

LesionDiffusion: Towards Text-controlled General Lesion Synthesis

Wenhui Lei^{1,2,†}, Hengrui Tian^{1,†}, Linrui Dai^{1,2†}, Hanyu Chen³, and Xiaofan Zhang^{1,2,✉}

¹ Shanghai Jiaotong University

² Shanghai Artificial Intelligence Laboratory

³ The First Hospital of China Medical University
xiaofan.zhang@sjtu.edu.cn

Abstract. Fully-supervised lesion recognition methods in medical imaging face challenges due to the reliance on large annotated datasets, which are expensive and difficult to collect. To address this, synthetic lesion generation has become a promising approach. However, existing models struggle with scalability, fine-grained control over lesion attributes, and the generation of complex structures. We propose LesionDiffusion, a text-controllable lesion synthesis framework for 3D CT imaging that generates both lesions and corresponding masks. By utilizing a structured lesion report template, our model provides greater control over lesion attributes and supports a wider variety of lesion types. We introduce a dataset of 1,505 annotated CT scans with paired lesion masks and structured reports, covering 14 lesion types across 8 organs. LesionDiffusion consists of two components: a lesion mask synthesis network (LMNet) and a lesion inpainting network (LNet), both guided by lesion attributes and image features. Extensive experiments demonstrate that LesionDiffusion significantly improves segmentation performance, with strong generalization to unseen lesion types and organs, outperforming current state-of-the-art models. Code is available at [here](#).

Keywords: Generation Model · Lesion Segmentation · CT inpainting.

1 Introduction

Deep learning has driven significant advances in medical image analysis, especially in lesion detection and diagnosis [9,11]. However, current fully-supervised lesion recognition methods rely heavily on large annotated datasets, which are costly to collect and share due to privacy concerns.

To address the scarcity of lesion data, synthetic lesion generation has emerged as a promising solution. In 3D CT imaging, techniques such as GANs [4], diffusion models [16,6], and physical simulations have been used to generate synthetic lesions for various conditions, including lung nodules [10], COVID-19 [21], and liver tumors [7]. However, these methods train each lesion type independently,

[†] Contributed equally to this work.

which limits their scalability and generalizability. Additionally, they lack fine-grained control over lesion attributes, and are primarily designed to simulate solid lesions, neglecting complex structures such as hollow lesions found in colorectal or gastric cancers [2]. This makes them hard to serve as a general-purpose, controllable 3D CT lesion synthesis model.

In this work, we aim to develop a general, text-controllable lesion synthesis model for 3D CT imaging that generates both lesions and their corresponding masks. Our approach addresses the limitations of previous lesion synthesis methods, which include: 1) the limited variety of lesion types, primarily due to the scarcity of publicly available CT datasets that include both lesion descriptions and corresponding masks; 2) insufficient control over lesion attributes; and 3) the challenges in generating complex lesion textures and masks.

To overcome these limitations and enhance the generalizability and controllability of lesion synthesis, we adopt a structured lesion report template, as proposed in [?], to serve as the textual condition. This template, which covers various aspects of lesion attributes (as illustrated in Fig. 1), enables precise control over lesion mask generation and corresponding image inpainting based on specified attributes. Additionally, we have collected and annotated 1,505 CT scans with paired lesion masks and structured reports, covering 14 types of lesions across 8 organs. Building on this dataset, we developed **LesionDiffusion** (Fig. 2), a text-controlled framework for generalizable lesion synthesis. The framework consists of two key components: 1) a lesion mask synthesis network (LMNet) that is guided by lesion bounding boxes (bbox) and mask attributes, and 2) a lesion inpainting network (LInet) that is guided by both image attributes and the lesion mask. Through extensive experiments, we demonstrate that LesionDiffusion significantly improves segmentation performance across a wide range of lesion types. More importantly, our approach shows exceptional generalization, even for unseen organs and lesion types, achieving a notable improvement in lesion segmentation and surpassing the performance of existing state-of-the-art lesion synthesis models.

2 Methods

2.1 Stage I: Lesion Mask Generation

Concurrent tumor inpainting methods, like DiffTumor [2], adopt an end-to-end patch-wise generation scheme of tumor CTs based on a tumor mask deformed from a unit sphere. Although this generation scheme is plausible in organs whose tumors are characterized by a round-like shape and an arbitrary position, it undermines generalizability as more tumor types do not share these features.

To effectively translate vague shape information into concrete lesion masks, mask-related attributes from structured reports—such as shape, organ type, and location—are used as conditional guidance, as illustrated in Fig.1 and Fig.2 (a).

Next, we develop a diffusion model, named LMNet, to generate lesion masks that adhere to these constraints. Formally, for a precise lesion annotation \mathbf{m}_0 ,




Fig. 1. Lesions and Structured Lesion Reports in the Training Stage. The LesionDiffusion model is trained on 14 types of lesions across 8 organs. The corresponding structured reports include 11 attributes, which are categorized into mask attributes and image attributes.

the forward noising process can be described as

$$\mathbf{m}_t = \sqrt{\bar{\alpha}_t} \mathbf{m}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad (1)$$

where \mathbf{m}_t is the noisy mask at timestep t and ϵ is the Gaussian noise, $\bar{\alpha}_t$ represents the noise variance schedule [6]. Denoting the annotated mask attributes as $\mathcal{A}_{\text{mask}}$, the \mathbf{m}_0 surrounded bounding box as \mathcal{B} and the target organ semantics as \mathcal{M} , we proposed a bbox-weighted loss function with a hyperparameter γ :

$$\mathcal{L}_{\text{LMNet}} = \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} [\|\epsilon \odot \mathcal{B} - \epsilon_{\theta}(\mathbf{m}_t; t, \mathcal{M}, \mathcal{B}, \mathcal{A}_{\text{mask}}) \odot \mathcal{B}\|_2^2 + \gamma \|\epsilon - \epsilon_{\theta}(\mathbf{m}_t; t, \mathcal{M}, \mathcal{B}, \mathcal{A}_{\text{mask}})\|_2^2] \quad (2)$$

This objective is proposed to encourage the model's focus on the constrained area specified by our bounding box \mathcal{B} , thus promoting accurate synthesis of the lesion masks. The target organ semantic is extracted from the surrounding organ semantics as a separate channel before concatenating with the bounding box \mathcal{B} and serving as a spatial condition for tumor site generation. Mask attributes $\mathcal{A}_{\text{mask}}$ are embedded using a pretrained text encoder in BiomedCLIP [22] and fused into the generation process via cross-attention modules to describe shape, location and other pertinent spatial information to the diffusion model.

2.2 Stage II: Lesion Image Inpainting

Similar to the mask-attributes-guided generation of the lesion mask, we utilize image attributes—such as density, variation, and surface characteristics—to perform lesion image inpainting in Stage II, as illustrated in Fig.1 and Fig. 2 (c).

Due to the large memory consumption of 3D CT images, we choose Latent-Diffusion [15] to train the diffusion model on a smaller latent space encoded by a pair of dedicated autoencoders \mathcal{E}_{CT} , \mathcal{D}_{CT} . This is done following a similar

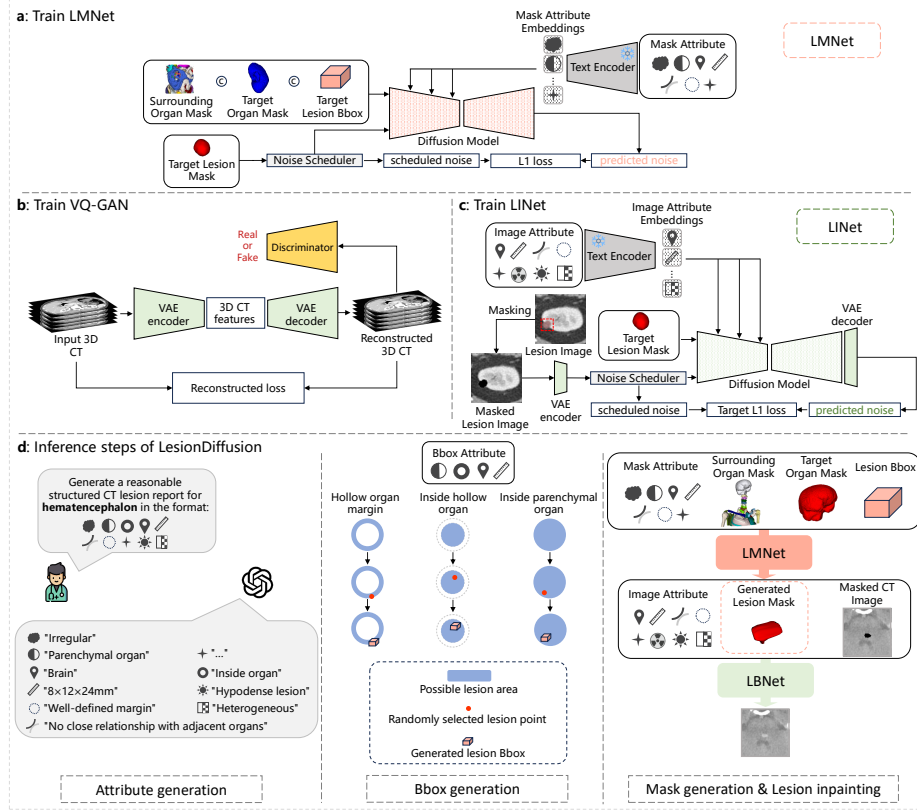


Fig. 2. Overview of LesionDiffusion framework. (a) LMNet is trained to generate lesion masks; (b) VQ-GAN is trained to compress 3D CT images into latent space and then reconstruct them; (c) LINet is trained to perform lesion inpainting in the latent space; (d) During the inference stage, the framework generalizes to any lesion type, involving lesion attribute generation, lesion bounding box generation, lesion mask generation, and inpainting.

paradigm as VQ-VAE [18], which encodes images with a discretized codebook \mathbf{c} and quantization step \mathbf{q}_c . We also adopt an adversarial module \mathcal{D}_{disc} to boost autoencoding performance similar to DiffTumor [2]. For each CT image \mathbf{x} , denoting its encoded latent as $\mathbf{z} = \mathcal{E}_{CT}(\mathbf{x})$ and its reconstructed counterpart as $\hat{\mathbf{x}} = \mathcal{D}_{CT}(\mathbf{q}_c(\mathbf{z}))$, the training objective of our autoencoder is defined as the combination of l_2 reconstruction loss, codebook loss [3] and discriminative loss from the adversarial module

$$\begin{aligned} \mathcal{L}_{AE} = & \|\mathbf{x} - \hat{\mathbf{x}}\|_2^2 + \|\Omega(\mathbf{z}) - \mathbf{q}_c(\mathbf{z})\|_2^2 + \|\Omega(\mathbf{q}_c(\mathbf{z})) - \mathbf{z}\|_2^2 \\ & + \log \mathcal{D}_{disc}(\mathbf{x}) + \log(1 - \mathcal{D}_{disc}(\hat{\mathbf{x}})) \end{aligned} \quad (3)$$

where Ω is the stop-gradient function. After training the paired autoencoder, we develop another diffusion model, dubbed LINet, on the latent space of \mathcal{E}_{CT} . Since the annotated lesion masks \mathbf{m} encompass the full extent of pathological regions—including both the lesion core and affected surrounding tissue—we choose to perform the diffusion process solely within the lesion mask. Formally, the forward diffusion process can be described as

$$\mathbf{z}_t = (\sqrt{\bar{\alpha}_t}\mathbf{z}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon) \odot \text{Resize}(\mathbf{m}) + \mathbf{z}_0 \odot (1 - \text{Resize}(\mathbf{m})) \quad (4)$$

where $\text{Resize}(\cdot)$ function is used to resize the semantic condition \mathbf{m} to be of the same spatial resolution as the image latent \mathbf{z}_0 , ϵ is the Gaussian noise, and $\bar{\alpha}_t$ represents the noise variance schedule. The denoising step uses LINet to retrieve inpainted image latent from noises conditioned on image attributes $\mathcal{A}_{\text{image}}$ encoded by BiomedCLIP. To optimize the performance of the diffusion model, we leverage the bounding box \mathcal{B}' surrounding the lesion area in the mask and use l_1 loss to measure the difference between the input noise ϵ and the denoised latents ϵ_θ within this box. Formally, the training objective is

$$\mathcal{L}_{\text{LINet}} = \mathbb{E}_{\epsilon \sim \mathcal{N}(0, I)} \left[\frac{1}{|\mathcal{B}'|} \|\epsilon \odot \mathcal{B}' - \epsilon_\theta(\mathbf{z}_t; t, \text{Resize}(\mathbf{m}), \mathcal{A}_{\text{image}}) \odot \mathcal{B}'\|_1 \right] \quad (5)$$

2.3 Inference process

In the inference stage of LesionDiffusion, given the lesion to be simulated, we first generate its structured report and then sample to obtain the bounding box (bbox). Finally, we use LMNet and LINet to obtain the lesion mask and inpainted CT images, as shown in Fig. 2.

Specifically, we propose two methods for generating the complete structured lesion report: (1) the user selects possible options for each attribute based on medical knowledge, followed by controlled combination; (2) the user queries a large language model (LLM) to generate the report. From the report, we extract four key attributes: "organ", "organ type", "lesion location", and "size", which guide a rule-based bounding box (bbox) generation pipeline. Based on the surrounding organ mask and these attributes, we categorize plausible lesion regions—such as “inside hollow organ,” “hollow organ margin,” or “inside parenchymal organ”—and sample bbox centers accordingly under spatial constraints. If the sampled location or size violates predefined rules, we resample or regenerate the report until a valid configuration is obtained.

Next, the lesion mask attribute, surrounding organ mask, target organ mask, and lesion bbox are fed into LMNet. We use DDIM [17] for fast sampling from the learned lesion mask distribution. Since the generation result is a continuous floating-point intensity map, we apply a simple capping threshold to obtain a binary tumor semantic mask. For LINet inference, firstly VQ-GAN encoder compresses CT image input into latent space. Then Gaussian noise replaces the latent codes within the mask region, resulting in $z_T = \epsilon \odot m + z_0 \odot (1 - m)$, and execute T time steps of reverse denoising to obtain the edited latent code, which is then decoded back into a CT image by the VQ-GAN decoder.

3 Experiment


Dataset Construction We compiled several datasets for generative training and downstream evaluation. Specifically, we assemble 1505 CT scans and masks collected from public and private sources: KiTS23 [5]: kidney tumor and cyst (489 scans); MSD [1]: colon tumor (126 scans), liver tumor (303 scans), lung tumor (96 scans), pancreas tumor (216 scans), pancreas cyst (65 scans); private data collected at The First Hospital of China Medical University: liver cyst (30 scans), gallbladder cancer (30 scans), gallstones (30 scans), esophageal cancer (30 scans), gastric cancer (30 scans), kidney stone (30 scans), bladder cancer (30 scans). All these scans are annotated with structured lesion reports by four radiologists. The dataset was partitioned into training and testing sets with a 4:1 ratio for each lesion, resulting in 1204 training samples and 301 testing samples. We also set a hematencephalon segmentation dataset, INSTANCE22 [12] (100 scans, train:test=80:20), for the generalization test, which is not included in the generative model training set. We further collected 10767 CT images at The First Hospital of China Medical University as templates for the inpainting stage. We utilize Totalsegmentator[19] to obtain organ segmentation for all CT scans.

Implementation Details All generative trainings are performed under a static learning rate of 1×10^{-4} with a batch size of 1 per GPU and a patch size of $128 \times 128 \times 128$ on 4 NVIDIA 4090 GPUs. The training processes use AdamW [14] optimizers with a momentum of 0.99 and a weight decay of 1×10^{-5} . The diffusion model uses cosine noise schedule with 1000 timesteps and DDIM denoising [17] of 200 timesteps to recover CT latents from Gaussian noise at a guidance scale of 2.5. Moreover, our image autoencoder downsamples the original image by $2\times$ and uses a codebook of 2048 entries. γ is set to 0.001. For the downstream segmentation tasks, we train each segmentation network using the official implementation of nnUNet-v2 [8]. Performance was evaluated for each lesion under three settings: (1) training the segmentation model using 1000 synthetic lesions generated from template data; (2) fine-tuning the nnUNet-V2 model, initially trained on synthetic data, with real data from the corresponding lesion’s training set; (3) training the segmentation model exclusively on the real dataset. The same settings were used when evaluating other synthesis models.

Downstream usability We first validate LesionDiffusion’s competence as a text-based dataset augmentator by comparing the performance of dedicated segmentation models trained with and without LesionDiffusion-synthesized samples. For reference, we also compare LesionDiffusion’s augmentation performance with another lesion-generating framework, DiffTumor [2]. As shown in Tab. 1, samples synthesized by LesionDiffusion can provide segmentation guidance comparable to that of real samples and far better than that of DiffTumor-synthesized samples. Specifically, we observe an averaged 6.9% DSC increase in segmentation models pre-trained with LesionDiffusion-synthesized lesions instead of DiffTumor. These

Table 1. Dice Similarity Coefficient(DSC) comparison under three settings real lesion data only, synthetic lesion data only, and synthetic lesion data augmented by real lesion data.

Methods	Stone		Cyst			Cancer										Avg.
	Gall.	Kid.	Liver	Pan.	Kid.	Eso.	Gall.	Lung	Liver	Pan.	Kid.	Sto.	Colon	Blad.		
Real	77.6	70.4	76.4	78.2	59.8	57.7	60.4	69.0	61.8	49.6	79.2	66.0	58.1	83.3	67.7	
DiffTumor [2]	58.1	17.5	61.6	66.2	55.1	44.8	49.2	54.8	44.6	30.1	59.7	56.4	41.6	65.2	50.4	
DiffTumor+Real	75.1	54.4	84.4	79.2	63.1	62.0	63.4	66.7	48.0	44.1	81.9	71.3	59.9	79.5	66.6	
LesionDiffusion	71.9	21.0	77.2	65.7	48.3	60.4	55.7	61.3	63.8	35.8	60.5	61.1	42.6	76.2	57.3	
LesionDiffusion+Real	77.2	67.0	83.0	80.4	62.5	61.8	66.0	68.9	64.7	51.7	82.6	70.0	61.4	84.2	70.1	

**Fig. 3. Generalization to hemencephalon.** (a) Downstream segmentation results for hemencephalon. (b) Examples synthesized by LesionDiffusion.

models also performs the best on average after task-specific finetuning, outperforming models trained only with real data by 2.4% and models pre-trained with DiffTumor-synthesized lesions by 3.5%.

Downstream Generalization We also expect the synthetic ability of LesionDiffusion to be generalizable across both seen and unseen downstream tasks. Thanks to our comprehensive structured medical report template which encompasses key attributes for describing lesion shape and texture, LesionDiffusion can be easily transferred to new tasks by using only an appropriate set of attribute inputs, as shown in Fig. 3. To quantify LesionDiffusion’s generalization ability, we use it to synthesize samples of hemencephalon for supervising a dedicated segmentation model. As illustrated in Fig. 3(a), despite LesionDiffusion’s original training set lacking brain CT images and hemorrhage lesion types, models trained on LesionDiffusion synthesized samples can still provide segmentation guidance comparable to that of real samples. In contrast, DiffTumor suffers significantly from this domain gap, with a 46% drop in DSC compared to LesionDiffusion. Moreover, after task-specific finetuning, model trained using LesionDiffusion synthesized lesions manages to provide a segmentation performance 15% better than that trained with real data and 5% better than that trained with DiffTumor.

Generation alignment with textual guidance Firstly, we quantitatively evaluate LMNet’s conditional alignment between generated tumor shape and prompted

attributes using two morphological metrics individually evaluating the convexity [13] and sphericity [20] of generated masks. Mathematically, a higher convexity and sphericity indicate that the lesion more closely resembles a spherical shape. As shown in Fig. 4(a), we can see a clear trend that masks generated with "round-like" prompts are higher in both metrics compared to those with "irregular" prompts. Additionally, we qualitatively evaluate LINet’s conditional alignment by alternating several image attributes in Fig. 4(b,c), the results show promising conformity with the input conditions.

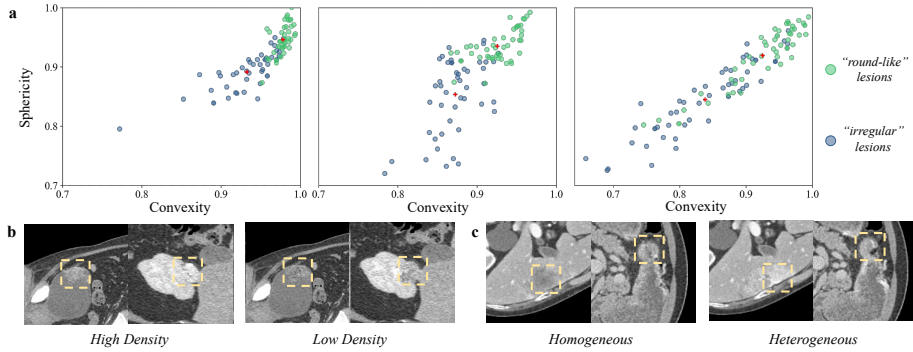


Fig. 4. Generation alignment with textual guidance. This figure illustrates the effect of shape, density, and density variation conditions on image inpainting results. (a) Green and blue dots represent the morphological metrics computed for round-like and irregular tumor volumes respectively, with red crosses indicating the metric mean values. From left to right, the examples correspond to an esophageal tumor, bladder tumor, and liver cyst. (b) Comparison of generation control under the condition of density. (c) Comparison of generation control under the condition of density variation.

Ablation study of textual condition Ablation studies of LesionDiffusion are performed by removing textual conditions in the generation process, forcing the model to only implicitly learn lesion features through the training process. As shown in Tab. 2, the results show different degrees of decrease in downstream performances with respect to all different tasks, with an averaged decrease of 15.4% DSC. This effect is particularly pronounced for organs containing multiple lesion types in the training set, such as the liver, kidney, and pancreas, and is primarily due to the absence of textual conditions, which causes the model to be uncertain about the specific lesions to generate in these organs.

4 Conclusion

In conclusion, we introduce LesionDiffusion, a text-controllable framework for inpainting lesions in 3D CT images. By utilizing a structured lesion report tem-

Table 2. Ablation studies of LesionDiffusion.

Settings	Stone		Cyst			Cancer										Avg.
	Gall.	Kid.	Liver	Pan.	Kid.	Eso.	Gall.	Lung	Liver	Pan.	Kid.	Sto.	Colon	Blad.		
LesionDiffusion	71.9	21.0	77.2	65.7	48.3	60.4	55.7	61.3	63.8	35.8	60.5	61.1	42.6	76.2	57.3	
LesionDiffusion <i>w.o.</i> text	62.4	15.1	42.8	48.6	24.9	60.2	50.7	57.5	44.3	20.1	41.9	33.7	31.8	53.1	41.9	

plate, the framework enables fine-grained control over lesion attributes and accommodates a diverse range of lesions. Our model demonstrates substantial improvements in segmentation performance and exhibits strong generalization to unseen lesions and organs. LesionDiffusion outperforms current state-of-the-art models, providing a scalable solution for lesion recognition in medical imaging. While promising, further research is needed to explore its generalization across a broader range of diseases, and the training dataset is still limited. In future work, we aim to expand the training dataset and validation scenarios to develop a controllable lesion generation foundation model.

Acknowledgments. This work was supported by National Natural Science Foundation of China (NSFC 62301311) and the Fundamental Research Funds for the Central Universities (project number YG2022QN029).

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

References

1. Antonelli, M., Reinke, A., Bakas, S., Farahani, K., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B., Ronneberger, O., Summers, R.M., et al.: The medical segmentation decathlon. *Nature communications* **13**(1), 4128 (2022)
2. Chen, Q., Chen, X., Song, H., Xiong, Z., Yuille, A., Wei, C., Zhou, Z.: Towards generalizable tumor synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11147–11158 (2024)
3. Esser, P., Rombach, R., Ommer, B.: Taming transformers for high-resolution image synthesis. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 12873–12883 (2021)
4. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial networks. *Communications of the ACM* **63**(11), 139–144 (2020)
5. Heller, N., Isensee, F., Trofimova, D., Tejpaul, R., Zhao, Z., Chen, H., Wang, L., Golts, A., Khapun, D., Shats, D., Shoshan, Y., Gilboa-Solomon, F., George, Y., Yang, X., Zhang, J., Zhang, J., Xia, Y., Wu, M., Liu, Z., Walczak, E., McSweeney, S., Vasdev, R., Hornung, C., Solaiman, R., Schoepfoerster, J., Abernathy, B., Wu, D., Abdulkadir, S., Byun, B., Spriggs, J., Struyk, G., Austin, A., Simpson, B., Hagstrom, M., Virnig, S., French, J., Venkatesh, N., Chan, S., Moore, K., Jacobsen, A., Austin, S., Austin, M., Regmi, S., Papanikolopoulos, N., Weight, C.: The kits21 challenge: Automatic segmentation of kidneys, renal tumors, and renal cysts in corticomedullary-phase ct (2023)

6. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
7. Hu, Q., Chen, Y., Xiao, J., Sun, S., Chen, J., Yuille, A.L., Zhou, Z.: Label-free liver tumor segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 7422–7432 (2023)
8. 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 (2021)
9. Jiang, Y., Lei, W., Zhang, X., Zhang, S.: Unleashing the potential of vision-language pre-training for 3d zero-shot lesion segmentation via mask-attribute alignment. *arXiv preprint arXiv:2410.15744* (2024)
10. Jin, D., Xu, Z., Tang, Y., Harrison, A.P., Mollura, D.J.: Ct-realistic lung nodule simulation from 3d conditional generative adversarial networks for robust lung segmentation. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part II* 11. pp. 732–740. Springer (2018)
11. Lei, W., Chen, H., Zhang, Z., Luo, L., Xiao, Q., Gu, Y., Gao, P., Jiang, Y., Wang, C., Wu, G., et al.: A data-efficient pan-tumor foundation model for oncology ct interpretation. *arXiv preprint arXiv:2502.06171* (2025)
12. Li, X., Luo, G., Wang, K., Wang, H., Liu, J., Liang, X., Jiang, J., Song, Z., Zheng, C., Chi, H., Xu, M., He, Y., Ma, X., Guo, J., Liu, Y., Li, C., Chen, Z., Siddiquee, M.M.R., Myronenko, A., Sanner, A.P., Mukhopadhyay, A., Othman, A.E., Zhao, X., Liu, W., Zhang, J., Ma, X., Liu, Q., MacIntosh, B.J., Liang, W., Mazher, M., Qayyum, A., Abramova, V., Lladó, X., Li, S.: The state-of-the-art 3d anisotropic intracranial hemorrhage segmentation on non-contrast head ct: The instance challenge (2023)
13. Lian, Z., Godil, A., Rosin, P.L., Sun, X.: A new convexity measurement for 3d meshes. In: *2012 IEEE Conference on Computer Vision and Pattern Recognition*. pp. 119–126 (2012)
14. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101* (2017)
15. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
16. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *International conference on machine learning*. pp. 2256–2265. PMLR (2015)
17. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502* (2020)
18. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. *Advances in neural information processing systems* **30** (2017)
19. Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D.T., Cyriac, J., Yang, S., et al.: Totalsegmentator: robust segmentation of 104 anatomic structures in ct images. *Radiology: Artificial Intelligence* **5**(5), e230024 (2023)
20. Weirich, M., Misiulia, D., Antonyuk, S.: Characterization of particle shape with an improved 3d light scattering sensor (3d-lss) for aerosols. *Sensors* **24**(3) (2024)
21. Yao, Q., Xiao, L., Liu, P., Zhou, S.K.: Label-free segmentation of covid-19 lesions in lung ct. *IEEE transactions on medical imaging* **40**(10), 2808–2819 (2021)

22. Zhang, S., Xu, Y., Usuyama, N., Xu, H., Bagga, J., Tinn, R., Preston, S., Rao, R., Wei, M., Valluri, N., Wong, C., Tupini, A., Wang, Y., Mazzola, M., Shukla, S., Liden, L., Gao, J., Crabtree, A., Piening, B., Bifulco, C., Lungren, M.P., Naumann, T., Wang, S., Poon, H.: Biomedclip: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs (2025)