 | 8.81 $\pm$ 3.88                 | 0.345 $\pm$ 0.103                 | 0.762 $\pm$ 0.143                 |
|       | nnU-Net             | 0.627 $\pm$ 0.130                 | 7.82 $\pm$ 6.90                 | 0.683 $\pm$ 0.071                 | 0.594 $\pm$ 0.183                 |
|       | nnU-Net (GTV Prior) | <b>0.862<math>\pm</math>0.026</b> | <b>0.95<math>\pm</math>0.22</b> | <b>0.845<math>\pm</math>0.053</b> | <b>0.882<math>\pm</math>0.028</b> |
|       | OncoAgent (Ours)    | <b>0.842<math>\pm</math>0.026</b> | <b>1.06<math>\pm</math>0.15</b> | <b>0.884<math>\pm</math>0.029</b> | 0.807 $\pm$ 0.053                 |
| PTV   | DDAU-Net            | 0.517 $\pm$ 0.088                 | 7.95 $\pm$ 3.46                 | 0.398 $\pm$ 0.098                 | 0.785 $\pm$ 0.119                 |
|       | nnU-Net             | 0.659 $\pm$ 0.118                 | 8.77 $\pm$ 7.16                 | 0.723 $\pm$ 0.059                 | 0.620 $\pm$ 0.171                 |
|       | nnU-Net (GTV Prior) | <b>0.893<math>\pm</math>0.021</b> | <b>1.02<math>\pm</math>0.25</b> | 0.873 $\pm$ 0.045                 | <b>0.916<math>\pm</math>0.020</b> |
|       | OncoAgent (Ours)    | <b>0.880<math>\pm</math>0.015</b> | <b>1.12<math>\pm</math>0.14</b> | <b>0.922<math>\pm</math>0.027</b> | 0.843 $\pm$ 0.033                 |

guideline-derived CTV boundaries, incorporating GTV prior knowledge—either via supervised multi-channel inputs (nnU-Net with GTV Prior) or zero-shot guideline reasoning (OncoAgent)—yielded substantial performance gains.

Notably, OncoAgent demonstrated performance highly comparable to the fully supervised state-of-the-art model, nnU-Net (GTV Prior), despite achieving this entirely zero-shot without the annotated training cases required by the DL baselines. Both models maintained high overall performance (DSC  $\geq 0.84$ ) and sub-voxel distance errors (MSD  $\approx 1.0$  mm). As indicated by the bold values in Table 1, there were no statistically significant differences ( $p \geq 0.05$ ) between the two models across primary metrics like DSC and MSD. While the supervised model showed slightly higher Precision, OncoAgent achieved comparable or superior Sensitivity, yielding similar overall performance.


### 3.3 Clinical Assessment and Qualitative Analysis

The blinded physician survey is summarized in Table 2, showing a strong clinical preference for OncoAgent. OncoAgent outperformed the nnU-Net (GTV Prior) model across all criteria, especially in Guideline Compliance (4.06 vs. 3.56) and Modification Effort (4.12 vs. 3.44).

Visual comparisons between the manual GT, nnU-Net (GTV Prior), and OncoAgent results are illustrated in Fig. 3. In the axial view (yellow arrows), the DL model often invades adjacent OARs (e.g., lung) as it lacks explicit end-points. On the other hand, OncoAgent strictly applies explicit boolean exclusion rules thereby avoiding anatomical violations; although this rigid subtraction mechanism led to a “slightly over-eroded” appearance near the heart compared to the GT, it successfully avoided critical OAR overlap and resulted in a safer target volume. Furthermore, in the sagittal view, OncoAgent accurately matched the GT by strictly adhering to the specified longitudinal ( $z$ -axis) margins, while the DL model arbitrarily over-expanded by approximately 6 mm superiorly.

**Table 2.** Blinded physician survey (5-point Likert scale). **Bold:** highest.

| Evaluation Criteria    | Method                  | Mean $\pm$ SD                     | Ratings $\geq 4$ (%) |
|------------------------|-------------------------|-----------------------------------|----------------------|
| Guideline Compliance   | nnU-Net (GTV Prior)     | 3.56 $\pm$ 0.51                   | 56.2%                |
|                        | <b>OncoAgent (Ours)</b> | <b>4.06 <math>\pm</math> 0.68</b> | <b>81.2%</b>         |
| Modification Effort    | nnU-Net (GTV Prior)     | 3.44 $\pm$ 0.63                   | 37.5%                |
|                        | <b>OncoAgent (Ours)</b> | <b>4.12 <math>\pm</math> 0.89</b> | <b>68.8%</b>         |
| Clinical Acceptability | nnU-Net (GTV Prior)     | 3.38 $\pm$ 0.72                   | 50.0%                |
|                        | <b>OncoAgent (Ours)</b> | <b>3.81 <math>\pm</math> 0.75</b> | <b>75.0%</b>         |

**Fig. 3.** Qualitative comparison of nnU-Net (GTV Prior) and OncoAgent in axial and sagittal views. Yellow arrows highlight areas of interest.

### 3.4 Zero-Shot Generalization and Cross-Guideline Adaptability

Unlike DL models requiring retraining for guideline updates, OncoAgent adapts immediately. We evaluated it on alternative esophageal guidelines (CROSS Trial [4], JASTRO 2024 [7]) and extended it to prostate cancer (RTOG 0126 [10]) without retraining. As shown in Table 3, OncoAgent attained high Tool Call F1 scores (0.73–1.00) across esophageal guidelines, demonstrating cross-guideline adaptability without task-specific fine-tuning. Extending beyond esophageal cancer to prostate, it still reached 0.64 despite unfamiliar terminology (e.g., “SV\_prox1cm”), suggesting preliminary cross-site generalization.

## 4 Discussion and Conclusion

### 4.1 Discussion

We have presented OncoAgent, the first AI agent framework for guideline-aware target volume auto-delineation, evaluated on mid-thoracic esophageal cancer. This approach shifts the focus from *learning CTV from annotated data* to *reasoning about CTV from clinical guidelines*, yielding three distinct advantages.

**Table 3.** Zero-shot transferability across different clinical guidelines and a novel site (prostate).

| Target Site & Guideline                               | Tool Call F1 $\uparrow$ |
|-------------------------------------------------------|-------------------------|
| Esophagus (IJROBP Expert Consensus <sup>†</sup> [17]) | 1.00                    |
| Esophagus (JASTRO 2024 Guideline [7])                 | 0.78                    |
| Esophagus (CROSS Trial [4])                           | 0.73                    |
| Prostate (RTOG 0126 [10])                             | 0.64                    |

<sup>†</sup>Primary guideline used in the main evaluation.

*Annotation-Free Deployment.* OncoAgent requires zero CTV-labeled training data, relying instead on a clinical guideline document and pre-trained OAR segmentation models. Despite this zero-shot approach, it achieved performance comparable to the fully supervised nnU-Net (GTV Prior), with no statistically significant differences in primary metrics such as DSC and MSD. Furthermore, in blinded clinical evaluations, physicians demonstrated a strong preference for OncoAgent over the supervised baseline, rating it higher in both Guideline Compliance (4.06 vs. 3.56) and Modification Effort (4.12 vs. 3.44).

*Instant Guideline Adaptability.* When guidelines evolve or institutional protocols vary, OncoAgent adapts through simple document updates, bypassing the need for costly re-annotation or model retraining. We demonstrated this by transitioning between the IJROBP expert consensus guideline [17], CROSS [4], and the JASTRO 2024 guideline [7], producing appropriately distinct CTV contours without any additional training. This capability directly addresses a critical challenge in clinical practice, where conventional DL models often lose their clinical applicability with each guideline update.

*Interpretability and Auditability.* OncoAgent’s intermediate plan—specifying margins, excluded OARs, and execution steps in human-readable form—enables clinicians to review and approve the delineation logic before contour generation. This transparency is critical for regulatory compliance and clinical trust, and is fundamentally absent in end-to-end DL approaches.

*Limitations.* First, OncoAgent’s performance depends on the quality of the underlying OAR segmentation model; however, recent advancements in robust auto-segmentation substantially mitigate this concern. Second, as with any LLM-based system, OncoAgent is susceptible to semantic misinterpretation and hallucination—for instance, misapplying a guideline clause to an inappropriate anatomical context—and cannot capture implicit clinical knowledge absent from the guideline text. Iterative self-refinement rejects structurally invalid plans before execution, as empirically supported by high Tool Call F1 scores (0.73–1.00) across diverse guidelines; nevertheless, clinician review of the interpretable execution plan remains essential before clinical use. Third, while OncoAgent’s strict guideline compliance yields geometrically precise contours, this may lack

the pragmatic safety margins that physicians manually draw to accommodate esophageal motion from respiration, swallowing, and cardiac pulsation; future work should integrate adaptive margin strategies to reconcile geometric precision with motion robustness. Fourth, this study is limited to 40 patients from a single institution; multi-center validation, extension to other tumor sites (e.g., head and neck, cervical cancer), and clinician feedback loops are important future directions.

## 4.2 Conclusion

OncoAgent demonstrates that LLM-based reasoning, combined with pre-trained segmentation tools and geometric operation tools, can achieve clinically competitive target volume delineation without any expert-annotated training data. Notably, in a blinded physician survey, OncoAgent demonstrated strong clinical preference over the supervised baseline, achieving higher ratings in Guideline Compliance, Modification Effort, and Clinical Acceptability. This framework bridges the gap between text-based clinical guidelines and automated target delineation, offering a transparent and adaptable alternative to conventional data-driven deep learning.

**Acknowledgments.** The authors thank Samsung Medical Center for providing the clinical data for this study. This study was approved by the Institutional Review Board of Samsung Medical Center (DRB No. D2025-0006, IRB No. SMC 2024-11-122).

**Disclosure of Interests.** Y. Kim, W. Cho, J. Lee, and H. Chae are employees of Oncosoft Inc. J. Kim is a CEO of Oncosoft Inc. The remaining authors have no competing interests to declare.

## References

1. Chang, X., Deng, W., Wang, X., Zhou, Z., Yang, J., Guo, W., Liu, M., Qi, X., Li, L., Zhang, K., et al.: Interobserver variability in target volume delineation in definitive radiotherapy for thoracic esophageal cancer: a multi-center study from china. *Radiat. Oncol.* **16**, 102 (2021)
2. Gao, L., Madaan, A., Zhou, S., Alon, U., Liu, P., Yang, Y., Callan, J., Neubig, G.: PAL: Program-aided language models. In: *Proc. Int. Conf. Mach. Learn.* vol. 202, pp. 10764–10799 (2023)
3. Geng, S., Josifoski, M., Peber, M., Cooper, H., Moskal, M., Jenkins, S., Berman, J., Ranchin, N., West, R., Horvitz, E., Nori, H.: Jsonschemabench: A rigorous benchmark of structured outputs for language models. *arXiv preprint arXiv:2501.10868* (2025)
4. van Hagen, P., Hulshof, M., van Lanschot, J., Steyerberg, E., van Berge Henegouwen, M., Wijnhoven, B., Richel, D., Nieuwenhuijzen, G., Hospers, G., Bonenkamp, J., et al.: Preoperative chemoradiotherapy for esophageal or junctional cancer. *N. Engl. J. Med.* **366**(22), 2074–2084 (2012)

5. International Commission on Radiation Units and Measurements: ICRU Report 62: Prescribing, Recording and Reporting Photon Beam Therapy (Supplement to ICRU Report 50). International Commission on Radiation Units and Measurements (1999)
6. Isensee, F., Jaeger, P., Kohl, S., Petersen, J., Maier-Hein, K.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nat. Methods* **18**(2), 203–211 (2021)
7. Japanese Society for Radiation Oncology (JASTRO): Guidelines for radiotherapy planning of esophageal cancer (2024)
8. Kihara, S., Koike, Y., Takegawa, H., Anetai, Y., Nakamura, S., Tanigawa, N., Koizumi, M.: Clinical target volume segmentation based on gross tumor volume using deep learning for head and neck cancer treatment. *Med. Dosim.* **48**(1), 20–24 (2023)
9. Madaan, A., Tandon, N., Gupta, P., Hallinan, S., Gao, L., Wiegrefe, S., Alon, U., Dziri, N., Prabhunoye, S., Yang, Y., Gupta, S., Majumder, B., Hermann, K., Welleck, S., Yazdanbakhsh, A., Clark, P.: Self-refine: Iterative refinement with self-feedback. *Adv. Neural Inf. Process. Syst.* **36** (2023)
10. Michalski, J., et al.: Effect of standard vs dose-escalated radiation therapy for patients with intermediate-risk prostate cancer: the nrg oncology rtog 0126 randomized clinical trial. *JAMA Oncol.* **4**(6), e180039 (2018)
11. Njeh, C.: Tumor delineation: the weakest link in the search for accuracy in radiotherapy. *J. Med. Phys.* **33**(4), 136–140 (2008)
12. OpenAI: GPT-5 technical report. Tech. rep., OpenAI (2025)
13. Schick, T., Dwivedi-Yu, J., Dessì, R., Raileanu, R., Lomeli, M., Zettlemoyer, L., Cancedda, N., Scialom, T.: Toolformer: language models can teach themselves to use tools. In: *Adv. Neural Inf. Process. Syst.* (2023)
14. Schroeder, W., Martin, K., Lorensen, B.: *The Visualization Toolkit: An Object-Oriented Approach to 3D Graphics*. Kitware (2006)
15. Shusharina, N., Söderberg, J., Edmunds, D., Löfman, F., Shih, H., Bortfeld, T.: Automated delineation of the clinical target volume using anatomically constrained 3d expansion of the gross tumor volume. *Radiother. Oncol.* **146**, 37–43 (2020)
16. Wang, L., Ma, C., Feng, X., Zhang, Z., Yang, H., Zhang, J., Chen, Z., Tang, J., Chen, X., Lin, Y., Zhao, W., Wei, Z., Wen, J.: A survey on large language model based autonomous agents. *Front. Comput. Sci.* **18**(6), 186345 (2024)
17. Wu, A., Bosch, W., Chang, D., Hong, T., Jabbour, S., Kleinberg, L., Mamon, H., Thomas, C., Goodman, K.: Expert consensus contouring guidelines for intensity modulated radiation therapy in esophageal and gastroesophageal junction cancer. *Int. J. Radiat. Oncol. Biol. Phys.* **92**(4), 911–920 (2015)
18. Yousefi, S., Sokooti, H., Elmahdy, M., Lips, I., Shalmani, M., Zinkstok, R., Dankers, F., Staring, M.: Esophageal tumor segmentation in ct images using a dilated dense attention unet (ddaunet). *IEEE Access* **9**, 68372–68383 (2021)
19. Zhou, H., Nova, A., Larochelle, H., Courville, A., Neyshabur, B., Sedghi, H.: Teaching algorithmic reasoning via in-context learning. *arXiv preprint arXiv:2211.09066* (2023)