From Pixel to Cancer:
Cellular Automata in Computed Tomography

Yuxiang Lai[1,2], Xiaoxi Chen[3], Angtian Wang[1],
Alan Yuille[1], and Zongwei Zhou[★1]

[1] John Hopkins University
[2] Southeast University
[3] University of Illinois Urbana-Champaign

Abstract. AI for cancer detection encounters the bottleneck of data scarcity, annotation difficulty, and low prevalence of early tumors. Tumor synthesis seeks to create artificial tumors in medical images, which can greatly diversify the data and annotations for AI training. However, current tumor synthesis approaches are not applicable across different organs due to their need for specific expertise and design. This paper establishes a set of generic rules to simulate tumor development. Each cell (pixel) is initially assigned a state between zero and ten to represent the tumor population, and a tumor can be developed based on three rules to describe the process of growth, invasion, and death. We apply these three generic rules to simulate tumor development—from pixel to cancer—using *cellular automata*. We then integrate the tumor state into the original computed tomography (CT) images to generate synthetic tumors across different organs. This tumor synthesis approach allows for sampling tumors at multiple stages and analyzing tumor-organ interaction. Clinically, a reader study involving three expert radiologists reveals that the synthetic tumors and their developing trajectories are convincingly realistic. Technically, we analyze and simulate tumor development at various stages using 9,262 raw, unlabeled CT images sourced from 68 hospitals worldwide. The performance in segmenting tumors in the liver, pancreas, and kidneys exceeds prevailing literature benchmarks, underlining the immense potential of tumor synthesis, especially for earlier cancer detection. The code and models are available at https://github.com/MrGiovanni/Pixel2Cancer

Keywords: Data Synthesis · Cellular Automata · Tumor Segmentation.

1 Introduction

Training AI models for image segmentation needs extensive datasets with detailed per-pixel annotations To reduce manual annotation efforts, *data synthesis* is a compelling approach for generating enormous synthetic data and annotations for both training [19,12] and evaluating [13] AI models. This is particularly

---

★ Correspondence to: Zongwei Zhou (zzhou82@jh.edu)

attractive for segmenting small tumors, which are clinically important to detect at the early stage but rarely occur in public datasets because they are difficult for radiologists to identify.

Preliminary studies have been successful in synthesizing specific disease conditions in medical images, such as synthesizing pulmonary inflammatory lesions in COVID-19 [19], lung nodules [7], abdominal tumors [15,2], and brain tumors [31]. However, key challenges remain: (i) generating highly realistic tumors, (ii) synthesizing tumors that generalize across multiple organs, and (iii) leveraging synthetic tumors for AI model training. Moreover, while some mathematical studies [3,22] have modeled tumor development in simplified hypothetical environments, they cannot generate realistic synthetic data for model training.

In this paper, we revisit a classic technique, cellular automata [30], incorporating novel designs to generate realistic tumors of various shapes. In cellular automata, the basic element is the *cell*, which our work refers to as a single *pixel* in the computed tomography (CT) image. We propose Pixel2Cancer, a framework that can generate synthetic tumors guided by three generic rules we have developed[1]. As a result, Pixel2Cancer can generate multiple stages of tumor development (illustrated in Figure 1), simulate the interaction between tumors and their surrounding environments, exhibit applicability across various organs, and potentially contribute to the estimation of tumor prognosis in longitudinal data.

We have evaluated the Pixel2Cancer from two perspectives. *Clinically*, reader studies involving three expert radiologists have been conducted, and results have been so convincing that even medical professionals with over ten years of experience can mistake them for real tumors (Table 1). *Technically*, we have generated synthetic tumors on healthy data to train the segmentation model. The performance in segmenting tumors in the liver, pancreas, and kidneys exceeds previous benchmarks and even model training on real data (Appendix Table 3), underlining the immense potential of tumor synthesis, especially for earlier cancer detection/diagnosis (Appendix Figure 4).

This performance is attributable to our key observation: Many tumors, regardless of origin, share underlying growth, invasion, and interaction dynamics. Our generic rules effectively capture these shared processes, enabling tumor synthesis across diverse organs. Benign tumors and early-stage cancers usually grow in a contained manner, pushing against surrounding tissues rather than invading them. In contrast, advanced cancers often exhibit invasive growth, infiltrating surrounding tissues and potentially spreading to distant parts of the body (metastasis)[2]. Furthermore, we found that the 9,262 raw CT images we used were captured during the venous phase, during which a significant proportion

---

[1] We develop these generic rules by imitating tumor behavior based on medical knowledge (detailed in §2), including tumor growth and death [3], tissue invasion [4], and interactions with the tumor microenvironments [27].

[2] Using liver cancer as an example, early-stage liver cancer commonly presents as small nodules with a relatively regular shape. In contrast, advanced liver cancer frequently demonstrates infiltration of intrahepatic blood vessels and bile ducts, or distant metastasis. Similar patterns are observed in other solid tumors as well [20].

of tumors appeared hypo-intense (dark). Therefore, we chose to simulate this typical presentation, closely aligning with clinical practice. This observation implies the existence of fundamental tumor development principles that transcend organ-specific variations. We hereby contrast Pixel2Cancer with **related work** to show our innovations and novelties.

1. **Requiring no manual annotation.** Learning-based approaches, such as GAN [6] and Diffusion [11], excel in learning tumor representations but require abundant paired tumor data for effective CT image generation. Moreover, the generation process requires extra manual efforts, including the creation of masks to indicate tumor locations and shapes [15].

2. **Simulating tumors development.** None of the existing synthetic approaches can simulate tumor development, especially in the intricate processes of proliferation and invasion [8]. The complexity of these processes, influenced by the surrounding environment [29], makes it difficult to generate synthetic tumors when interacting with organ tissues and structures, particularly when generalizing to different organs.

3. **Synthesizing tumors across organs.** Modeling-based approaches leverage specialized design and domain expertise to simulate tumor appearances. Although the utilization of modeling-based synthetic tumors eliminates the need for manual annotation [12], they require significant effort for designing proper tumor characteristics customized to a specific organ, therefore limiting generalization, especially across different organs.

2   Pixel2Cancer

Overview. We begin by quantifying an organ (e.g., liver) based on CT values, assigning each pixel to four levels: level 1 (softest) to level 4 (hardest), with vessels/boundaries as level 0. Next, a tumor is generated from an arbitrary pixel of the organ tissue (levels > 0). Inspired by cellular automata, we design three generic rules in §2.1: (1) growth, (2) invasion, and (3) death. These rules oversee the tumor development and interaction with the quantified organ (see §2.2), maintaining a dynamically updated tumor population map. Finally, we translate this map back to the original CT image (see §2.3), modifying the CT values based on the original values and the tumor population—a greater tumor population denotes a smaller CT value (visually darker). As shown in Figure 1, our Pixel2Cancer enables the generation of tumors at various stages in CT images.

2.1   Designing Three Rules for Tumor Development

Since CT pixel scales are larger than cell scales, we introduce a *tumor population map* to quantify the tumor population state within each cell (*pixel*). The state ranges from 0 (no tumor) to 10 (maximal tumor population), updating over time based on the states of neighboring cells. A key challenge in cellular automata is designing the rules overseeing state evolution. To initiate tumor development,
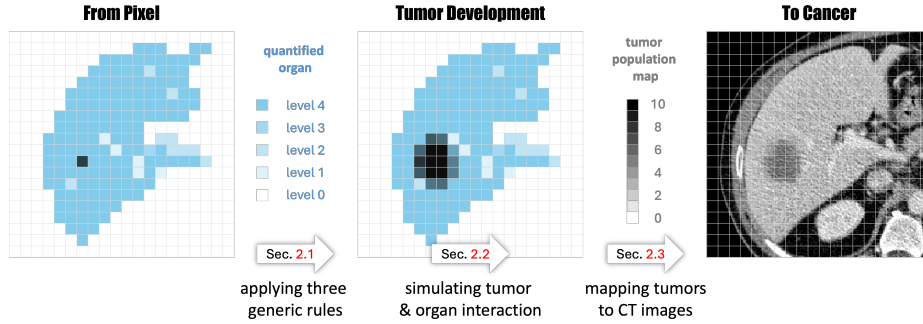
Fig. 1. Initially, we quantify the organ from CT intensity to create the quantified organ and select a pixel as the starting point. Subsequently, we apply three generic rules to simulate tumor development based on the quantified organ, including growth, interaction, and death (§2.1). Then, using these generic rules, Pixel2Cancer simulates tumor development and records simulation results using a tumor population map (§2.2). Finally, we generate tumors in CT scans through the mapping function according to the tumor population map and CT intensity (§2.3).

we randomly select a cell in the organ region and assign its state $= 1$. We then defined three generic rules to simulate tumor shape, texture, and behaviors:

- **R1: Growth.** Cells with a state greater than 0 but less than 10 can accumulate with a probability (self-state $+ 1$). This rule simulates the proliferative behavior of the tumor.
- **R2: Invasion.** Cells with a state greater than 0 can incentivize one of its neighbors to grow (neighbor-state $+ 1$). This rule simulates the invasive behavior of the tumor.
- **R3: Death.** Cells with a state of 10 surrounded by neighbors with 10 can die with a probability (self-state $\leftarrow -1$). This rule simulates the effect of crowding in the tumor, which leads to the death of tumor cells.

R1 and R3 are straightforward, but R2's complexity arises from its potential to influence neighboring cells. Tumor invasion depends not only on its population but also on the nature of neighboring cells (e.g., tissue, vessels, boundaries, etc.). To address this, the next section will describe how we model the interaction between organs and tumors during the invasion process (R2).

2.2   Simulating Interaction Between Organs and Tumors

Quantified Organ. CT images typically exhibit a wide range of intensities, typically ranging from -1000 to 1000 Hounsfield Units (HU), which poses challenges for simulating interactions between synthetic tumors and surrounding tissues. To simplify the simulation process, we convert the organ in CT images into four levels, resulting in the quantified organ. As shown in Figure 1, the quantified map categorizes the CT intensity into four levels, spanning from
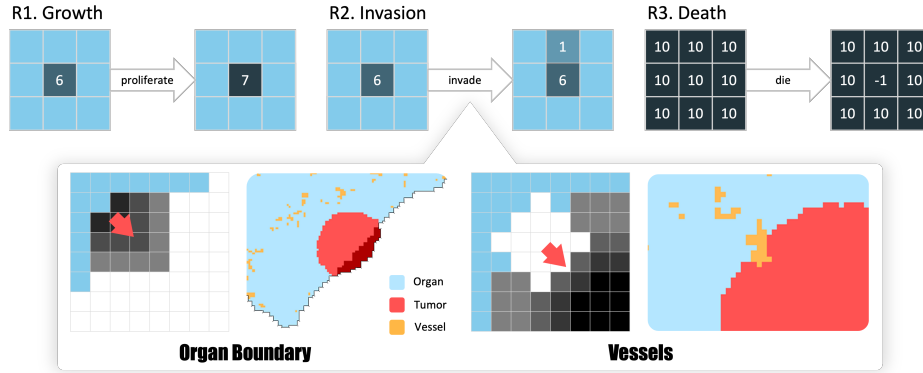
Fig. 2. R1. Growth: Tumor cells can proliferate themselves (self-state$+1$) with probability. **R2. Invasion:** Tumor cells can invade neighboring cells (neighbor-state $+ 1$). In this process, we simulate interactions among tumors, organ tissues, vessels, and boundaries. At the bottom line, we present cases where tumors are compressed by organ boundaries and vessels. **R3. Death:** Tumor cells surrounded by a full population of neighboring cells (state $= 10$) will undergo cell death (self-state $\leftarrow -1$).

low to high intensity (level 1 to 4), with vessels/boundaries assigned to level 0. Subsequently, Pixel2Cancer utilizes this quantified organ as the simulation environment to simulate growth and interaction (Figure 2).

Growth and Interaction: As mentioned in §2.1, the tumor will proliferate itself on the tumor population map (self-state $+1$). Increasing populations lead to greater space occupation, resulting in interactions with and invasion of surrounding tissues. Crowded tumor cells will randomly choose directions and attempt to invade. When interacting with organs, tumors can directly invade normal tissues (level 1,2,3) at different rates. However, the presence of physical barriers such as vessels, boundaries (level 0), and high-density organ tissues (level 4) can impede tumor growth. In this case, crowded tumor cells (state $= 10$) exert pressure as they attempt to invade surrounding tissues. The pressure levels range from 0 to n, indicating the number of invasion attempts (Figure 2). When the pressure of the target direction reaches the threshold, tumor cells can invade these high-density tissues. In Figure 2, we also simulate tumor cell death, influenced by factors such as the vascular feeding environment and cell size [23]. Our model corresponds to cells surrounded by maximal population neighbors (neighbor-states $= 10$), and certain evolutionary steps (of sufficient size) can lead to cell death.

2.3 Mapping Tumors to CT Images

After simulating tumor development, the resulting tumor population map needs to be mapped back into CT images (Figure 1). The mapping process is based on a mapping function, eliminating the need for annotations and training processes. It integrates information from the tumor population map and original CT inten-

Table 1. Visual Turing Test involved three experts each evaluating 150 CT images, with 50 per organ. They were tasked with categorizing each CT image as either *real*, or *synthetic*. A lower specificity score indicates a higher number of synthetic tumors being identified as real. We also provide typical examples of synthetic tumors recognized correctly and incorrectly. The red dot indicates the human expert believes this is a synthetic tumor, while the green dot indicates they believe it is a real tumor.

	metric	liver	pancreas	kidneys
R1	sensitivity (%)	100	95.0	95.5
3-year	specificity (%)	27.3	22.7	26.7
experience	accuracy (%)	60.9	57.1	67.6
R2	sensitivity (%)	94.7	87.5	90.0
7-year	specificity (%)	47.8	47.4	56.3
experience	accuracy (%)	69.1	65.7	75.0
R3	sensitivity (%)	100	100	100
10-year	specificity (%)	45.4	55.6	57.9
experience	accuracy (%)	68.4	72.4	75.8

positives: real tumors ($N = 25$); negatives: synthetic tumors ($N = 25$).

synthetic liver tumors   synthetic pancreatic tumors   synthetic kidney tumors
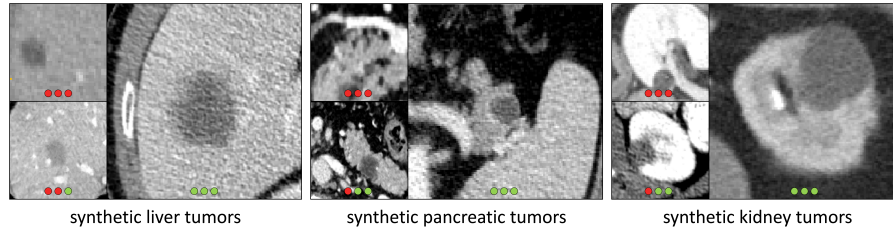
sity. Tumor population maps guide the mapping process to modify CT intensity values, reflecting different tumor states, including dead cells (state $= -1$), tumor cells (state $= 10$), normal cells, and blurry boundaries where two cell types mix ($0 <$ state < 10). To identify the intensity of each tumor pixel, we incorporate typical intensity differences between organ tissues and tumors. Then, we can calculate the intensity by considering the proportion of tumor populations within pixels. To create realistic tumor textures, we introduce 3D Gaussian noise with predefined mean intensity and standard deviation as organ parenchyma.

3   Fake Tumors, Real Results

3.1   Experiment

Datasets: For *training*, we synthesize tumors on healthy CT data. The healthy liver dataset comprises images with healthy livers assembled from CHAOS [16], BTCV [17], Pancreas-CT [25], and healthy subjects in LiTS. The healthy pancreas dataset comprises images from Pancreas-CT and BTCV. The healthy kidney dataset consists of images from BTCV, WORD [18], and Abdomenct-1k [21]. For *evaluation*, Liver tumor annotations are from LiTS [1]. Pancreas tumors are annotated in MSD (Task07 Pancreas) [26]. Kidney annotations are from KiTS [10]. We use 5-fold cross-validation to evaluate the performance.

Table 2. Performance of liver, pancreatic, and kidney tumor segmentation. Pixel2Cancer outperforms Hu et al. [12] and even surpass models trained on real tumors. Higher DSC and NSD values imply more precise overall segmentation; lower SD and HD values suggest more accurate segmentation boundaries.

| organ	tumor	U-Net		Swin UNETR		nnU-Net	
		DSC/NSD (%) ↑	SD/HD (mm) ↓	DSC/NSD (%) ↑	SD/HD (mm) ↓	DSC/NSD (%) ↑	SD/HD (mm) ↓
liver	real tumors	56.7/58.0	23.2/61.4	53.5/55.9	21.3/57.8	56.2/55.3	24.6/58.3
	Hu et al.	54.5/57.6	23.8/58.8	52.3/56.5	22.9/56.9	53.7/56.1	22.5/57.2
	Pixel2Cancer	58.9/63.7	17.9/52.4	56.7/62.5	18.7/51.3	57.9/63.2	18.9/52.7
pancreas	real tumors	57.8/56.5	13.1/47.7	56.7/52.8	24.6/53.9	56.8/52.1	14.5/44.6
	Hu et al.	54.1/52.2	15.7/49.3	53.6/54.9	22.5/47.4	54.6/52.4	17.1/48.0
	Pixel2Cancer	60.9/57.1	12.4/43.5	59.3/59.5	20.4/40.7	59.8/56.9	13.3/41.4
kidney	real tumors	71.3/62.8	27.2/64.3	70.7/61.2	19.8/57.1	65.2/58.1	25.6/59.3
	Hu et al.	63.2/55.4	35.1/69.0	61.7/52.3	26.2/61.6	55.5/49.9	27.9/62.7
	Pixel2Cancer	73.2/65.0	13.6/40.9	73.9/63.5	15.9/45.7	67.6/60.1	14.8/42.2

DSC - dice similarity coefficient; NSD - normalized surface dice.
SD - surface distance; HD - Hausdorff distance.

Implementation: *Cellular Automata:* We conduct simulation by implementing 3D Cellular Automata package using CUDA. In the 3D space, the CA kernel size is $(3,3,3)$, with cell growth frequency set to 1 in each iteration. We employed 1024 blocks with 64 threads for multiprocessing. Compared to standard CPU implementations, our CUDA-based approach achieves a speedup of $500\times$. *Tumor segmentation:* Our segmentation code is based on the MONAI framework using U-Net [24], Swin-UNETR [9], and nnU-Net [14] as backbones. During training, random patches of size $96 \times 96 \times 96$ are cropped from 3D CT images. Models are trained for $2,000$ epochs with a base learning rate of 0.0002, using a batch size of 2 per GPU and employing linear warm-up and cosine annealing learning rate schedule. For inference, we use the sliding window strategy with 0.75 overlapping.

Hyper-parameter in Pixel2Cancer: We hereby specify the hyper-parameters when generalizing Pixel2Cancer to new organs. The quantification process requires setting the typical *highest and lowest HU values* for organs. The growth and interaction process requires setting the *invasion pressure threshold* based on organ softness and the *maximum growth steps* based on organ size. The mapping function requires setting the typical *tumor HU values* for different organs.

3.2   Results

Clinical Validation: We performed a Visual Turing Test [5] on 150 CT images, including 25 images with real tumors and 25 images with synthetic tumors in each organ test set. The results in Table 1 show performance metrics for different radiologists. Radiologist R1 (3 years experience) only achieves specificity below 30%. R2 (7 years) exhibits specificity around 50%, confusing half of the synthetic tumors. Even R3 (10 years) misidentifies 47.1% of synthetic tumors as real. This demonstrates the realistic tumor simulation achieved by Pixel2Cancer.
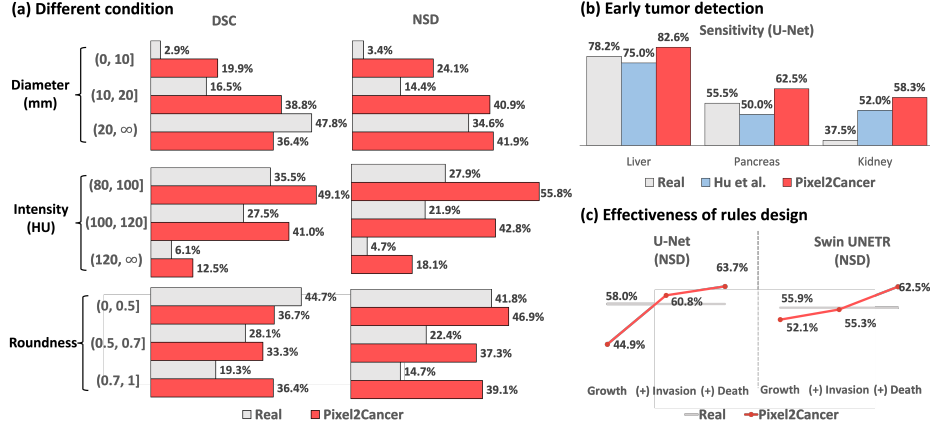
Fig. 3. (a) Different condition: We evaluate the influence of different tumor sizes, intensities, and roundness on liver tumors to assess the robustness of Pixel2Cancer. **(b) Early tumor detection:** We report the Sensitivity of U-Net tumor detection among real tumors, Hu et al. [12], and Pixel2Cancer in three organs. **(c) Effectiveness of rules:** We evaluated the performance of each generic rule on liver tumors, demonstrating the effectiveness of our rule design in tumor simulation.

Tumor Segmentation Performance: We benchmark Pixel2Cancer against the state-of-the-art modeling-based method [12] and real-tumor method. Table 2 demonstrates Pixel2Cancer's superiority in liver segmentation with DSC of 58.9%, and NSD of 63.7%. We also achieve superior performance with real tumors in the pancreas and kidney, with DSCs of 60.9% and 73.2%, respectively. Compared to models trained on real liver tumors, Pixel2Cancer outperforms by 5.7% in NSD. We also surpass Hu et al. by 4.4% in DSC and 6.1% in NSD.

Superiority in Boundary Segmentation: Pixel2Cancer generates synthetic tumors with absolutely precise tumor masks, while real data annotations often have inaccuracies at boundaries, leading to *label noise* and boundary segmentation inaccuracy. In Table 2, we apply distance metrics (NSD, SD, and HD). Our Pixel2Cancer synthesis approach surpasses real liver tumors, achieving 6.1% improvement in NSD and reductions of 24.7% in SD and 10.8% in HD. These results highlight the precision of Pixel2Cancer in boundary segmentation, indicating its potential for surgical guidance, particularly in tumor excision procedures.

Ability in Small Tumor Detection: Early detection of small tumors is essential for timely cancer diagnosis, but real datasets often lack sufficient instances due to the lack of CT scans in the early stages. Figure 3(b) and Appendix Figure 4 illustrates the performance of small tumor detection. Training solely on our synthetic tumor outperforms models trained on real tumors and Hu et al. [12], highlighting the potential of synthetic tumors for early cancer detection.

Ablation Studies: (1) We evaluated the impact of various tumor conditions on the performance of the model. Pixel2Cancer exhibits greater variability and

robustness across different conditions compared to real tumors (Figure 3(a)). (2) We evaluated the effectiveness of generic rules on liver tumors (Figure 3 (c)) Basic growth rules yielded 52.1% NSD. Incorporating interaction and invasion matched real data performance. Adding the tumor death rule outperformed real tumors by 4.9% in NSD, demonstrating the effectiveness of these generic rules.

4   Conclusion & Discussion

In this paper, we highlight the potential of Pixel2Cancer for tumor synthesis, data augmentation, small tumor detection, and accurate boundary segmentation. Despite simulating tumor behaviors and interactions with organ tissues, including the bending of vessels and the deformation of organ boundaries, there are still various mechanisms of accompanying changes in organs caused by tumors that we have not yet considered, including focal pancreas parenchymal atrophy hepatic capsular retraction [28], and splenomegaly. In future work, we plan to devise organ rules within our approach to accurately simulate changes induced by tumors when generating synthetic tumors.

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

References

1. Bilic, P., Christ, P.F., Vorontsov, E., Chlebus, G., Chen, H., Dou, Q., Fu, C.W., Han, X., Heng, P.A., Hesser, J., et al.: The liver tumor segmentation benchmark (lits). arXiv preprint arXiv:1901.04056 (2019) 6
2. Chen, Q., Chen, X., Song, H., Xiong, Z., Yuille, A., Wei, C., Zhou, Z.: Towards generalizable tumor synthesis. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (2024), https://github.com/MrGiovanni/DiffTumor 2
3. Forster, J.C., Douglass, M.J., Harriss-Phillips, W.M., Bezak, E.: Development of an in silico stochastic 4d model of tumor growth with angiogenesis. Medical Physics **44**(4), 1563–1576 (2017) 2
4. Frieboes, H.B., Jin, F., Chuang, Y.L., Wise, S.M., Lowengrub, J.S., Cristini, V.: Three-dimensional multispecies nonlinear tumor growth—ii: tumor invasion and angiogenesis. Journal of theoretical biology **264**(4), 1254–1278 (2010) 2
5. Geman, D., Geman, S., Hallonquist, N., Younes, L.: Visual turing test for computer vision systems. Proceedings of the National Academy of Sciences **112**(12), 3618–3623 (2015) 7
6. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems **27** (2014) 3
7. Han, C., Kitamura, Y., Kudo, A., Ichinose, A., Rundo, L., Furukawa, Y., Umemoto, K., Li, Y., Nakayama, H.: Synthesizing diverse lung nodules wherever massively: 3d multi-conditional gan-based ct image augmentation for object detection. In: 2019 International Conference on 3D Vision (3DV). pp. 729–737. IEEE (2019) 2

8. Harpold, H.L., Alvord Jr, E.C., Swanson, K.R.: The evolution of mathematical modeling of glioma proliferation and invasion. Journal of Neuropathology & Experimental Neurology **66**(1), 1–9 (2007) 3

9. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D.: Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In: International MICCAI Brainlesion Workshop. pp. 272–284. Springer (2022) 7

10. Heller, N., Sathianathen, N., Kalapara, A., Walczak, E., Moore, K., Kaluzniak, H., Rosenberg, J., Blake, P., Rengel, Z., Oestreich, M., et al.: The kits19 challenge data: 300 kidney tumor cases with clinical context, ct semantic segmentations, and surgical outcomes. arXiv preprint arXiv:1904.00445 (2019) 6

11. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems **33**, 6840–6851 (2020) 3

12. Hu, Q., Chen, Y., Xiao, J., Sun, S., Chen, J., Yuille, A.L., Zhou, Z.: Label-free liver tumor segmentation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7422–7432 (2023), https://github.com/MrGiovanni/SyntheticTumors 1, 3, 7, 8, 12

13. Hu, Q., Yuille, A., Zhou, Z.: Synthetic data as validation. arXiv preprint arXiv:2310.16052 (2023), https://github.com/MrGiovanni/SyntheticValidation 1

14. 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) 7

15. Jin, Q., Cui, H., Sun, C., Meng, Z., Su, R.: Free-form tumor synthesis in computed tomography images via richer generative adversarial network. Knowledge-Based Systems **218**, 106753 (2021) 2, 3

16. Kavur, A.E., Gezer, N.S., Barış, M., Aslan, S., Conze, P.H., Groza, V., Pham, D.D., Chatterjee, S., Ernst, P., Özkan, S., et al.: Chaos challenge-combined (ct-mr) healthy abdominal organ segmentation. Medical Image Analysis **69**, 101950 (2021) 6

17. Landman, B., Xu, Z., Igelsias, J., Styner, M., Langerak, T., Klein, A.: 2015 miccai multi-atlas labeling beyond the cranial vault workshop and challenge (2015), doi:10.7303/syn3193805 6

18. Luo, X., Liao, W., Xiao, J., Song, T., Zhang, X., Li, K., Wang, G., Zhang, S.: Word: Revisiting organs segmentation in the whole abdominal region. arXiv preprint arXiv:2111.02403 (2021) 6

19. Lyu, F., Ye, M., Carlsen, J.F., Erleben, K., Darkner, S., Yuen, P.C.: Pseudo-label guided image synthesis for semi-supervised covid-19 pneumonia infection segmentation. IEEE Transactions on Medical Imaging (2022) 1, 2

20. M. Cunha, G., Fowler, K.J., Roudenko, A., Taouli, B., Fung, A.W., Elsayes, K.M., Marks, R.M., Cruite, I., Horvat, N., Chernyak, V., et al.: How to use li-rads to report liver ct and mri observations. RadioGraphics **41**(5), 1352–1367 (2021) 2

21. Ma, J., Zhang, Y., Gu, S., Zhu, C., Ge, C., Zhang, Y., An, X., Wang, C., Wang, Q., Liu, X., et al.: Abdomenct-1k: Is abdominal organ segmentation a solved problem. IEEE Transactions on Pattern Analysis and Machine Intelligence (2021) 6

22. Menze, B.H., Van Leemput, K., Honkela, A., Konukoglu, E., Weber, M.A., Ayache, N., Golland, P.: A generative approach for image-based modeling of tumor growth. In: Information Processing in Medical Imaging: 22nd International Conference, IPMI 2011, Kloster Irsee, Germany, July 3-8, 2011. Proceedings 22. pp. 735–747. Springer (2011) 2

23. Rieger, H., Welter, M.: Integrative models of vascular remodeling during tumor growth. Wiley Interdisciplinary Reviews: Systems Biology and Medicine **7**(3), 113–129 (2015) 5
24. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 234–241. Springer (2015) 7
25. Roth, H., Farag, A., Turkbey, E.B., Lu, L., Liu, J., Summers, R.M.: Data from pancreas-ct (2016). https://doi.org/10.7937/K9/TCIA.2016.TNB1KQBU 6
26. 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. arXiv preprint arXiv:1902.09063 (2019) 6
27. Subramanian, S., Gholami, A., Biros, G.: Simulation of glioblastoma growth using a 3d multispecies tumor model with mass effect. Journal of mathematical biology **79**, 941–967 (2019) 2
28. Tan, G.X.V., Miranda, R., Sutherland, T.: Causes of hepatic capsular retraction: a pictorial essay. Insights into Imaging **7**(6), 831–840 (2016) 9
29. Tanase, M., Waliszewski, P.: On complexity and homogeneity measures in predicting biological aggressiveness of prostate cancer; implication of the cellular automata model of tumor growth. Journal of surgical oncology **112**(8), 791–801 (2015) 3
30. Wolfram, S.: Statistical mechanics of cellular automata. Reviews of modern physics **55**(3), 601 (1983) 2
31. Wyatt, J., Leach, A., Schmon, S.M., Willcocks, C.G.: Anoddpm: Anomaly detection with denoising diffusion probabilistic models using simplex noise. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 650–656 (2022) 2

Table 3. Early tumor performance. In this evaluation, tumors are generally small or medium-sized, presenting predominantly challenging cases for tumor segmentation. This challenge is particularly pronounced in the pancreas and kidney, where tumors are more difficult to segment compared to liver tumors.

| organ	tumor	label/non-label	U-Net DSC / NSD (%) ↑	U-Net SD / HD (mm) ↓	Swin UNETR DSC / NSD (%) ↑	Swin UNETR SD / HD (mm) ↓
liver	real tumors	101/0	46.7/48.0	23.2/61.4	43.5/45.9	21.3/57.8
	Hu et al. [12]	0/116	44.5/47.6	23.8/58.8	42.3/46.5	22.9/55.4
	Pixel2Cancer	0/116	**47.2/52.9**	**17.9/52.4**	**45.4/51.6**	**18.7/51.3**
pancreas	real tumors	96/0	34.3/33.9	22.2/**47.7**	28.3/33.6	22.3/45.9
	Hu et al. [12]	0/104	27.3/28.0	28.4/52.0	25.8/27.3	29.5/50.9
	Pixel2Cancer	0/104	**36.5/35.9**	**21.9**/52.2	**31.5/34.6**	**21.7/41.8**
kidney	real tumors	96/0	17.9/16.1	102.8/150.6	**29.4**/27.6	93.8/136.8
	Hu et al. [12]	0/120	14.8/17.4	**70.1/112.9**	23.5/25.9	83.2/**122.3**
	Pixel2Cancer	0/120	**18.1/19.3**	85.2/131.7	28.6/**28.2**	**80.9**/127.2
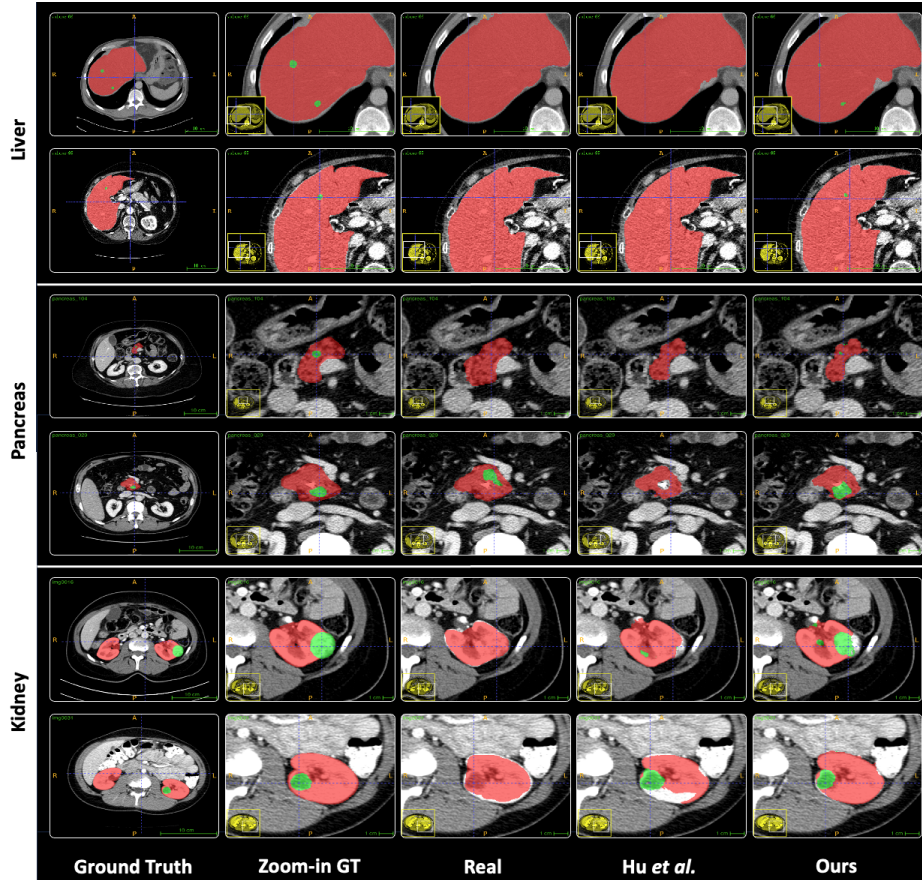
Fig. 4. Examples of early tumor detection. Qualitative visualizations of segmentation models for liver, pancreas, and kidney tumor detection.
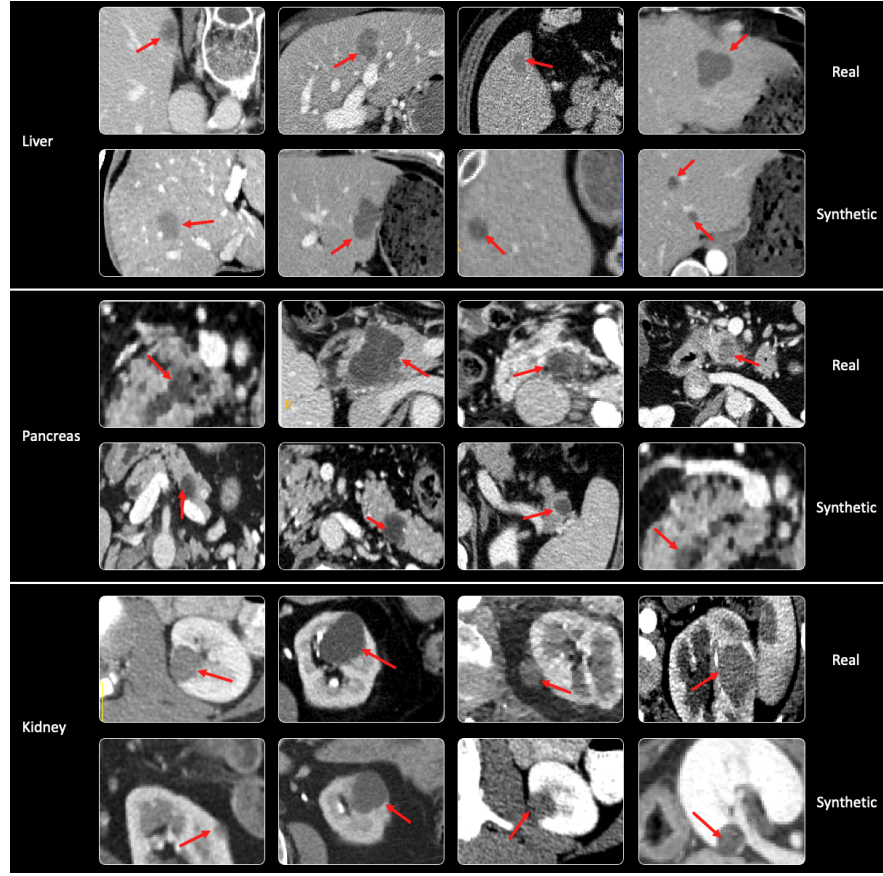
Fig. 5. Examples of Visual Turning Test. Tubular results are presented in Table 1. Our Pixel2Cancer can be used to augment available healthy CT volumes.
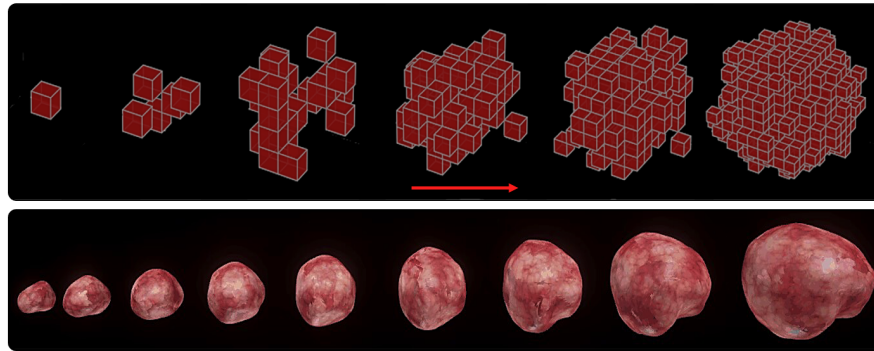
Fig. 6. Tumor development visualization. This rendering visualization depicts the iteration of Pixel2Cancer and the progression of tumors. We illustrate iterations and the gradual development of tumors.