



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

LinGuinE: Longitudinal Guidance Estimation for Volumetric Tumour Segmentation

Nadine Garibli*, Bence Csiba, Konstantinos Sidiropoulos^[0000–0003–1014–967X],
Yi Wei, and Mayank Patwari^[0000–0001–5197–5260]

AstraZeneca, 1 Francis Crick Avenue, Cambridge, United Kingdom

*Corresponding author: nadine.garibli@astrazeneca.com

Abstract. Longitudinal volumetric tumour segmentation is critical for radiotherapy planning and response assessment, yet this problem is under-explored and most methods produce single-timepoint semantic masks, lack lesion correspondence, and offer limited radiologist control. We introduce LinGuinE (**Longitudinal Guidance Estimation**), a PyTorch framework that combines image registration and guided segmentation to deliver lesion-level tracking and volumetric masks across all scans in a longitudinal study from a single radiologist prompt. LinGuinE is temporally direction agnostic, requires no training on longitudinal data, and allows any registration and semi-automatic segmentation algorithm to be repurposed for the task. We evaluate various combinations of registration and segmentation algorithms within the framework. LinGuinE achieves state-of-the-art segmentation and tracking performance across four datasets with a total of 456 longitudinal studies. Tumour segmentation performance shows minimal degradation with increasing temporal separation. We conduct ablation studies to determine the impact of auto-regression, pathology specific finetuning, and the use of real radiologist prompts. We release our code (<https://github.com/n-e-garibli/LinGuinE>) and substantial public benchmarking for longitudinal segmentation, facilitating future research.

Keywords: Longitudinal Study · Tumour Tracking · Volumetric Segmentation

1 Introduction

Volumetric segmentation of gross tumour volumes is an important first step in radiotherapy and interventional planning [21]. Volumetric assessments also show promise for evaluating response to oncological therapy [20]. Response is normally studied over multiple scans, acquired at different timepoints, of the same patient.

Presently, deep learning methods achieve strong automatic segmentation of tumour volumes [17], however most approaches, including those that incorporate multiple timepoints [19, 11, 14], are designed to produce semantic segmentations of a single volume at a single timepoint. Such fully automatic methods do not perform lesion tracking, which is essential for measuring treatment response and

assessing disease progression. Moreover, fully automatic methods do not allow radiologist input, such as choosing lesions of focus, which reduces the chance of clinical adoption [3]. Promptable semi-automatic segmentation, popularised by Segment Anything [12] and adapted to medical imaging to incorporate radiologist guidance (e.g., points or bounding boxes) [16], typically improves performance over fully automatic pipelines and supports segmentation of individual tumour volumes. However, semi-automatic methods also operate at single time-points, providing limited support for the analysis of an entire longitudinal study.

There exist several methods that track tumours in longitudinal studies by predicting their locations but do not produce segmentations [4, 25]. To the best of our knowledge, only three methods handle simultaneous tracking and segmentation. Hering et al. [8] use image registration to define a region of interest (ROI) around each tumour based on a radiologist-provided location in a previous time-point and invokes an nnUNet to generate segmentations on follow up scans [9]. Therefore, radiologist guidance is limited to an ROI in the follow ups, capping segmentation performance at that of the automatic nnUNet. LesionLocator [18] addresses this by introducing an end-to-end trainable model, combining registration and promptable segmentation for tumours on CT. However, if performance on a dataset is limited, the model cannot be finetuned without a relevant labelled longitudinal dataset with tracked lesions, which are difficult to obtain. Hein et al. [7] perform automatic segmentation and assign tumour correspondence with a CNN. Hence, this method cannot leverage radiologist inputs and requires training the CNN on longitudinal data. Additionally, existing methods are evaluated in the context of a single modality and cancer type, and on private datasets with a limited number of timepoints per study [7, 8, 18], which prohibits in-depth analysis of temporal performance degradation.

To address these limitations and provide a well-performing out-of-the-box solution, we present LinGuinE (**L**ongitudinal **G**uidance **E**stimation). LinGuinE is a customizable framework for the tracking and segmentation of tumours in longitudinal studies from a prompt provided at a single timepoint. LinGuinE combines prompt propagation using image registration with guided segmentation, and is the first framework that allows arbitrary registration and semi-automatic segmentation algorithm (with no requirement for training on longitudinal data) to be directly repurposed. We demonstrate that LinGuinE:

1. Achieves state-of-the-art results from minimal radiologist input.
2. Is the first temporally direction agnostic framework, allowing radiologists to provide prompts at any timepoint, facilitating retrospective analysis.
3. Outperforms existing methods by combining public algorithms, making LinGuinE a state-of-the-art solution in the absence of longitudinal training data.
4. Is more robust to temporal degradation in performance than other methods.

To facilitate further research, our contribution includes a PyTorch package (<https://github.com/n-e-garibli/LinGuinE>) and the first benchmarking of longitudinal algorithms on public data, multiple cancer types, and studies with more than three timepoints. The segmentation masks generated with various configurations of LinGuinE will be made publicly available.

2 Methods

2.1 The LinGuinE Framework

Let S denote a set of volumetric scans from the same patient. Let $I_i \in S$ be a scan at timepoint i and P_i^k be the 3D coordinate center of a tumour k on I_i . LinGuinE produces a segmentation for the same tumour on every scan in S ,

$$\text{LinGuinE}(I_i, P_i^k) \rightarrow \{M_j(k) \mid I_j \in S\} \quad (1)$$

where we call I_i the source scan, I_j the destination scan at timepoint j , and $M_j(k)$ the segmentation mask of tumour k in the destination scan. Any scan in S can be selected as the source. Consider also a segmentation algorithm that consumes a guidance point to produce a mask at any timepoint, such that,

$$\text{GuidedSeg}(I, P^k) \rightarrow M(k) \quad (2)$$

LinGuinE starts from $M_i(k)$, which can be provided by a radiologist directly, or obtained by invoking the segmentation algorithm from an input radiologist click, as in equation (2). P_i^k is then sampled as the center of $M_i(k)$ and is propagated from the coordinate space of I_i to the coordinate space of I_j using image registration. We denote a propagated point as $P_{i \rightarrow j}^k$. The intuition behind point propagation is that tumours maintain an anatomically similar location across timepoints. The registration algorithm is a configurable component of the framework. In this work, we evaluate LinGuinE using publicly available GradICON models (UniGradICON (UGI) [23] and, to show the customization benefit in the context of lung cancer, the LungGradICON (LGI) which is designed specifically for lung registration [24]). Upon computing $P_{i \rightarrow j}^k$, LinGuinE invokes a guided segmentation model to produce $M_j(k)$ on I_j .

$$\text{GuidedSeg}(I_j, P_{i \rightarrow j}^k) \rightarrow M_j(k) \quad (3)$$

After inference, LinGuinE keeps the connected component closest to $P_{i \rightarrow j}^k$ as the final $M_j(k)$ to avoid false positives in nearby tumours. The segmentation of a longitudinal study is then defined as;

$$\text{LinGuinE}(I_i, P_i^k) \rightarrow \{\text{GuidedSeg}(I_j, P_{i \rightarrow j}^k) \mid I_j \in S\} \quad (4)$$

However, the location, shape, and size of a tumour can change over time and $P_{i \rightarrow j}^k$ may land outside the tumour in I_j , even with accurate registration. Hence, performance may be enhanced by leveraging the segmentation model to pull a propagated point closer to the center of a tumour with a technique called ‘boosting’. In boosted configurations, the final prompt to the model is the center of an initial lesion segmentation obtained using $P_{i \rightarrow j}^k$. The below equation then replaces equation (3) to obtain $M_j(k)$. This is an optional feature in LinGuinE.

$$\text{GuidedSeg}(I_j, \text{Center}(\text{GuidedSeg}(I_j, P_{i \rightarrow j}^k))) \rightarrow M_j(k) \quad (5)$$

2.2 Guided Volumetric Segmentation

Any guided segmentation algorithm can be used in the LinGuinE framework. Although ROI-based guidance, as leveraged by Hering et al. [8], offers coarse spatial information, point prompts encode precise locations. Recent advances in the field led to the availability of many point-prompted models for tumour segmentation, and we argue that directly repurposing them yields strong performance in a longitudinal setting. In this work, we evaluate nnInteractive [10], Vista3D [6], and the tumour segmentation component of LesionLocator (LLSeg) [18] within LinGuinE. To demonstrate the benefit of pathology specific optimization, we fine-tune the aforementioned models for lung tumour segmentation.

The datasets used for finetuning include 1) 418 CT scans from the publicly available NSCLC-Radiomics dataset [1], collected at the MAASTRO clinic in the Netherlands, 2) 63 CT volumes from task three of the Medical Imaging Decathlon [22], and 3) A selection of 635 CT scans from a confidential phase 3 study of stage IV non-small cell lung cancer (NSCLC). From each dataset, 16% of the scans were set aside for validation and an additional 16% for testing, leaving 68% for training. For each training run, the same image preprocessing as expected by the pretrained weights was used. The models were trained for up to 200 epochs on an NVIDIA A10G GPU using the PyTorch lightning framework. The checkpoint with the lowest validation Dice-CrossEntropy loss was selected.

2.3 Test Datasets

We evaluate LinGuinE on four datasets, unseen by any of the tested segmentation algorithms. The first is 4DCBCT [2], consisting of CBCT studies and primary tumour segmentations from 10 NSCLC patients over the course of 9 weeks of chemoradiotherapy. Each patient has up to 9 timepoints and 10 scans per time point (one per breathing cycle phase). We use the scan at the 90% phase, corresponding to an inhale. This defines a subset of 50 scans.

The second dataset, Phase-3, is a private dataset originating from a confidential phase 3 monoclonal antibody study targeting metastatic NSCLC. This dataset encompasses data from 30 countries and 2,398 patients. We randomly selected and annotated up to 10 lesions per patient in 354 CT scans from 95 patients, spanning over 52 weeks. Each patient has at least four timepoints segmented by two radiologists with over 15 years of combined clinical experience.

The third is autoPET Longitudinal CT (APLCT) [15], containing whole-body CT at two timepoints across 300 studies of melanoma patients undergoing treatment. The images were acquired at the University Hospital Tübingen.

Finally, UniToChest [5] is comprised of 715 thoracic CT volumes from 623 lung cancer patients, collected and annotated at Citta della Salute e della Scienze Hospital. We selected 51 patients with multiple timepoints and at least one tumour larger than 5mm in axial diameter. At least one tumour per patient was then associated across time by a senior radiologist to create a test set of 78 unique lesions across 109 scans. Our associated ground truth masks will be made publicly available to support benchmarking and further research.




Fig. 1. LinGuinE segmentations of two longitudinal studies in the 4DCBCT dataset. The center of the screening label was propagated to obtain segmentations (red contours) on subsequent timepoints. The ground truth is denoted by yellow contours.

3 Experimental Results

3.1 Segmentation Performance

To demonstrate state-of-the-art performance, we propagated radiologist guidance from the baseline scan to all subsequent scans per patient as in Figure 1. We evaluated multiple registration-segmentation combinations; the best-performing configurations are reported in Table 1. We used Dice to evaluate overall segmentation quality and the Euclidean Distance, ED (mm), between the centers of the ground truth and predicted mask to evaluate lesion tracking. We computed metrics on a per-lesion basis, hence, mis-tracked tumours receive a Dice of zero. Therefore, the metric distributions are not Gaussian and we report the median to more fairly capture performance.

LinGuinE configurations consistently outperform the pretrained LesionLocator model with respect to Dice on every dataset and ED on all lung datasets. While LesionLocator seemingly attains the best ED on APLCT, this metric is undefined when the model produces a false negative, which is a frequent cause of low Dice. This excludes the worst performing cases. Hence, we find LesionLocator to only be an appropriate choice when longitudinal training data is available. For comparison with Hering et al., we trained an nnUNet on the same datasets used for finetuning the guided models for lung cancer. We do not evaluate Hering et al. on APLCT due to lack of relevant data to train a suitable nnUNet.

We found that for lung cancer datasets, using a segmentation algorithm finetuned for lung cancer with LGI in a LinGuinE context tends to yield the top results. This highlights the flexibility of LinGuinE that allows for performance enhancing, task-specific customisations. Nonetheless, as seen in Table 1, LinGuinE can outperform other methods by solely leveraging pretrained general-purpose algorithms like LLSeg, nnInteractive, and UGI - despite the nnUNet used for the Hering et al. comparison being designed specifically for lung cancer.

Table 1. Median Dice and ED (mm) on the test datasets with various configurations of LinGuinE denoted as (registration, segmentation) compared to other methods. We only include a boosted configuration if it yielded the top metric for a dataset. ‘Finetuned’ refers to models finetuned to segment lung tumours. APLCT is not evaluated with lung-specific algorithms, since it contains tumours in all parts of the body. Best metrics are in **bold**. *Italics* indicates that the metric was undefined in $>20\%$ of cases due to false negative predictions.

	Phase-3		4DCBCT		UniToChest		APLCT	
	Dice	ED	Dice	ED	Dice	ED	Dice	ED
Existing Method								
LesionLocator	0.164	<i>2.786</i>	0.098	16.331	0.520	<i>1.770</i>	0.303	1.035
Hering et al.	0.570	9.015	0.728	5.722	0.606	5.226	-	-
LinGuinE Configuration								
UGI, Vista3D	0.528	3.611	0.334	6.427	0.414	4.391	0.080	7.462
UGI, nnInteractive	0.596	3.819	0.719	4.712	0.556	5.390	0.023	14.942
UGI, LLSeg	0.745	1.730	0.680	4.094	0.642	5.131	0.433	4.131
LGI, Vista3D	0.539	3.533	0.350	4.056	0.629	2.624	-	-
LGI, nnInteractive	0.612	3.784	0.704	4.470	0.699	2.461	-	-
LGI, LLSeg	0.759	1.809	0.717	4.105	0.786	1.262	-	-
LGI, Finetuned Vista3D	0.732	1.917	0.773	3.125	0.698	1.876	-	-
LGI, Finetuned nnInter.	0.712	2.051	0.702	5.214	0.780	1.345	-	-
LGI, Finetuned LLSeg	0.736	1.934	0.753	3.667	0.752	1.942	-	-
LinGuinE Configuration with Boosting								
UGI, LLSeg	0.753	1.865	0.662	6.088	0.632	5.111	0.491	3.620
LGI, LLSeg	0.755	1.738	0.672	5.511	0.773	1.198	-	-
LGI, Finetuned LLSeg	0.760	1.635	0.744	5.169	0.775	1.410	-	-

3.2 Temporal Direction Agnosticism and Degradation

For temporal analyses, we consider datasets with more than three timepoints per patient (Phase-3 and 4DCBCT) with their respective best configuration of LinGuinE based on Dice as in Table 1. To demonstrate that LinGuinE can propagate guidance in any temporal direction, we evaluated Dice distributions when using three different starting scans for each patient: (i) the earliest, (ii) the midpoint, and (iii) the latest. The set of segmentations for this three-way comparison is smaller than in other experiments, as we exclude any scan used as a starting point and any tumour not present in all starting scans for a patient. In Phase-3, 149 tumour segmentations from 77 scans across 61 patients were used. We found no statistically significant difference (three-way Friedmann chi-squared test, $p = 0.4$) regardless of the source timepoint. Similar results were noted ($p = 0.3$) in 4DCBCT, across 20 scans from 10 patients.

To analyse temporal degradation, we consider Dice distributions for tumours at various timepoints. As seen in Figure 2, there is minimal correlation between LinGuinE performance and the elapsed time between the baseline (used as the




Fig. 2. LinGuinE displays minimal performance degradation with elapsed time between the baseline and follow-up. Points represent mean Dice at various timepoints with the standard error on the mean displayed as error bars. The line of best fit for each method is shown for trend comparison - LinGuinE has the smallest slope.

source) and the follow up, indicating robustness over the duration of the entire study compared to other methods, which display a clear downward trend. Across all patient studies, we find a mean decline in Dice of 2.591% and 0.457% per week in LesionLocator and Hering et al., respectively. LinGuinE displays a significantly smaller decline of 0.126%. Moreover, the temporal direction agnosticism of our framework enables radiologists to mitigate this by selecting midpoint scans as the source, reducing the elapsed time between scan pairs as needed.

3.3 Ablation Studies

Radiologist inputs To assess the impact of propagating a point sampled from a mask versus a click, two radiologists semi-automatically re-segmented primary lung tumours at one random timepoint for 53 patients in the Phase-3 dataset. Radiologists were instructed to place at least one click on each tumour and were not informed of the experiment’s aim, ensuring authentic interaction. We used those clicks as inputs into the best LinGuine configuration for Phase-3 to obtain 146 tumour segmentations in 132 scans on the remaining timepoints for those patients. The Dice was 0.806. For the same tumours, when full ground truth masks were used to define P_i , it improved to 0.811. The improvement is statistically significant (paired t-test, $p = 0.03$), suggesting that LinGuinE works best with a full segmentation as input. Nonetheless, the performance with a click input is similar to the reported inter-radiologist median Dice of 0.810 for manual lung tumour segmentation [13].

Impact of prompt boosting As seen in Table 1, enabling prompt boosting yields LinGuinE configurations with the best ED metrics on three out of four datasets. The benefit is inconsistent—for example, in 4DCBCT the best performance was achieved without boosting. Conversely, in APLCT boosting improves Dice by 13.4% and ED by 12.4% in the best-performing configuration of UGI

with LLSeg. We hypothesize that boosting is most beneficial when tumour tracking is challenging and registration struggles, but this requires further validation.

Auto-regressive propagation LesionLocator is the only prior method that analyzes an entire longitudinal study rather than a scan pair. Unlike LinGuinE, LesionLocator is autoregressive, propagating the previous timepoint prediction to account for changing tumour morphology. To evaluate the impact of auto-regression, we alter the LinGuinE workflow so that P_i^k is first propagated to the I_j closest in time to I_i . Then, I_j becomes the source and P_j^k is propagated to the next closest scan in time. P_j^k is defined as the center of $M_j(k)$, where $M_j(k)$ is obtained via equation (3). Starting from the baseline scan and using the same algorithms as the best LinGuinE configuration, we compare this autoregressive workflow to the original in equation (4) with paired t-tests. In Phase-3, the autoregressive approach yields similar metrics (Dice=0.758, ED=1.620mm) with insignificant differences (Dice $p = 0.3$, ED $p = 0.1$). In 4DCBCT, we found a performance degradation (Dice=0.734, ED=4.447mm), which is statistically significant (Dice $p = 0.02$, ED $p = 0.003$). In UniToChest, we consider a subset of seven patients with more than two timepoints. The Dice drops from 0.746 to 0.741 (insignificant, $p = 0.4$) and the ED remains at 0.983mm.

Hence, we find no benefit in auto-regression. Our approach enables temporal direction agnosticism and in some scenarios avoids compounding errors by directly considering radiologist input before each propagation. We note that the temporal degradation observed for LesionLocator in Figure 2 differs from what was observed in the original paper [18], and could possibly be attributable to the larger number of time points considered. By contrast, Hering et al. displays behaviour more akin to LinGuinE, and we hypothesize that this is because we invoked Hering et al. in a non-autoregressive manner.

4 Conclusion

In this study, we presented LinGuinE, a framework for longitudinal semi-automatic volumetric segmentation and tracking of tumours without the requirement of training on longitudinal data. LinGuinE demonstrates accurate segmentation on four unseen datasets with a total of 456 longitudinal studies, outperforming other methods. Our method scales to an arbitrary number of timepoints, is temporally agnostic to the initial scan, and segments the tumour boundaries with less temporal degradation than existing approaches. LinGuinE can potentially be applied as a facilitating tool for the introduction of volumetric biomarkers as a response criterion in oncological therapy. We release LinGuinE as a PyTorch package to support this objective and to assist in the creation of annotated longitudinal datasets with tracked tumours. Further work is needed to handle disappearing or coalescing lesions, and to establish LinGuinE as a segmentation tool for different cancer types and imaging modalities.

Acknowledgments. The authors thank Dr. Michael James Parker and Dr. James Michael Parkin for expert radiological review and annotation support for the imaging data used in this study. The authors also thank Dr. Ahmed H. Shahin and Dr. Keerthy Kusumam for their review of the manuscript, and Dr. Mark Barber for support in securing funding.

Disclosure of Interests. This work was funded by AstraZeneca, and all authors were employees of AstraZeneca during the conduct of the study and manuscript preparation.

References

1. Aerts, H.J., Velazquez, E.R., Leijenaar, R.T., Parmar, C., Grossmann, P., Cavalho, S., Bussink, J., Monshouwer, R., Haibe-Kains, B., et al.: Decoding tumour phenotype by noninvasive imaging using a quantitative radiomics approach. *Nature Communications* **5** (6 2014). <https://doi.org/10.1038/ncomms5006>
2. Balik, S., Weiss, E., Jan, N., Roman, N., Sleeman, W.C., Fatyga, M., Christensen, G.E., Zhang, C., Murphy, M.J., Lu, J., Keall, P., Williamson, J.F., Hugo, G.D.: Evaluation of 4-dimensional computed tomography to 4-dimensional cone-beam computed tomography deformable image registration for lung cancer adaptive radiation therapy. *International Journal of Radiation Oncology Biology Physics* **86**(2), 372–379 (6 2013). <https://doi.org/10.1016/j.ijrobp.2012.12.023>
3. Bernstein, M.H., Atalay, M.K., Dibble, E.H., Maxwell, A.W., Karam, A.R., et al.: Can incorrect artificial intelligence (AI) results impact radiologists, and if so, what can we do about it? A multi-reader pilot study of lung cancer detection with chest radiography. *European Radiology* **33**(11), 8263–8269 (11 2023). <https://doi.org/10.1007/s00330-023-09747-1>
4. Cai, J., Tang, Y., Yan, K., Harrison, A.P., Xiao, J., Lin, G., Lu, L.: Deep Lesion Tracker: Monitoring Lesions in 4D Longitudinal Imaging Studies. In: *Computer Vision and Pattern Recognition*. pp. 15154–15164 (2021)
5. Chaudhry, H.A.H., Renzulli, R., Perlo, D., Santinelli, F., Tibaldi, S., Cristiano, C., Grosso, M., Limerutti, G., Fiandrotti, A., Grangetto, M., Fonio, P.: Unitochest: A lung image dataset for segmentation of cancerous nodules on ct scans. In: *Image Analysis and Processing – ICIAP 2022*. pp. 185–196. Springer International Publishing, Cham (2022)
6. He, Y., Guo, P., Tang, Y., Myronenko, A., Nath, V., Xu, Z., Yang, D., Zhao, C., Belue, M., Harmon, S., Xu, D., Li, W.: Vista3d: A unified segmentation foundation model for 3d medical imaging. In: *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 20863–20873 (2025). <https://doi.org/10.1109/CVPR52734.2025.01943>
7. Hein, S.P., Schultheiss, M., Gafita, A., Zaum, R., Yagubbayli, F., Tauber, R., Rauscher, I., Pfeiffer, F., Weber, W.A.: Ai lesion tracking in pet/ct imaging: a proposal for a siamese-based cnn pipeline applied to psma pet/ct scans. *European Journal of Nuclear Medicine and Molecular Imaging* **53**, 429–441 (7 2025)
8. Hering, A., Peisen, F., Amaral, T., Gatidis, S., Eigentler, T., Othman, A., Moltz, J.H.: Whole-body soft-tissue lesion tracking and segmentation in longitudinal CT imaging studies. In: *Proceedings of the Fourth Conference on Medical Imaging with Deep Learning. Proceedings of Machine Learning Research*, vol. 143, pp. 312–326. PMLR (07–09 Jul 2021), <https://proceedings.mlr.press/v143/hering21a.html>

9. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnU-Net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods* **18**(2), 203–211 (2 2021). <https://doi.org/10.1038/s41592-020-01008-z>
10. Isensee, F., Rokuss, M., Krämer, L., Dinkelacker, S., Ravindran, A., others Florian Stritzke, Hamm, B., Wald, T., Langenberg, M., Ulrich, C., Deissler, J., Floca, R., Maier-Hein, K.: nninteractive: Redefining 3d promptable segmentation (2025), <https://arxiv.org/abs/2503.08373>
11. Kim, J.N., Liu, Y., Soans, R., et al.: Omnimamba4d: Spatio-temporal mamba for longitudinal ct lesion segmentation. In: 2025 IEEE 22nd International Symposium on Biomedical Imaging (ISBI). pp. 1–5 (2025). <https://doi.org/10.1109/ISBI60581.2025.10980921>
12. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Dollár, P., Girshick, R.: Segment anything (4 2023)
13. Kothari, G., Woon, B., Patrick, C.J., Korte, J., Wee, L., et al.: The impact of inter-observer variation in delineation on robustness of radiomics features in non-small cell lung cancer. *Scientific Reports* **12**(12822) (jul 2022). <https://doi.org/10.1038/s41598-022-16520-9>
14. Krüger, J., Opfer, R., Gessert, N., Ostwaldt, A.C., Manogaran, P., Kitzler, H.H., Schlaefer, A., Schippling, S.: Fully automated longitudinal segmentation of new or enlarged multiple sclerosis lesions using 3D convolutional neural networks. *NeuroImage: Clinical* **28**, 102445 (1 2020). <https://doi.org/10.1016/j.nicl.2020.102445>
15. Kustner, T., Peisen, F., Gatidis, S., Wagner, A., Megne, O., Othman, A., Sanner, A., et al.: Longitudinal-ct (2025), <https://doi.org/10.57754/FDAT.qwsry-7t837>
16. Luo, X., Wang, G., Song, T., Zhang, J., Aertsen, M., Deprest, J., et al.: MIDeepSeg: Minimally Interactive Segmentation of Unseen Objects from Medical Images Using Deep Learning. *Medical Image Analysis* **72**(102102) (8 2021). <https://doi.org/10.1016/j.media.2021.102102>
17. Primakov, S.P., Ibrahim, A., van Timmeren, J.E., Wu, G., Keek, S.A., et al.: Automated detection and segmentation of non-small cell lung cancer computed tomography images. *Nature Communications* **13**(1) (12 2022). <https://doi.org/10.1038/s41467-022-30841-3>
18. Rokuss, M., Kirchhoff, Y., Akbal, S., Kovacs, B., Roy, S., Ulrich, C., Wald, T., Rotkopf, L.T., Schlemmer, H.P., Maier-Hein, K.: Lesionlocator: Zero-shot universal tumor segmentation and tracking in 3d whole-body imaging. In: 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 30872–30885 (2025). <https://doi.org/10.1109/CVPR52734.2025.02875>
19. Rokuss, M., Kirchhoff, Y., Roy, S., Kovacs, B., Ulrich, C., Wald, T., Denner, S., Isensee, F., Vollmuth, P., Kleesiek, J., Maier-Hein, K.: Longitudinal segmentation of ms lesions via temporal difference weighting. *arXiv:2409.13416* (2024)
20. Sargent, D.J., Rubinstein, L., Schwartz, L., Dancey, J.E., Gatsonis, C., Dodd, L.E., Shankar, L.K.: Validation of novel imaging methodologies for use as cancer clinical trial end-points. *European Journal of Cancer* **45**(2), 290–299 (1 2009). <https://doi.org/10.1016/j.ejca.2008.10.030>
21. Savjani, R.R., Lauria, M., Bose, S., Deng, J., et al.: Automated Tumor Segmentation in Radiotherapy. *Seminars in Radiation Oncology* **32**(4), 319–329 (2022). <https://doi.org/10.1016/j.semradonc.2022.06.002>
22. Simpson, A.L., Antonelli, M., Bakas, S., Bilello, M., Farahani, K., van Ginneken, B., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B., et al.: A large annotated medical image dataset for the development and evaluation of segmentation algorithms (2019), <https://arxiv.org/abs/1902.09063>

23. Tian, L., Greer, H., Kwitt, R., Vialard, F.X., San José Estépar, R., et al.: uni-GradICON: A Foundation Model for Medical Image Registration . In: proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. vol. LNCS 15002. Springer Nature Switzerland (October 2024)
24. Tian, L., Greer, H., Vialard, F.X., Kwitt, R., et al.: Gradicon: Approximate diffeomorphisms via gradient inverse consistency. In: 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 18084–18094 (2023). <https://doi.org/10.1109/CVPR52729.2023.01734>
25. Vizitiu, A., Mohaiu, A.T., Popdan, I.M., Ghesu, F.C., Comaniciu, D.: Multi-scale self-supervised learning for longitudinal lesion tracking with optional supervision. In: proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2023. pp. 573–582. Springer Nature Switzerland, Cham (2023)