

RadGPT: Constructing 3D Image-Text Tumor Datasets

Pedro R. A. S. Bassi^{1,2,3} Mehmet Can Yavuz⁴ Kang Wang⁴ Xiaoxi Chen⁵ Wenxuan Li¹
 Sergio Decherchi³ Andrea Cavalli^{2,3,6} Yang Yang⁴ Alan Yuille¹ Zongwei Zhou^{1,*}

¹Johns Hopkins University ²University of Bologna ³Italian Institute of Technology
⁴University of California, San Francisco ⁵University of Illinois Urbana-Champaign
⁶École Polytechnique Fédérale de Lausanne

Code, dataset, and models: <https://github.com/MrGiovanni/RadGPT>

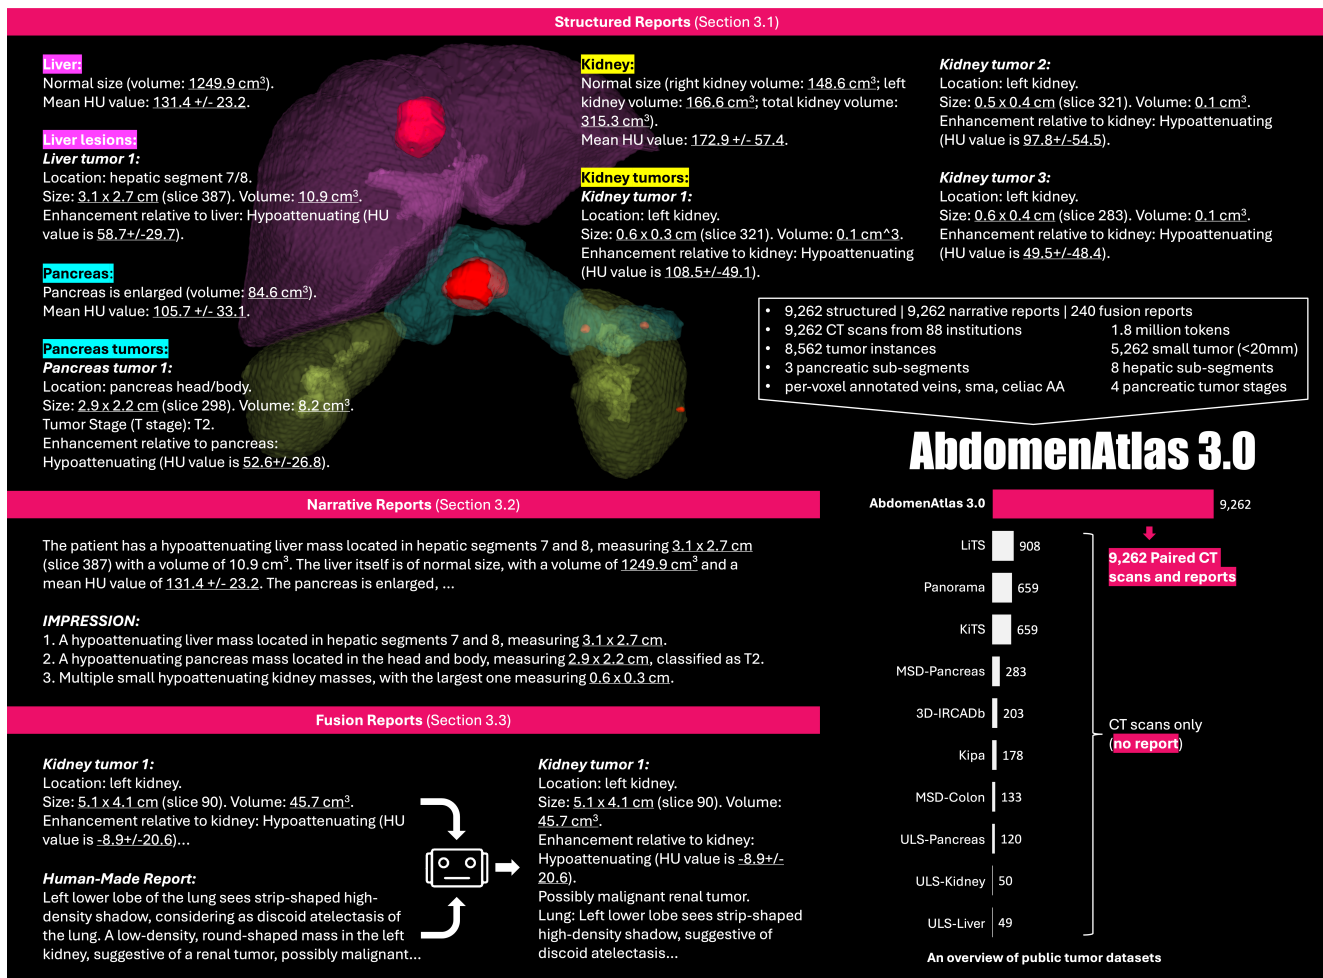


Figure 1. **AbdomenAtlas 3.0** is a large-scale, image-text tumor dataset of 9,262 3D CT scans. Each CT scan has per-voxel tumor annotations and reports, including 3,036 liver tumors, 354 pancreatic tumors and 4,239 kidney tumor, 5,262 of which are small tumor (≤ 2 cm). In addition, AbdomenAtlas 3.0 provides detailed annotations for pancreatic cancer staging (T1–T4), as well as per-voxel segmentations of liver sub-segments (1–8) and pancreatic sub-segments (head, body, and tail).

Abstract

With over 85 million CT scans performed annually in the United States, creating tumor-related reports is a challenging and time-consuming task for radiologists. To address this need, we present RadGPT, an Anatomy-Aware Vision-Language AI Agent for generating detailed reports from CT scans. RadGPT first segments tumors, including benign cysts and malignant tumors, and their surrounding anatomical structures, then transforms this information into both structured reports and narrative reports. These reports provide tumor size, shape, location, attenuation, volume, and interactions with surrounding blood vessels and organs. Extensive evaluation on unseen hospitals shows that RadGPT can produce accurate reports, with high sensitivity/specificity for small tumor (<2 cm) detection: 80/73% for liver tumors, 92/78% for kidney tumors, and 77/77% for pancreatic tumors. For large tumors, sensitivity ranges from 89% to 97%. The results significantly surpass the state-of-the-art in abdominal CT report generation.

RadGPT generated reports for 17 public datasets. Through radiologist review and refinement, we have ensured the reports' accuracy, and created the first publicly available image-text 3D medical dataset, comprising over 1.8 million text tokens and 2.7 million images from 9,262 CT scans, including 2,947 tumor scans/reports of 8,562 tumor instances. Our reports can: (1) localize tumors in eight liver sub-segments and three pancreatic sub-segments annotated per-voxel; (2) determine pancreatic tumor stage (T1-T4) in 260 reports; and (3) present individual analyses of multiple tumors—rare in human-made reports. Importantly, 948 of the reports are for early-stage tumors.

1. Introduction

Each year, over 85 million CT scans are performed in the United States [55, 62], growing 6% per year, and significantly outpacing the 0.7% annual growth rate of the medical imaging workforce [16]. This disparity puts radiologists under significant time pressure, making it challenging to generate detailed, accurate radiology reports. These reports involve identifying tumors, measuring them, determining their location, classifying types¹, and assessing surrounding tissue interactions—a process that is labor-intensive, time-consuming, and prone to human error. Recent studies showed that radiologists have a sensitivity rate of about 65–85% for detecting small liver tumors [52] and as low as 33–44% for small pancreatic tumors [24], leading to frequent oversight. These missed detections can delay timely treatment and diminish the quality of AI training, as many small

¹We use the term 'tumor' to refer to an abnormal mass of tissue, which may be malignant (cancer) or benign (e.g., cysts).

tumors go unrecorded in radiology reports. Automated report generation is needed to improve the accuracy and efficiency of tumor detection beyond human-made reports.

Automated report generation from abdominal CT scans is particularly important as many high-mortality cancer like liver, pancreatic, and colorectal cancer originate in the abdomen and require early detection and precise staging [13, 19, 40, 41, 70]. However, as of this writing, there are no publicly available datasets or models for automated report generation that detail tumor-related findings using abdominal CT scans². Pre-existing report generation approaches, such as those based on Vision-Language Models (VLMs), are mainly designed for 2D images [12, 35, 54, 59, 67, 68] or Chest CT scans [5, 27, 65, 66], using different imaging protocols from abdominal scans and relying on weak annotations from CLIP-like [50] image-text pairs, which limits their ability to provide precise location, measurement, and tracking for individual tumors (evidenced in Table 2).

To address this gap, we present Radiology Generative Pre-trained Transformers, **RadGPT**, an Anatomy-Aware Vision-Language AI Agent for generating detailed reports from abdominal CT scans. This approach seeks to improve the accuracy of tumor quantification and bring consistency to report generation. RadGPT can convert per-voxel annotation of tumors and organs into two types of reports: (1) *structured* reports and (2) *narrative* reports. As shown in Figure 1, the structured reports follow a radiologist-provided template, and RadGPT can adapt this reports into narrative reports that match the style (word choice and text organization) of a target institution (hospital or clinic). Both formats have been validated by radiologists and offer comprehensive and detailed information, such as tumor size, location, attenuation, volume, descriptions of adjacent organs (e.g., fatty liver, enlarged spleen), patient demographics, contrast enhancement protocols, and pancreatic tumor T stage, which consider tumor size and interaction with blood vessels, key considerations for surgical removal. In addition to generating accurate reports, the proposed RadGPT offers two key functions:

- Enhancing human-made reports:** RadGPT can leverage per-voxel tumor segmentation and human-made reports or clinical notes (i.e., informal and short reports) to generate *fusion reports* (§3.3). They inherit the precision and detail of segmentation (e.g., multiple precise tumor size and volume measurements), and the diagnostic breadth of human-made reports/notes, encompassing multiple conditions beyond tumors (§4.5). This synergy enhances clarity and enriches human-made reports/notes with quantitative insights.
- Diagnostic evaluation of AI-made reports:** RadGPT

²Two concurrent studies are underway, but the datasets and models from Merlin [9] are not publicly available, and M3D [4] provides CT scans in a 2D PNG/JPEG format (0–255), unsuitable for clinical use (§2)

introduces a quantitative diagnostic assessment for report generation—a feature missing in current approaches for abdominal CT report generation (§2). First, a LLM extracts numerical and categorical labels (e.g., tumor presence/absence) from AI- and human-made reports, with 96% accuracy in zero-shot (Figure 3). By comparing the labels from AI- and human-made reports, we can evaluate AI’s diagnostic accuracy using metrics such as sensitivity and specificity (§3.4).

Finally, we used RadGPT to generate reports for 17 public datasets, and through annotation, review, and refinement by 12 radiologists, we created **AbdomenAtlas 3.0** (Figure 1), the largest public image-text abdominal CT dataset. It contains 1,843,262 text tokens from reports and 2,789,975 images from 9,262 CT scans. Each scan features structured and narrative reports detailing tumors’ sizes, shapes, locations and appearances. The dataset includes a total of 2,947 reports describing 8,562 tumor instances across the liver (929 reports, 3,036 instances), pancreas (344 reports, 354 instances), and kidneys (1,674 reports, 4,239 instances), all per-voxel annotated. Additionally, 6,061 reports document no tumors in these organs, serving as control. Notably, a large number of reports documenting tumors document *small* tumors (≤ 2 cm)—347 (37.4%) for liver, 83 (24.1%) for pancreas, and 466 (27.8%) for kidney—critical for early cancer detection but rare and challenging to collect. Tumors are localized using per-voxel annotated liver sub-segments (1–8) and pancreatic sub-segments (head, body, tail). AbdomenAtlas 3.0 includes 260 reports with pancreatic tumor *staging* (T1–T4), enabled by per-voxel annotations of key vessels (e.g., SMA). AbdomenAtlas 3.0 is the first public dataset with pancreatic cancer staging and pancreas sub-segments. It also provides 240 human-AI fusion reports, combining our structured/narrative reports with clinical notes (see §3.3). Metadata like contrast enhancement, slice thickness, and in-plane spacing further enhance AbdomenAtlas 3.0.

2. Related Work

Per-voxel tumor annotations are scarce. Public abdominal CT datasets typically focus on one tumor type—such as liver [8], pancreas [3], and kidney [25] tumors—and present only a few hundreds per-voxel tumor annotations (Table 1). This limited quantity undermines robust AI training and evaluation. Thus, supported by radiologists, we *tripled* the number of per-voxel tumor annotations in the 17 public datasets inside AbdomenAtlas 3.0 (Table 1).

Even scarcer than per-voxel tumor annotations are reports. No public abdominal CT dataset has real-world radiology reports (i.e., from a medical institution). Actually, only one public abdominal CT dataset, M3D-Cap [4], has some text annotations. They are captions, sourced from the Radiopedia website [22], along with CT scans. However,

these scans do not meet clinical standards. They are 2D JPG/PNG image series instead of standard 3D NIFTI or DICOM volumes. Thus, they lack information that is essential for precise tumor measurement by AI, like inter-image spacing and Hounsfield units (HU) [70]. Conversely, the CT scans in AbdomenAtlas 3.0 have this information, being sourced from 88 medical institutions in standard format.

Due to the scarcity of abdominal CT reports in public datasets, only two models for abdominal CT report generation exist: M3D [4] (public) and Merlin [9] (unreleased). However, both were solely evaluated with text similarity metrics, like BLEU and ROUGE [39]. These standard LLM evaluation metrics can be highly influenced by changes in the report style, even when underlying diagnoses do not change (§4.4). In contrast, RadGPT enables evaluating AI-made reports using diagnostic sensitivity and specificity, which are clinically relevant metrics [10, 64].

3. AbdomenAtlas 3.0 & RadGPT

Table 1 shows advantages of **AbdomenAtlas 3.0** over the 17 public datasets. It is the only dataset with reports, and has three report types: structured reports (§3.1), narrative reports (§3.2), and fusion reports (§3.3). All reports were generated by **RadGPT**, an anatomy-aware vision-language AI agent illustrated in Figure 2. First, RadGPT transformed the ground-truth per-voxel annotations of tumors and organs into structured reports, using deterministic algorithms. These algorithms ensure the diagnoses in the report align with the ground-truth per-voxel annotations. Then, RadGPT used LLMs to convert the structured reports into narrative reports, and to create *fusion reports*, which fuse structured reports with clinical notes.

3.1. Creating Structured Reports

Structured reports use a radiologist-designed template, enhancing clarity and aiding medical decisions [1] (Figure 1). To fill out this template, RadGPT performs key tasks: (1) Sub-segments organs to localize tumors (§3.1.1). (2) Measures tumor size, volume, and attenuation from segmentations (§3.1.2). (3) Uses tumor and blood vessel segmentations for staging (§3.1.3).

3.1.1. Sub-segment Organs to Localize Tumors

Human-made reports use organ sub-segments to locate tumors. Location is key for prognosis, tracking tumor progression, and treatment planning. E.g., the possibility of tumor removal depends on its location [57]. To locate liver and pancreas tumors in structured reports, RadGPT sub-segments the organs and checks which sub-segments intersect with the tumor. Besides releasing the first public dataset with pancreas sub-segments, we release the *first* public AI models that sub-segment pancreas and liver.

dataset	CTs	institutions	countries	annotated liver tumors	annotated pancreatic tumors	annotated kidney tumors
FLARE'23 [2022] [link]	4,100	35	1	0 → 376	0 → 34	0 → 763
KiTS'23 [2020] [link]	489	1	1	0 → 1	0	452
LiTS [2019] [link]	131	7	5	50	0	0
TCIA-Pancreas-CT [2015] [link]	42	1	1	0	0	0
CT-ORG [2020] [link]	140	8	6	0 → 41	0	0 → 17
Trauma Det. [2023] [link]	4,714	23	13	0 → 46	0 → 15	0 → 38
BTCV [2015] [link]	47	1	1	0	0	0
CHAOS [2018] [link]	20	1	1	0	0 → 1	0
AbdomenCT-1K [2021] [link]	1,050	12	7	0 → 82	0 → 100	0 → 113
MSD CT Tasks (6) [2021] [link]	945	1	1	251	191	0 → 224
WORD [2021] [link]	120	1	1	0 → 28	0	0 → 26
AMOS [2022] [link]	200	2	1	0 → 54	0 → 3	0 → 41
AbdomenAtlas 3.0 (ours)	9,262	88	19	301 → 929	191 → 344	452 → 1,674

dataset	liver sub-segments	pancreas sub-segments	peripancreatic blood vessels ³	tumor stage	radiology reports	text tokens
FLARE'23 [2022] [link]	✗	✗	✗	✗	0	0
KiTS'23 [2020] [link]	✗	✗	✗	✗	0	0
LiTS [2019] [link]	✓	✗	✗	✗	0	0
TCIA-Pancreas-CT [2015] [link]	✗	✗	✗	✗	0	0
CT-ORG [2020] [link]	✗	✗	✗	✗	0	0
Trauma Det. [2023] [link]	✗	✗	✗	✗	0	0
BTCV [2015] [link]	✗	✗	✗	✗	0	0
CHAOS [2018] [link]	✗	✗	✗	✗	0	0
AbdomenCT-1K [2021] [link]	✗	✗	✗	✗	0	0
MSD CT Tasks (6) [2021] [link]	✗	✗	✗	✗	0	0
WORD [2021] [link]	✗	✗	✗	✗	0	0
AMOS [2022] [link]	✗	✗	✗	✗	0	0
AbdomenAtlas 3.0 (ours)	✓	✓	✓	✓	18,524	1,843,262

→ represents the number of CT scans with tumor annotations in the original dataset, followed (→) by our updated number of CT scans with tumor annotations, including the additional annotations AbdomenAtlas 3.0 provided with radiologist support.

Table 1. **Besides being the only public abdominal CT dataset with paired radiology reports, AbdomenAtlas 3.0 offers 3.12× more annotated tumors than the combined total of its constituent datasets.** The table highlights how AbdomenAtlas 3.0 enhances public datasets with reports and tumor annotations. It includes 936 CT scans with liver tumors, 342 with pancreatic tumors, and 1,416 with kidney tumors, most newly annotated with radiologist support. Each sample includes per-voxel annotations and reports. AbdomenAtlas 3.0 is also the first dataset to provide per-voxel segmentations of pancreas sub-segments and peripancreatic blood vessels.

For liver sub-segmentation, we propose using ground-truth per-voxel annotations of the entire liver to help the AI find liver sub-segments. First, we offset the liver intensity (by 200 HU), following its ground-truth per-voxel annotation. Second, using these CT scans with offsets as input, we trained an nnU-Net [29] for liver sub-segmentation. The sub-segments follow the Couinaud standard [18], which divides the liver into eight sub-segments that can be independently removed in surgeries. Couinaud sub-segment annotations are publicly available for 131 LiTS CT scans [8, 69], which we used for training. Given the small size of this dataset, we fine-tuned an nnU-Net pre-trained on 9,262 CT scans in AbdomenAtlas 1.1 [36]. After training, we inferred the nnU-Net on AbdomenAtlas 3.0. The HU value offsetting ensured the precise alignment between the generated sub-segments and the existing liver ground-truth per-voxel annotations. AbdomenAtlas 3.0 is the *second* [69] but the *largest public dataset with liver sub-segments*.

For pancreas sub-segmentation, there is no public dataset with per-voxel annotations of pancreas head, body, and tail. Thus, to sub-segment the pancreas, we used a radiology landmark: the superior mesenteric artery (SMA). We trained an nnU-Net to segment the SMA (using private

data), and developed a deterministic algorithm that uses the segmented SMA to sub-segment the pancreas (algorithm 1). The algorithm follows standard radiology guidelines. First, it uses the SMA to find the pancreatic neck, since it curves around the SMA. The neck locates the head-body boundary. Then, the body-tail boundary is set at the midpoint along their length [58]. AbdomenAtlas 3.0 is the *first public dataset with pancreas sub-segment annotations*.

3.1.2. Segment Tumors and Measure Like Radiologists

Standardized tumor measurement is key for cancer prognosis and treatment planning [38, 47]. Human-made radiology reports commonly use the World Health Organization (WHO) tumor measurement standard, which measures tumors with two diameters: the longest tumor diameter in any axial plane (D), and its perpendicular diameter in the same plane (d) [47]. Thus, to measure tumors like radiologists in our structured reports, RadGPT first needs tumor segmentations. AbdomenAtlas 3.0 presents ground-truth segmentations of liver, kidney and pancreas tumors. For fully-automated report generation, RadGPT segments tumors with DiffTumor, a public AI model [14, 15]. From segmentations, RadGPT extracts tumor measurements us-

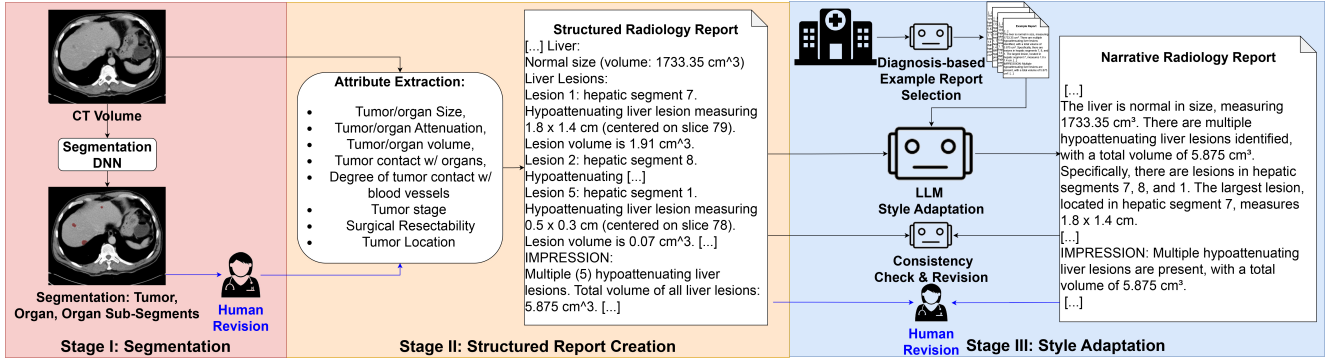


Figure 2. RadGPT adopts 3-stage pipeline for report generation. Blue arrows denote quality assurance steps used for creating AbdomenAtlas 3.0 but omitted in fully-automated generation. **Stage I. Segmentation.** We utilize DiffTumor [14] and a proprietary-trained nnU-Net [29] to segment 26 anatomical structures important for cancer detection and staging³. For AbdomenAtlas 3.0, we asked radiologists to revise segmentations, or substituted them by the ground-truth from public datasets when available. **Stage II. Structured Report Generation.** Deterministic algorithms (§3.1.1-3.1.3) extract radiologist-selected attributes (important for cancer detection, staging and treatment planning) from CT scans and per-voxel segmentations to fill a template created by radiologists, generating structured reports for liver, kidney, and pancreatic cancer assessment. The rule-based deterministic algorithms ensure the reports are fully coherent with the segmentations and explainable. **Stage III. Style Adaptation.** A LLM adapts structured reports to a target hospital’s narrative style using in-context learning with example reports from the hospital. An new LLM-based example selection prioritizes reports with similar diagnoses (§3.2). The LLM is instructed to maintain medical information accuracy and perform double checks for consistency. For AbdomenAtlas 3.0, radiologists conducted final report revisions. Additionally, we use an LLM to fuse structured reports with human-made reports, creating comprehensive fusion reports that combine segmentation-based precision with the generality of human-made reports (Section 3.3).

ing a deterministic algorithm, which we designed to implement the WHO measuring standard (Alg. 2). Additionally, our structured reports present tumor and organ *volume* and *attenuation* (HU values), both extracted from segmentation. AbdomenAtlas 3.0 has ground-truth segmentations for multiple organs³. For fully-automated report generation, RadGPT segments organs with an nnU-Net trained in AbdomenAtlas 1.1 [7, 37]. Organ volumes allows our structured reports to diagnose enlarged organs, and attenuation allows diagnosing fatty liver (average HU < 40 [32]) and pancreas (pancreas-to-spleen attenuation < 0.7 [21])—a condition related to diabetes and pancreatic cancer [21]. Meanwhile, tumor attenuation helps diagnose tumor type.

3.1.3. Stage Pancreatic Cancer using Segmentation

Cancer T stage indicates the tumor size and its interaction with nearby structures. It is vital for surgical planning and survival, especially for pancreatic adenocarcinoma (PDAC), a highly aggressive cancer [1]. However, staging is time-consuming for radiologists. To provide PDAC T stage in our structured reports, we measure the tumor (§3.1.2) and analyze its interaction with nearby blood vessels (SMA, CHA, CA, SA³) [1]. Thus, RadGPT first segments these

³AbdomenAtlas 3.0 is the first dataset with per-voxel annotations for the blood vessels key for pancreatic tumor staging: the celiac axis (CA), superior mesenteric artery (SMA), superior mesenteric vein (SMV), common hepatic artery (CHA), and portal vein. It also has per-voxel annotations for other 22 structures important for cancer detection/staging: liver tumors, kidney tumors, pancreas tumors, liver, kidney, pancreas, spleen, adrenal glands, stomach, duodenum, bile duct, intestines, aorta, and postcava.

vessels and tumors, using an nnU-Net trained on private data and DiffTumor, respectively. Then, a deterministic algorithm uses the segmentations to measure the tumor-vessel angle of contact (Alg. 3). Large angles (>180°) make surgical removal challenging, increasing stage. For interpretability, the reports justify stages with tumor size and tumor-vessel degree of contact. AbdomenAtlas 3.0 is *first public CT dataset with PDAC T stage annotations*.

3.2. Creating Narrative Reports

Structured reports use rigid templates to improve clarity and clinical decision-making [1]. However, rigid templates may conflict with the reporting style of an institution. Thus, RadGPT can create narrative reports that mimic the style of a target institution. In AbdomenAtlas 3.0, they mimic human-made reports at UCSF (Figure 1). The narrative reports are created through style adaptation with in-context learning: we provide a pre-trained LLM⁴ with a structured report and 10 human-made reports from the target institution, and the LLM adapts the structured report to the style of the human-made reports. We ask the LLM *not* to change any diagnosis. Thus, narrative reports contain all the detailed information found in structured reports (§3.1).

However, the style of human-made reports varies significantly with its diagnoses. E.g., pancreatic tumor reports differ from liver tumor or pneumonia reports [48]. Thus, we verify diagnoses to give the LLM human-made reports

⁴Llama-3.1 with 70B parameters and AWQ quantization [20]

Large Tumors >2 cm								
Model	Liver Tumor (HCC)		Kidney Tumor (RCC)		Pancreas Tumor (PDAC)		CRC Liver Metastases	
	Sensitivity (%)	Specificity (%)	Sensitivity (%)	Specificity (%)	Sensitivity (%)	Specificity (%)	Sensitivity (%)	Specificity (%)
CT2Rep	0.0 (0/301)	100.0 (244/244)	10.0 (22/219)	98.0 (239/244)	3.8 (4/105)	96.7 (236/244)	9.4 (2/21)	98.1 (155/158)
M3D	14.6 (44/301)	88.9 (217/244)	9.1 (20/219)	92.6 (226/244)	4.8 (5/105)	96.3 (235/244)	12.1 (7/58)	85.4 (135/158)
RadGPT (ours)	89.4 (269/301)	73.4 (179/244)	97.3 (213/219)	78.3 (191/244)	91.4 (96/105)	76.6 (187/244)	100.0 (58/58)	70.3 (111/158)
Small Tumors \leq 2 cm								
Model	Liver Tumor (HCC)		Kidney Tumor (RCC)		Pancreas Tumor (PDAC)		CRC Liver Metastases	
	Sensitivity (%)	Specificity (%)	Sensitivity (%)	Specificity (%)	Sensitivity (%)	Specificity (%)	Sensitivity (%)	Specificity (%)
CT2Rep	0.0 (0/142)	100.0 (244/244)	4.0 (2/50)	98.0 (239/244)	2.1 (8/385)	96.7 (236/244)	3.4 (2/58)	98.1 (155/158)
M3D	6.3 (9/142)	88.9 (217/244)	2.0 (1/50)	92.6 (226/244)	2.8 (11/385)	96.3 (235/244)	14.3 (3/21)	85.4 (135/158)
RadGPT (ours)	79.6 (113/142)	73.4 (179/244)	92.0 (46/50)	78.3 (191/244)	76.9 (296/385)	76.6 (187/244)	100.0 (21/21)	70.3 (111/158)

Table 2. **In tumor detection, fully-automated reports by RadGPT surpass reports created by end-to-end report generation models.** Here, we use RadGPT as a fully-automated two-stage pipeline (Figure 2). RadGPT surpasses the two available end-to-end 3D CT report generation models: M3D [4], the state-of-the-art in abdominal CT report generation, and CT2Rep [23], the state-of-the-art in chest CT report generation, which we adapted to abdominal CT scans. We performed *out-of-distribution* (OOD) evaluation at UCSF, a hospital not seen in training. The results indicate that per-voxel segmentation (step 1 in the RadGPT pipeline) may strongly benefit report generation.

with the correct style. We first use an LLM to categorize human-made reports according to their tumors (liver, pancreas, kidney, or none). Then, when adapting a structured report into narrative, we give the LLM human-made reports with the same tumor as the structured report (e.g., liver).

After adapting a structured report to a narrative report, the LLM performed a quality check. It extracted diagnoses and quantitative information (e.g., tumor size and stage) from both reports and checked for consistency. We prompted the LLM to correct in the narrative report any information diverging from the structured report, and to remove any diagnosis not present in the structured report.

3.3. Creating Human-AI Fusion Reports

The structured and narrative reports in AbdomenAtlas 3.0 focus on tumors. Our reports can quantify multiple tumors in a CT scan so they are much more detailed than human-made reports (§4.5). However, human-made reports cover multiple diagnoses unrelated to tumors. To get the best of both worlds, RadGPT prompts a zero-shot LLM⁴ to fuse the details in structured reports with the many diagnoses in human-made reports or clinical notes (Figures 11 and 1), generating *fusion reports*. AbdomenAtlas 3.0 has 240 of them: 209 used clinical notes for TotalSegmentator CT scans [4], and 31 used notes from collaborating radiologists.

3.4. Evaluating Diagnoses in AI-Made Reports

We propose an automatic strategy to evaluate the clinical utility of AI-made reports. Specifically, it evaluates whether the diagnoses in the reports are correct. We prompt an LLM⁴ (see prompts in Appendix B.3) to identify whether reports mention liver, kidney, or pancreatic tumors. Then, we convert the LLM answers into categorical labels. We create labels for both AI-made reports and human-made reports (ground-truth). By comparing them, we calculate tumor detection sensitivity and specificity for the AI-made reports. Besides producing clinically relevant metrics (§4.4),

this evaluation strategy is practical: using zero-shot inference, it does not need fine-tuning and is easily adaptable to multiple hospitals. Also, with a simple prompt modification, it can evaluate multiple diagnoses beyond tumors.

4. Experiment & Result

In AbdomenAtlas 3.0, RadGPT leverages ground-truth tumor annotations, ensuring reports match the accuracy of these annotations. However, this section evaluates RadGPT as a fully automated two-stage pipeline, without manual annotations, as shown in Figure 2. **Stage I:** AI models segment tumors and organs, using DiffTumor [14, 15] for tumors and SuPreM [37] for organs. **Stage II:** RadGPT converts these segmentations into structured reports. We compare our two-stage approach (RadGPT) to end-to-end report generation approaches (M3D [4] and CT2Rep [23]). To ensure realistic evaluation [6, 7], we test on out-of-distribution hospitals not included in the training data.

4.1. RadGPT Enhances Tumor Detection

To enable automated evaluation on a large test dataset, we use an LLM (Llama-3.1) to assess the reports generated by RadGPT (§3.4). First, we verify the LLM’s ability to determine whether a report indicates the presence of a tumor. As shown in Figure 3, the zero-shot LLM achieves 96% accuracy on this task, demonstrating its reliability for evaluating tumor detection sensitivity and specificity.

The LLM evaluation showed that the reports generated by RadGPT strongly surpassed existing abdominal CT report generation models. Table 2 compares RadGPT reports to M3D [4], the only public AI model for abdominal CT report generation, and to CT2Rep [23], the first AI model for report generation in chest CT scans. Here, we adapted CT2Rep to the abdomen region (Appendix B.1). Both CT2Rep and M3D had difficulty detecting tumors (low sensitivity), and RadGPT strongly outperformed them




Figure 3. **Confusion matrices.** (a) A zero-shot LLM (Llama 3.1) has 96% accuracy and 0.953 F1-score in determining whether radiology reports indicate tumors. Thus, the LLM can accurately calculate tumor detection sensitivity and specificity for AI-made reports (§3.4). The LLM accuracy rivals established labelers, like those in CheXpert [28] and CheX-ray14 [61]. Results were manually evaluated by radiologists. They used 447 randomly selected UCSF reports including kidney, pancreas, and liver tumors. (b) PDAC staging confusion matrix for RadGPT, the first public DNN for staging abdominal CT tumors. Results consider a private dataset with ground-truth tumor stage annotations ($N=42$).

for small and large tumors in the liver, pancreas, and kidneys. This performance difference shows the benefits of using segmentation to improve report generation: DiffTumor produces accurate tumor segmentations, which RadGPT translates into reports. By releasing AbdomenAtlas 3.0, the first abdominal CT dataset with triplets of CT scans, reports, and per-voxel annotations, our objective is to catalyze further research on segmentation-driven report generation.

RadGPT is the first AI model to perform cancer staging on abdominal CT. Figure 3 shows the result of RadGPT on the staging of pancreatic adenocarcinoma. The RadGPT fully-automated reports achieved accuracy of 71.43% in determining tumor T stages 1 to 3. The results show that AI is a promising tool for assisting cancer staging, a key but time-consuming task for radiologists.

4.2. RadGPT Accurately Measures Tumor Size

An expert radiologist manually evaluated structured reports generated by RadGPT. He analyzed each reported tumor, evaluating its measurement and checking if the tumor is a false-positive (tumor not present in the CT volume) or a true-positive (present). The radiologist deemed 75.6% of the tumors reported by RadGPT true-positives, and 93.8% of them were correctly measured (Table 3). RadGPT only made measuring mistakes for pancreatic tumors (PDAC), but even the radiologist could not measure 3 PDACs.

4.3. RadGPT Locates Tumors in Organs

RadGPT uses organ sub-segments to locate tumors. It achieved a Dice similarity coefficient (DSC) of 0.85 in segmenting eight liver sub-segments, according to the test set

Tumor	# Correct Tumor Detection	# Tumor Detection	# Correct Tumor Size	Tumor Detection Precision (%)	Size Accuracy (%)
HCC	12	13	13	92.3	100.0
PDAC	8	16	6	50.0	75.0
RCC	11	12	11	91.7	100.0
Total	31	41	30	75.6	93.8

Table 3. **RadGPT achieved 75.6% tumor detection precision and 93.8% tumor size measurement accuracy, according to manual evaluation.** A radiologist evaluated the reports RadGPT created for 23 CT scans from UCSF. A reported tumor measurement was considered correct if it deviated by 10% or less from the radiologist’s measurement (both use the WHO measuring standard [47]). As manual evaluation is time-consuming, the radiologist evaluated 23 reports. Using an LLM for automatically evaluating tumor measurements is challenging: it requires pairing tumors in AI-made reports and ground-truth reports. E.g., if both reports mention 3 tumors in liver sub-segment 2, it is hard to understand which measurements we should compare.

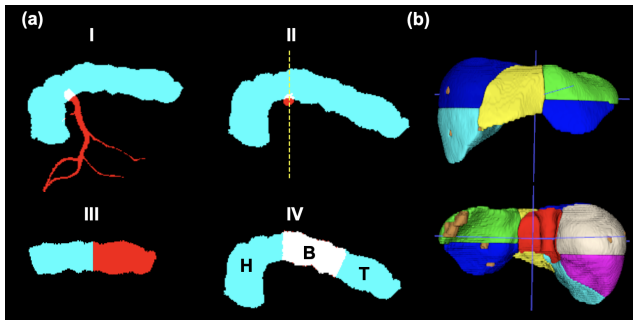


Figure 4. **Pancreas and liver sub-segments.** (a) RadGPT segments the pancreas based on radiology standards [58]. (I-II) the SMA separates the pancreas head (H) from the body (B), and (III) the remaining pancreas is divided at its midpoint into the body and tail (T). (b) Our liver sub-segmentation model achieved a DSC of 0.85 in segmenting eight liver sub-segments on a public test set [69]. Sub-segments are in different colors and tumors in brown. Sub-segments are essential for RadGPT to localize tumors.

from Zhang *et al.* [69]. For pancreas sub-segmentation, we do not have a ground-truth or dataset for testing, because AbdomenAtlas 3.0 is the first public dataset to present pancreas sub-segments (head, body, and tail). However, our algorithm to sub-segment the pancreas closely follows radiologist-accepted standards (Figure 4), and we asked radiologists to qualitatively evaluate our annotations.

4.4. RadGPT Enables Diagnostic Evaluation

In Table 4, we use standard text similarity metrics (common in LLM evaluation) to evaluate the reports generated by RadGPT, CT2Rep, and M3D. RadGPT achieves the highest scores in all metrics. These results align with the superiority of RadGPT in our diagnostic evaluation (Table 2). Thus, diagnostic accuracy may improve LLM metrics.

Model	BLEU	METEOR	ROUGE-1	R.-2	R.-L	BERT
M3D	0.0	2.5	6.9	1.0	4.5	44.2
CT2Rep	0.0	4.4	7.7	0.7	5.3	51.3
RadGPT-S	1.0	11.0	14.8	1.8	8.9	54.9
RadGPT-N	0.7	11.2	24.7	2.8	12.4	53.4

Table 4. **Report style impacts standard evaluation metrics.** With the same dataset as Table 2, we evaluate the fully-automated reports from RadGPT with text similarity metrics. RadGPT narrative (N) and structured (S) reports differ in style only, having the same sensitivity and specificity for tumor detection. However, by mimicking the style of ground-truth reports (§3.2), narrative reports have 67.1% higher ROUGE-1.

However, these metrics are also affected by the style of the report. Table 4 includes structured and narrative reports by RadGPT. They contain the same diagnoses and details, but different style (§3.2). In Table 4, ROUGE [39] was the metric most affected by the style variation and METEOR the least. The results show that standard LLM metrics are influenced by style, but the extent of this influence varies between metrics. Conversely, our proposed LLM-based sensitivity and specificity metrics are only influenced by the diagnostic accuracy of the reports. In addition, sensitivity and specificity provide clinicians with an objective and easily interpretable evaluation of AI-made reports, directly evaluating the clinical usefulness of a vision-language model.

4.5. RadGPT Can Enhance Human-made Reports

Human-made reports often have inconsistent documentation of critical findings, complicating clinical decision-making by referring physicians [1]. To improve consistency and interpretability, experts advocate for structured reports. AbdomenAtlas 3.0 and RadGPT address this human report problem providing structured and narrative reports.

Both types of reports provide measurements of organ and tumor volume and attenuation, rarely found in human-made reports. Of 90 randomly selected UCSF human-made reports, none had volumes or HU values. Volumetric assessments better capture tumor mass, detect size changes earlier, and improve prognostic accuracy [11, 46, 63]. Organ volume and attenuation for kidney, spleen, pancreas, and liver are valuable for monitoring and diagnosing diseases. E.g., reduced pancreatic volume can indicate diabetes [45] and low liver attenuation can indicate fat infiltration [49]. RadGPT and AbdomenAtlas 3.0 include volumes and HU values for organs and tumors. With time constraints, radiologists typically measure only the largest tumors in a CT. Of 90 human-made tumor reports from UCSF, 33 (36.7%) missed size or location of one or more tumor. Conversely, AbdomenAtlas 3.0 and RadGPT provide precise measurements and locations for all detected tumors. In summary, our structured reports have details usually absent in human-made reports: locations and measurements for all

detected tumors, and tumor/organ volumes and HU values.

Human-AI fusion reports enhance human-made reports or clinical notes with the precision and detail of RadGPT structured reports. AbdomenAtlas 3.0 has clinical notes for 240 CT scans (209 are public, for TotalSegmentator CT scans [4]). However, these notes lack quantitative details in reports from RadGPT: of the 63 TotalSegmentator notes that mention tumors, none includes measurements, and many mention diverse tumors but do not describe them individually. However, these clinical notes describe diagnoses beyond the scope of RadGPT (which focuses on tumor-related diagnoses), such as calcified arterial plaques. We used an LLM to merge clinical notes with detailed structured and narrative reports in AbdomenAtlas 3.0, resulting in 240 fusion reports that integrate general clinical findings with precise tumor data (see Figure 11).

5. Conclusion & Future Work

AbdomenAtlas 3.0 is the first public dataset providing high-quality abdominal CT scans with reports and per-voxel tumor annotations, encompassing 9,262 CT scans from 88 institutions. It uniquely includes pancreas sub-segments, peripancreatic blood vessels, and pancreatic cancer stages, features absent in existing public datasets. RadGPT is an Anatomy-Aware Vision-Language AI Agent that transforms per-voxel annotations into structured reports using deterministic algorithms. These reports align with the accuracy of pathology findings annotated by a team of radiologists in AbdomenAtlas 3.0. Additionally, RadGPT enables fully automated report generation, surpassing existing approaches in detecting tumors from abdominal CT scans. Together, AbdomenAtlas 3.0 and RadGPT bridge the gap between tumor segmentation and report generation, offering valuable resources and tools to advance AI in automated abdominal CT interpretation.

We are committed to expanding AbdomenAtlas 3.0 to include reports for more types of tumors. Additionally, we plan to host benchmarks using AbdomenAtlas 3.0 with two train/test splits. **IID Split:** Randomly sets aside 10% of the dataset for testing, where training and testing data come from the same institutions, following standard AI evaluation practices. **OOD Split:** Uses data from 23 unseen institutions (4,500 CT scans) for testing, providing a large test set to evaluate AI generalization to new environments. This benchmark will assess report generation models using standard text similarity metrics but will prioritize tumor detection sensitivity and specificity, enhanced by our proposed LLM-based diagnostic evaluation.

Acknowledgments. This work was supported by the Lustgarten Foundation for Pancreatic Cancer Research and the Patrick J. McGovern Foundation Award. We thank the funding of Italian Institute of Technology and the HPC infrastructure at Italian Institute of Technology.

References

- [1] Mahmoud M Al-Hawary, Isaac R Francis, Suresh T Chari, Elliot K Fishman, David M Hough, David S Lu, Michael Macari, Alec J Megibow, Frank H Miller, Koenraad J Mortelet, et al. Pancreatic ductal adenocarcinoma radiology reporting template: consensus statement of the society of abdominal radiology and the american pancreatic association. *Radiology*, 270(1):248–260, 2014. 3, 5, 8
- [2] Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Bennett A Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger, Ronald M Summers, Bram van Ginneken, et al. The medical segmentation decathlon. *arXiv preprint arXiv:2106.05735*, 2021. 4
- [3] Michela Antonelli, Annika Reinke, Spyridon Bakas, Keyvan Farahani, Annette Kopp-Schneider, Bennett A Landman, Geert Litjens, Bjoern Menze, Olaf Ronneberger, Ronald M Summers, et al. The medical segmentation decathlon. *Nature communications*, 13(1):1–13, 2022. 3
- [4] Fan Bai, Yuxin Du, Tiejun Huang, Max Q-H Meng, and Bo Zhao. M3d: Advancing 3d medical image analysis with multi-modal large language models. *arXiv preprint arXiv:2404.00578*, 2024. 2, 3, 6, 8
- [5] Shruthi Bannur, Kenza Bouzid, Daniel C Castro, Anton Schwaighofer, Sam Bond-Taylor, Maximilian Ilse, Fernando Pérez-García, Valentina Salvatelli, Harshita Sharma, Felix Meissen, et al. Maira-2: Grounded radiology report generation. *arXiv preprint arXiv:2406.04449*, 2024. 2
- [6] Pedro RAS Bassi, Sergio SJ Dertkigil, and Andrea Cavalli. Improving deep neural network generalization and robustness to background bias via layer-wise relevance propagation optimization. *Nature Communications*, 15(1):291, 2024. 6
- [7] Pedro RAS Bassi, Wenxuan Li, Yucheng Tang, Fabian Isensee, Zifu Wang, Jieneng Chen, Yu-Cheng Chou, Yannick Kirchhoff, Maximilian Rokuss, Ziyang Huang, Jin Ye, Junjun He, Tassilo Wald, Constantin Ulrich, Michael Baumgartner, Saikat Roy, Klaus H. Maier-Hein, Paul Jaeger, Yiwen Ye, Yutong Xie, Jianpeng Zhang, Ziyang Chen, Yong Xia, Zhaohu Xing, Lei Zhu, Yousef Sadegheih, Afshin Bozorgpour, Pratibha Kumari, Reza Azad, Dorit Merhof, Pengcheng Shi, Ting Ma, Yuxin Du, Fan Bai, Tiejun Huang, Bo Zhao, Haonan Wang, Xiaomeng Li, Hanxue Gu, Haoyu Dong, Jichen Yang, Maciej A. Mazurowski, Saumya Gupta, Linshan Wu, Jiabin Zhuang, Hao Chen, Holger Roth, Daguang Xu, Matthew B. Blaschko, Sergio Decherchi, Andrea Cavalli, Alan L. Yuille, and Zongwei Zhou. Touchstone benchmark: Are we on the right way for evaluating ai algorithms for medical segmentation? *Conference on Neural Information Processing Systems*, 2024. 5, 6
- [8] Patrick Bilic, Patrick Ferdinand Christ, Eugene Vorontsov, Grzegorz Chlebus, Hao Chen, Qi Dou, Chi-Wing Fu, Xiao Han, Pheng-Ann Heng, Jürgen Hesser, et al. The liver tumor segmentation benchmark (lits). *arXiv preprint arXiv:1901.04056*, 2019. 3, 4
- [9] Louis Blankemeier, Joseph Paul Cohen, Ashwin Kumar, Dave Van Veen, Syed Jamal Safdar Gardezi, Magdalini Paschali, Zhihong Chen, Jean-Benoit Delbrouck, Eduardo Reis, Cesar Truys, et al. Merlin: A vision language foundation model for 3d computed tomography. *arXiv preprint arXiv:2406.06512*, 2024. 2, 3
- [10] Kai Cao, Yingda Xia, Jiawen Yao, Xu Han, Lukas Lambert, Tingting Zhang, Wei Tang, Gang Jin, Hui Jiang, Xu Fang, et al. Large-scale pancreatic cancer detection via non-contrast ct and deep learning. *Nature medicine*, 29(12):3033–3043, 2023. 3
- [11] Z. Cao and et al. Volumetric versus linear measurements in lung cancer prognosis. *Annals of Surgical Oncology*, 26(12):3757–3763, 2019. 8
- [12] Juan Manuel Zambrano Chaves, Shih-Cheng Huang, Yanbo Xu, Hanwen Xu, Naoto Usuyama, Sheng Zhang, Fei Wang, Yujia Xie, Mahmoud Khademi, Ziyi Yang, et al. Towards a clinically accessible radiology foundation model: open-access and lightweight, with automated evaluation. *arXiv preprint arXiv:2403.08002*, 2024. 2
- [13] Jieneng Chen, Yingda Xia, Jiawen Yao, Ke Yan, Jianpeng Zhang, Le Lu, Fakai Wang, Bo Zhou, Mingyan Qiu, Qihang Yu, et al. Cancerunit: Towards a single unified model for effective detection, segmentation, and diagnosis of eight major cancers using a large collection of ct scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21327–21338, 2023. 2
- [14] Qi Chen, Xiaoxi Chen, Haorui Song, Zhiwei Xiong, Alan Yuille, Chen Wei, and Zongwei Zhou. Towards generalizable tumor synthesis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 4, 5, 6
- [15] Qi Chen, Yuxiang Lai, Xiaoxi Chen, Qixin Hu, Alan Yuille, and Zongwei Zhou. Analyzing tumors by synthesis. *arXiv preprint arXiv:2409.06035*, 2024. 4, 6
- [16] Noel CF Codella, Ying Jin, Shrey Jain, Yu Gu, Ho Hin Lee, Asma Ben Abacha, Alberto Santamaria-Pang, Will Guyman, Naiteek Sangani, Sheng Zhang, et al. Medimageinsight: An open-source embedding model for general domain medical imaging. *arXiv preprint arXiv:2410.06542*, 2024. 2
- [17] Errol Colak, Hui-Ming Lin, Robyn Ball, Melissa Davis, Adam Flanders, Sabeena Jalal, Kirti Magudia, Brett Marinelli, Savvas Nicolaou, Luciano Prevedello, Jeff Rudie, George Shih, Maryam Vazirabad, and John Mongan. Rsn2023 abdominal trauma detection. <https://kaggle.com/competitions/rsna-2023-abdominal-trauma-detection>, 2023. Kaggle. 4
- [18] Claude Couinaud. *Le foie: études anatomiques et chirurgicales*. Masson, 1957. 4
- [19] Manjiri Dighe, Joseph R Grajo, and Leslie Lee. *Abdominal Imaging: Case Review Series*. Elsevier Health Sciences, 2021. 2
- [20] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. 5
- [21] Yasunari Fukuda, Daisaku Yamada, Hidetoshi Eguchi, Tomoki Hata, Yoshifumi Iwagami, Takehiro Noda, Tada-fumi Asaoka, Koichi Kawamoto, Kunihito Gotoh, Shogo Kobayashi, et al. Ct density in the pancreas is a promising imaging predictor for pancreatic ductal adenocarcinoma. *Annals of surgical oncology*, 24:2762–2769, 2017. 5

- [22] F Gaillard et al. Radiopaedia: building an online radiology resource. European Congress of Radiology-RANZCR ASM 2011, 2011. 3
- [23] Ibrahim Ethem Hamamci, Sezgin Er, and Bjoern Menze. Ct2rep: Automated radiology report generation for 3d medical imaging, 2024. 6, 12
- [24] Keiji Hanada, Akihito Okazaki, Naomichi Hirano, Yoshihiro Izumi, Yuji Teraoka, Juri Ikemoto, Kozue Kanemitsu, Fumiaki Hino, Toshikatsu Fukuda, and Shuji Yonehara. Diagnostic strategies for early pancreatic cancer. *Journal of gastroenterology*, 50:147–154, 2015. 2
- [25] Nicholas Heller, Niranjana Sathianathan, Arveen Kalapara, Edward Walczak, Keenan Moore, Heather Kaluzniak, Joel Rosenberg, Paul Blake, Zachary Rengel, Makinna Oestreich, 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. 3
- [26] Nicholas Heller, Sean McSweeney, Matthew Thomas Peterson, Sarah Peterson, Jack Rickman, Bethany Stai, Resha Tejpaul, Makinna Oestreich, Paul Blake, Joel Rosenberg, et al. An international challenge to use artificial intelligence to define the state-of-the-art in kidney and kidney tumor segmentation in ct imaging., 2020. 4
- [27] Stephanie L Hyland, Shruthi Bannur, Kenza Bouzid, Daniel C Castro, Mercy Ranjit, Anton Schwaighofer, Fernando Pérez-García, Valentina Salvatelli, Shaury Srivastav, Anja Thieme, et al. Maira-1: A specialised large multimodal model for radiology report generation. *arXiv preprint arXiv:2311.13668*, 2023. 2
- [28] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silvana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghighi, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 590–597, 2019. 7
- [29] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18(2):203–211, 2021. 4, 5
- [30] Yuanfeng Ji, Hongliang Bai, Le Gu, Fucang Liu, Yingda Xia, Haofeng Lu, Zhennan Chen, Yunlong Gao, Hongyang Pan, Weilin Wang, et al. Amos: A large-scale abdominal multi-organ benchmark for versatile medical image segmentation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 111–120. Springer, 2022. 4
- [31] Erin S. Kilpatrick and colleagues. Thresholds for low and high kidney volume in adult men and women. *National Center for Biotechnology Information*, 2018. 16
- [32] Yoshihisa Kodama, Chaan S Ng, Tsung T Wu, Gregory D Ayers, Steven A Curley, Eddie K Abdalla, Jean Nicolas Vauthy, and Chusilp Charmsangavej. Comparison of ct methods for determining the fat content of the liver. *American Journal of Roentgenology*, 188(5):1307–1312, 2007. 5
- [33] Shigeru Kondoh and colleagues. Ct imaging characteristics of the pancreas: Normal and abnormal findings. *PubMed*, 2018. 16
- [34] Bennett Landman, Zhoubing Xu, J Igelsias, Martin Styner, T Langerak, and Arno Klein. Miccai multi-atlas labeling beyond the cranial vault—workshop and challenge. In *Proc. MICCAI Multi-Atlas Labeling Beyond Cranial Vault—Workshop Challenge*, page 12, 2015. 4
- [35] Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [36] Wenxuan Li, Chongyu Qu, Xiaoxi Chen, Pedro RAS Bassi, Yijia Shi, Yuxiang Lai, Qian Yu, Huimin Xue, Yixiong Chen, Xiaorui Lin, et al. Abdomenatlas: A large-scale, detailed-annotated, & multi-center dataset for efficient transfer learning and open algorithmic benchmarking. *Medical Image Analysis*, page 103285, 2024. 4
- [37] Wenxuan Li, Alan Yuille, and Zongwei Zhou. How well do supervised models transfer to 3d image segmentation? In *International Conference on Learning Representations*, 2024. 5, 6
- [38] Wenxuan Li, Pedro R. A. S. Bassi, Tianyu Lin, Yu-Cheng Chou, Xinze Zhou, Yucheng Tang, Fabian Isensee, Kang Wang, Qi Chen, Xiaowei Xu, Xiaoxi Chen, Lizhou Wu, Qilong Wu, Yannick Kirchhoff, Maximilian Rokuss, Saikat Roy, Yuxuan Zhao, Dexin Yu, Kai Ding, Constantin Ulrich, Klaus Maier-Hein, Yang Yang, Alan L. Yuille, and Zongwei Zhou. Scalemai: Accelerating the development of trusted datasets and ai models. *arXiv preprint arXiv:2501.03410*, 2025. 4
- [39] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*. Association for Computational Linguistics, 2004. 3, 8
- [40] Jie Liu, Yixiao Zhang, Jie-Neng Chen, Junfei Xiao, Yongyi Lu, Bennett A Landman, Yixuan Yuan, Alan Yuille, Yucheng Tang, and Zongwei Zhou. Clip-driven universal model for organ segmentation and tumor detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21152–21164, 2023. 2
- [41] Jie Liu, Yixiao Zhang, Kang Wang, Mehmet Can Yavuz, Xiaoxi Chen, Yixuan Yuan, Haoliang Li, Yang Yang, Alan Yuille, Yucheng Tang, et al. Universal and extensible language-vision models for organ segmentation and tumor detection from abdominal computed tomography. *Medical Image Analysis*, page 103226, 2024. 2
- [42] Xiangde Luo, Wenjun Liao, Jianghong Xiao, Tao Song, Xiaofan Zhang, Kang Li, Guotai Wang, and Shaoting Zhang. Word: Revisiting organs segmentation in the whole abdominal region. *arXiv preprint arXiv:2111.02403*, 2021. 4
- [43] Jun Ma, Yao Zhang, Song Gu, Cheng Zhu, Cheng Ge, Yichi Zhang, Xingle An, Congcong Wang, Qiyan Wang, Xin Liu, et al. Abdomenct-1k: Is abdominal organ segmentation a solved problem. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 4
- [44] Jun Ma, Yuxin Zhang, Yuankai Gu, Haozhe Zhang, Zonghao Han, Jiancheng Dong, Yong Xie, Yefeng Zheng, and Yaowei Chen. Fast and low-gpu-memory abdominal organ segmentation from 3d ct images with nnformer and swin-unetr. In *International Challenge on Abdominal Multi-Organ*

- Segmentation from CT Images, FLARE 2022*, pages 83–94. Springer, 2022. 4
- [45] Matthew Macauley, Sheena Gorman, Raymond M Greer, Ian Lipton, John McMahon, Fergus Cameron, Denny Cuthbertson, Peter Colman, Thomas W Jones, and Elizabeth A Davis. Reduced pancreatic volume in type 1 diabetes. *Diabetes Care*, 40(9):e77–e78, 2017. 8
- [46] A. McErlean and et al. Evaluation of tumor volume as a response metric in clinical trials and patient management. *Journal of Surgical Oncology*, 121(4):523–531, 2015. 8
- [47] A.B. Miller, B. Hoogstraten, M. Staquet, and A. Winkler. Reporting results of cancer treatment. *Cancer*, 47(1):207–214, 1981. 4, 7
- [48] Radiological Society of North America (RSNA). Radreport: Reporting templates. <https://www.radreport.org>, 2025. Accessed: 2025-01-06. 5
- [49] Perry J Pickhardt, Peter M Graffy, Ryan Zea, Scott J Lee, Jiamin Liu, Veit Sandfort, and Ronald M Summers. Automated ct biomarkers for opportunistic prediction of future cardiovascular events and mortality in an asymptomatic screening population: a retrospective cohort study. *The Lancet Digital Health*, 2(4):e192–e200, 2020. 8
- [50] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 2
- [51] Blaine Rister, Darvin Yi, Kaushik Shivakumar, Tomomi Nobashi, and Daniel L Rubin. Ct-org, a new dataset for multiple organ segmentation in computed tomography. *Scientific Data*, 7(1):1–9, 2020. 4
- [52] Lewis R Roberts, Claude B Sirlin, Feras Zaiem, Jehad Almasri, Larry J Prokop, Julie K Heimbach, M Hassan Murad, and Khaled Mohammed. Imaging for the diagnosis of hepatocellular carcinoma: a systematic review and meta-analysis. *Hepatology*, 67(1):401–421, 2018. 2
- [53] Holger R Roth, Le Lu, Amal Farag, Hoo-Chang Shin, Jiamin Liu, Evrim B Turkbey, and Ronald M Summers. Deeporgan: Multi-level deep convolutional networks for automated pancreas segmentation. In *International conference on medical image computing and computer-assisted intervention*, pages 556–564. Springer, 2015. 4
- [54] Khaled Saab, Tao Tu, Wei-Hung Weng, Ryutarō Tanno, David Stutz, Ellery Wulczyn, Fan Zhang, Tim Strother, Chunjong Park, Elahe Vedadi, et al. Capabilities of gemini models in medicine. *arXiv preprint arXiv:2404.18416*, 2024. 2
- [55] Laura Schöckel, Gregor Jost, Peter Seidensticker, Philipp Lengsfeld, Petra Palkowitsch, and Hubertus Pietsch. Developments in x-ray contrast media and the potential impact on computed tomography. *Investigative radiology*, 55(9):592–597, 2020. 2
- [56] Andrew Taylor, William Dodds, Sandra Erickson, and Edward Stewart. Ct of acquired abnormalities of the spleen. *AJR American Journal of Roentgenology*, 157(6):1213–1219, 1991. 16
- [57] Gianluca Tomasello, Michele Ghidini, Antonio Costanzo, Antonio Ghidini, Alessandro Russo, Sandro Barni, Rodolfo Passalacqua, and Fausto Petrelli. Outcome of head compared to body and tail pancreatic cancer: a systematic review and meta-analysis of 93 studies. *Journal of gastrointestinal oncology*, 10(2):259, 2019. 3
- [58] Alexandre Triay Bagur, Paul Aljabar, Gerard R Ridgway, Michael Brady, and Daniel P Bulte. Pancreas mri segmentation into head, body, and tail enables regional quantitative analysis of heterogeneous disease. *Journal of Magnetic Resonance Imaging*, 56(4):997–1008, 2022. 4, 7
- [59] Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaeckermann, Mohamed Amin, Pi-Chuan Chang, Andrew Carroll, Charles Lau, Ryutarō Tanno, Ira Ktena, et al. Towards generalist biomedical ai. *NEJM AI*, 1(3):AIoa2300138, 2024. 2
- [60] Vanya V Valindria, Nick Pawlowski, Martin Rajchl, Ioannis Lavdas, Eric O Aboagye, Andrea G Rockall, Daniel Rueckert, and Ben Glocker. Multi-modal learning from unpaired images: Application to multi-organ segmentation in ct and mri. In *2018 IEEE winter conference on applications of computer vision (WACV)*, pages 547–556. IEEE, 2018. 4
- [61] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M Summers. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2097–2106, 2017. 7
- [62] Mateusz Winder, Aleksander Jerzy Owczarek, Jerzy Chudek, Joanna Pilch-Kowalczyk, and Jan Baron. Are we overdoing it? changes in diagnostic imaging workload during the years 2010–2020 including the impact of the sars-cov-2 pandemic. In *Healthcare*, page 1557. MDPI, 2021. 2
- [63] H.C. Woodruff and et al. The role of tumor volume in assessing treatment response in cancer imaging. *Cancer Imaging*, 15(1):5–12, 2015. 8
- [64] Yingda Xia, Qihang Yu, Linda Chu, Satomi Kawamoto, Seyoun Park, Fengze Liu, Jieneng Chen, Zhuotun Zhu, Bowen Li, Zongwei Zhou, et al. The felix project: Deep networks to detect pancreatic neoplasms. *medRxiv*, 2022. 3
- [65] Shawn Xu, Lin Yang, Christopher Kelly, Marcin Sieniek, Timo Kohlberger, Martin Ma, Wei-Hung Weng, Atilla Kiraly, Sahar Kazemzadeh, Zakkai Melamed, et al. Elixr: Towards a general purpose x-ray artificial intelligence system through alignment of large language models and radiology vision encoders. *arXiv preprint arXiv:2308.01317*, 2023. 2
- [66] Yanwu Xu, Li Sun, Wei Peng, Shuyue Jia, Katelyn Morrison, Adam Perer, Afroz Zandifar, Shyam Visweswaran, Motahhare Eslami, and Kayhan Batmanghelich. Medsyn: Text-guided anatomy-aware synthesis of high-fidelity 3d ct images. *IEEE Transactions on Medical Imaging*, 2024. 2
- [67] Kai Zhang, Rong Zhou, Eashan Adhikarla, Zhiling Yan, Yixin Liu, Jun Yu, Zhengliang Liu, Xun Chen, Brian D Davison, Hui Ren, et al. A generalist vision–language foundation model for diverse biomedical tasks. *Nature Medicine*, pages 1–13, 2024. 2

- [68] Sheng Zhang, Yanbo Xu, Naoto Usuyama, Hanwen Xu, Jaspreet Bagga, Robert Tinn, Sam Preston, Rajesh Rao, Mu Wei, Naveen Valluri, et al. Biomedclip: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs. *arXiv preprint arXiv:2303.00915*, 2023. [2](#)
- [69] Xukun Zhang, Sharib Ali, Tao Liu, Xiao Zhao, Zhiming Cui, Minghao Han, Shuwei Ma, Jingyi Zhu, Yanlan Kang, Le Wang, et al. Robust and smooth couinaud segmentation via anatomical structure-guided point-voxel network. *Computers in Biology and Medicine*, 182:109202, 2024. [4](#), [7](#)
- [70] Zongwei Zhou, Michael B Gotway, and Jianming Liang. Interpreting medical images. In *Intelligent Systems in Medicine and Health*, pages 343–371. Springer, 2022. [2](#), [3](#)

Appendix

Table of Contents

A AbdomenAtlas 3.0 Dataset	2
A.1 Visualizations	3
A.2 Word Cloud	10
B Technical Details of RadGPT	11
B.1 Training CT2Rep	12
B.2 Segmentation Post-processing	12
B.3 LLM Prompts	14
B.4 Organ size standards	16
C Detailed Tumor Statistics	17

A. AbdomenAtlas 3.0 Dataset




Figure 5. **Tumor size distribution in AbdomenAtlas 3.0. A large number of CT scans in AbdomenAtlas 3.0 present small tumors ($\leq 2\text{ cm}$): 943.** The figure’s top row shows histograms of all annotated tumors, while the bottom row focuses on the largest tumor in each organ. Notably, even considering only the largest tumor per organ, AbdomenAtlas 3.0 still includes a substantial number of small tumors ($\leq 2\text{ cm}$): 504 for kidney, 358 for liver, and 81 for pancreas. These small tumor reports are crucial for training vision-language AI models to detect early-stage cancers, where identifying subtle abnormalities is critical for early detection and treatment.

A.1. Visualizations

A.1.1. Cancer Staging and Blood Vessels




Figure 6. **Our pancreatic tumor (PDAC) staging report for a stage T4 tumor.** To determine the PDAC T stage, radiologists measure the tumor's size and evaluate its interactions with critical nearby blood vessels. RadGPT automatically replicates this process by utilizing per-voxel annotations of the PDAC and surrounding major blood vessels. The figure highlights these segmentations, and the report shows the angles of contact between the tumor and the blood vessels. In this case, the PDAC fully encases the superior mesenteric artery (SMA), which is a vital vessel supplying blood to the intestines. Surgical removal of a tumor encasing the SMA is not feasible because the artery cannot be reconstructed or bypassed without severe risk to the patient's survival. This involvement classifies the tumor as surgically unresectable and a stage T4 tumor.

A.1.2. Pancreas Sub-Segments




Figure 7. CT scan with 2 pancreatic tumors (yellow), and illustration of pancreas sub-segmentation into head (white, left), body (yellow, middle) and tail (white, right). RadGPT used the sub-segments to localize both PDAC tumors in the pancreas head. AbdomenAtlas 3.0 is the first to present pancreas sub-segments annotated per voxel. This information is crucial for writing radiology reports, as localizing pancreatic tumors in the pancreas head, body or tail is key for determining if the tumor can be surgically removed, and for tracking tumors in time.




Figure 8. CT scan with a pancreatic cyst (yellow), and illustration of pancreas sub-segmentation into head (white, left), body (yellow, middle) and tail (white, right). RadGPT used the sub-segments to localize the cyst in the pancreas head. AbdomenAtlas 3.0 is the first to present pancreas sub-segments annotated per voxel. This information is crucial for writing radiology reports, as localizing pancreatic tumors in the pancreas head, body or tail is key for determining if the tumor can be surgically removed, and for tracking tumors in time.

A.1.3. Liver Sub-segments

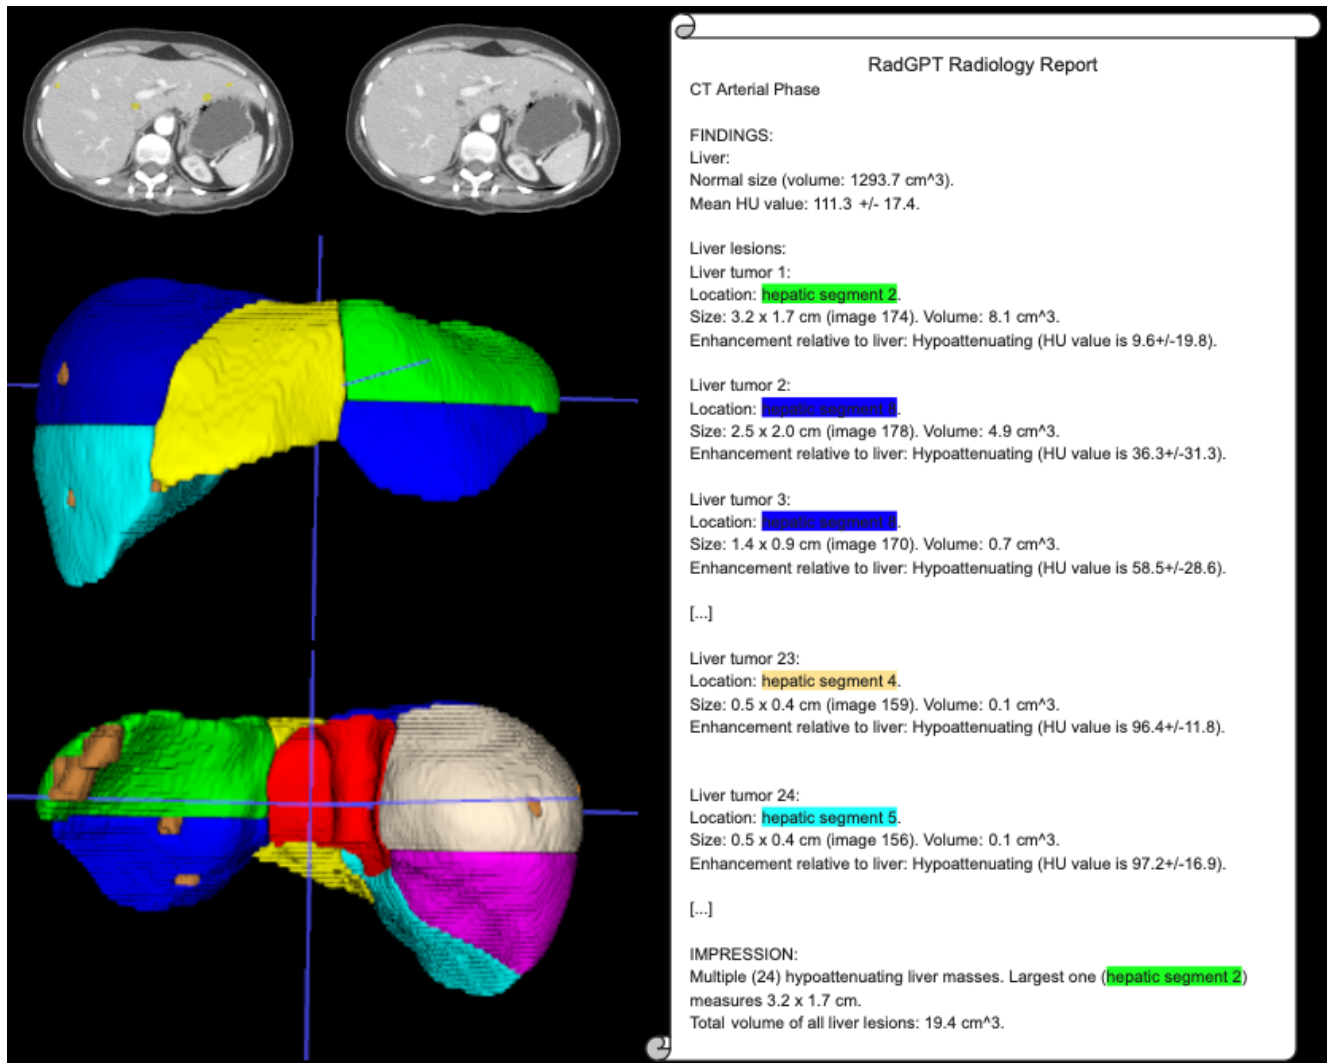


Figure 9. CT scan with 24 liver tumors (brown), showing how we segment the liver into eight sub-segments for tumor localization. Notably, unlike our report, most human-made reports would not describe 24 tumors in detail, due to the time required for this task. Liver sub-segments are functionally independent, and can be surgically removed without influencing nearby segments. Thus, localizing tumors into these segments is important for tracking tumors and for surgical planning.

A.1.4. Kidney Tumor Report




Figure 10. CT scan showing a large kidney tumor (yellow) and our report.

A.1.5. Human-AI Fusion Reports

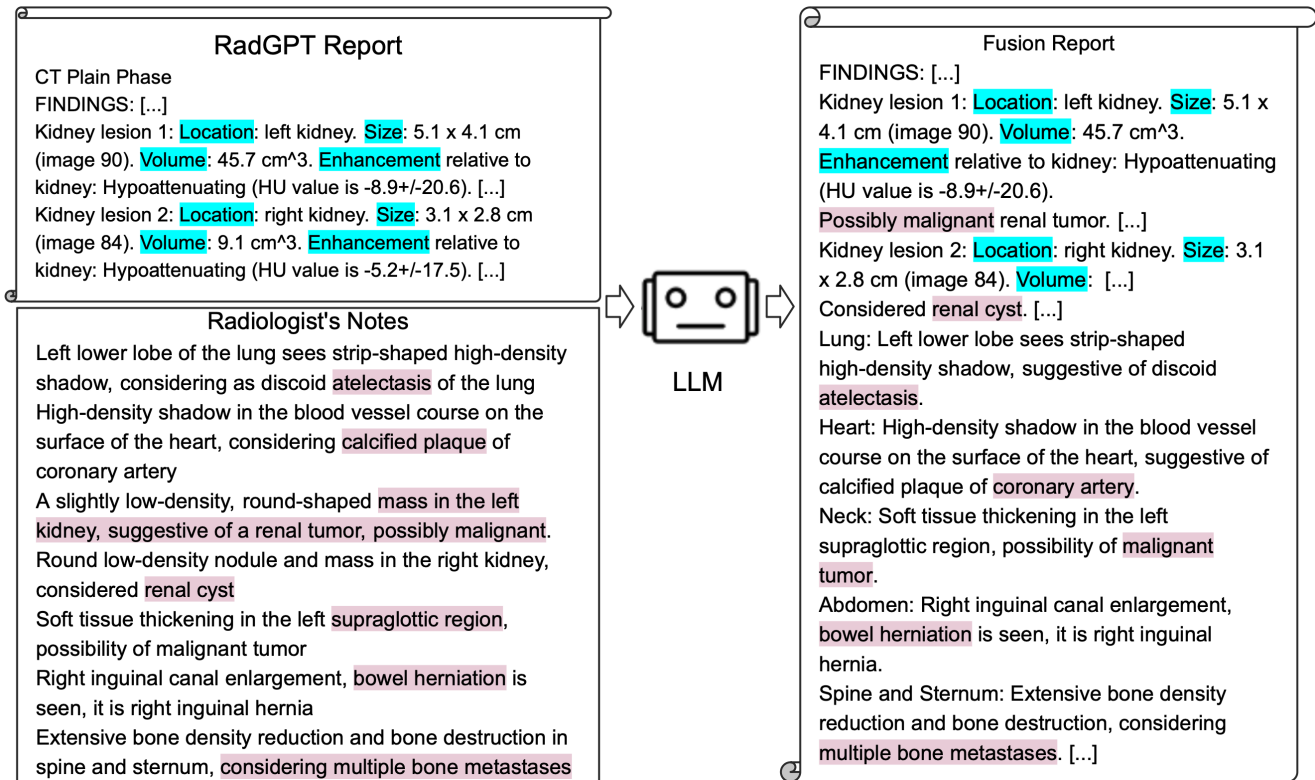


Figure 11. **In our human-AI fusion reports, LLMs combine detailed quantitative data from RadGPT’s reports with the generality of human-made reports or clinical notes.** In the image, the LLM extracted kidney tumor sizes and volumes from the RadGPT report, while incorporating tumor type and non-cancer-related findings from the radiologist’s notes. Reports created by RadGPT include detailed information often absent in human-made reports, such as tumor volumes and individual measurements for a large number of tumors when present. In contrast, human-made reports provide general findings that can expand RadGPT reports to many insights beyond cancer.

A.1.6. Narrative Reports

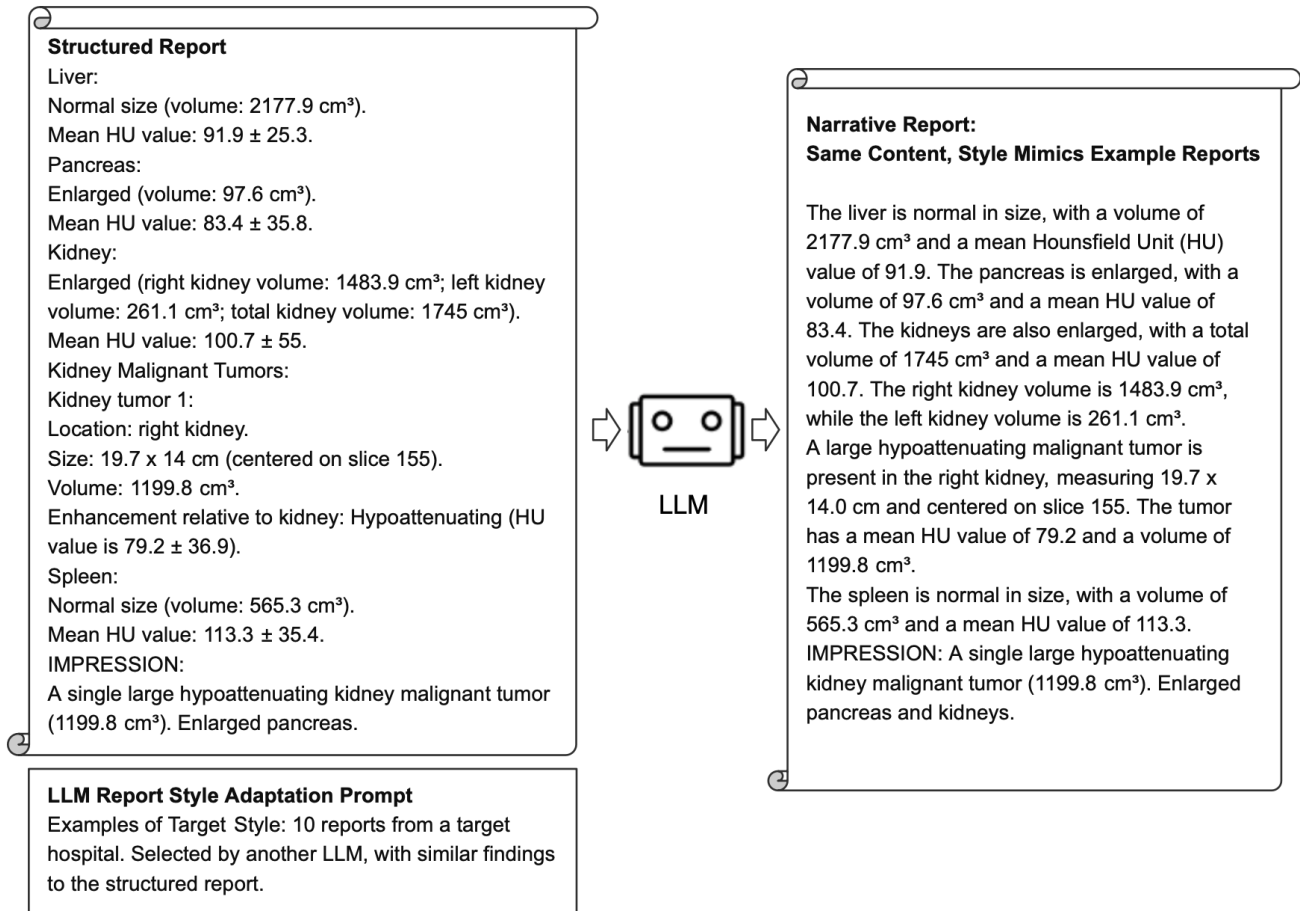


Figure 12. **Example of Narrative Report: we use LLM to convert structured reports into narrative reports that follow the writing style of a target hospital's.** The LLM receives 10 example reports from the hospitals as examples of style, and is instructed not to change the medical content of the structured report during style adaptation. Since reports targeting diverse abnormalities vary strongly in style, working ans structure, we use another LLM to pre-classify the hospital's human-made reports into diagnostic categories (e.g., liver tumor). Thus, during style adaptation, we use as examples only reports that focus on the same diagnosis as the structured report. E.g., if the structured report mentions liver tumors, the examples also will concentrate on liver tumors.

B. Technical Details of RadGPT

Algorithm 1 Pancreas Sub-segmentation Using SMA

```
1: Erase parts of the SMA annotation below the pancreas annotation.
2: Perform PCA on a random subset of the pancreas voxels and rotate the pancreas around its center of mass, aligning its principal component with the x-axis. Rotate the SMA annotation together with the pancreas.
3: Project the SMA onto the x-axis; consider the x-plane at projection's midpoint as the boundary between pancreatic head and body.
4: For the remaining pancreas (excluding head), split body and tail at the x-axis midpoint, using another x-plane.
5: for each x-plane (slice) from tail to head do
6:   Identify connected components in the current slice.
7:   if first pancreas slice then
8:     Classify all components as body.
9:   else
10:    Classify components overlapping with the body in the previous slice as body; reclassify others as head. This is important for cases where the pancreatic head bottom crosses the SMA.
11:   end if
12: end for
13: Undo rotations and translations; save head, body, and tail segmentations.
```

Algorithm 2 WHO-based Tumors Size Measurement

```
1: Interpolate the tumor segmentation mask to a standard 1x1x1 mm spacing.
2: for each CT slice  $s$  containing tumor  $A$  do
3:   Extract the tumor borders by subtracting the tumor segmentation slice  $s$  by itself after binary erosion.
4:   Calculate the diameter  $D_s$  as the longest line between any two points in the tumor borders in  $s$ .
5: end for
6: Select the slice  $s_{max}$  with the largest diameter  $D_{max}$ .
7: In the selected slice  $s_{max}$ :
8:   Draw two lines  $L_1$  and  $L_2$  parallel to the diameter  $D_{max}$ .
9:   Set these two parallel lines  $L_1$  and  $L_2$  as far as possible from each other while touching the tumor borders.
10:  Calculate the distance  $d$  between lines  $L_1$  and  $L_2$ .
11: Report the tumor size as  $D_{max} \times d$ , converting from mm to cm.
```

Algorithm 3 Automatic Tumor Staging

```
1: # Make tumor borders overlap with vessels and organs
2: Apply binary dilation (3x3x3) on tumor mask.
3: for each vessel in {SMA, CHA, CA, SA} do
4:   if no overlap with tumor then
5:     Set contact = no and continue
6:   end if
7:   # Isolate main vessel branch
8:   for each slice along z-axis from top to bottom do
9:     Retain the largest connected component touching the previous slice's main component, or the largest if within the first 5% of slices.
10:   end for
11:   Apply binary erosion and dilation (5x5x5), overlap with original vessel segmentation, and retain the largest 3D component.
12:   # Check main branch contact with tumor
13:   if no overlap with tumor then
14:     Set contact = no and continue
15:   end if
16:   # Align vessel over x-axis and analyze contact with tumor
17:   Skeletonize main branch and align rotate volume, aligning principal component (PCA) with x-axis.
18:   for each x-coordinate along the x-axis do
19:     Check intersection with tumor; if none, continue
20:     Align 5mm vessel segment around x-axis using skeleton PCA and crop to 2.5mm
21:     # Calculate percentage of border contact with tumor to estimate contact angle (vessels are not perfectly round)
22:     Extract vessel borders for each slice and calculate border-tumor overlap percentage
23:     Compute contact angle as  $\text{overlap percentage} \times 360$ ; update max_contact for vessel if new maximum angle is found.
24:   end for
25: end for
26: # Define T stage based on vessel contact and tumor size thresholds
27: if max_contact for {SMA, CA, CHA}  $\geq 180$  then
28:   Stage = T4
29: else
30:   Determine stage by tumor size: T1a  $\leq 0.5$ cm, T1b 0.5–1cm, T1c 1–2cm, T2 2–4cm, T3  $> 4$ cm
31: end if
```

B.1. Training CT2Rep

We trained CT2Rep using only CT scans and structured reports, ignoring the per-voxel annotations in AbdomenAtlas 3.0. Our training strategy for CT2Rep closely followed the code and hyper-parameters published by the model authors [23]. Possibly, careful search of hyper-parameters and training algorithms for the abdominal region could improve the model’s performance. We introduced minimal changes, needed to adapt the model to the abdominal region: we adopted sub-word tokenization to handle decimals frequently present in our reports; we standardized the CT spacing to 1.5 x 1.5 x 1.5 mm, a choice that reduces computational costs while facilitating tumor measurements by maintaining isotropy; to accommodate longer reports, we increased the model’s maximum sequence length to 600; and, for hold-out validation (we used 30% of AbdomenAtlas 3.0 as the validation set), we used validation loss rather than sequential decoding and BLEU scoring, which significantly reduced validation time. These adjustments, while minimal, were designed to tailor the model for the unique challenges of abdominal CT report generation.

B.2. Segmentation Post-processing

Segmentation models can produce noise: voxels incorrectly labeled as tumors or organs. This may cause false positive cancer detections when RadGPT generates reports from nnU-Net or DiffTumor outputs. To address this, we propose a noise reduction algorithm (Alg. 4). Segmentation noise usually appears as small structures. Thus, we reduce it with binary erosion. Afterwards, to restore the original shape of true tumors and organs, we applied binary dilation followed by a voxel-wise AND with the original tumor segmentation. To further avoid false positives, we perform organ-wise thresholding: we only consider an organ has tumors if the total volume of its tumor voxels is above a small threshold, defined to maximize per-class F1-Score on a validation dataset. For our results section, RadGPT thresholds are: 1 mm³ in the pancreas, 150 mm³ in the kidneys, 100 mm³ in the liver, and 50 mm³ for metastases. Figure 14 shows specificity and sensitivity for multiple thresholds. Algorithm 4 and thresholding are not necessary when we generate AbdomenAtlas 3.0 reports from radiologist revised segmentations or ground-truth segmentation masks. However, it is recommended when using RadGPT without human revision (Figure 2). Figure 14 displays performance variation for diverse thresholds.

Algorithm 4 Segmentation Noise Reduction

- 1: Copy the segmentation output.
 - 2: Apply binary erosion to the segmentation to erase small structures, considered noise. We use a 3x3x3 structuring element, erasing any structure smaller than a 3x3x3 cube.
 - 3: Perform binary dilation on the eroded segmentation. We use a 4x4x4 structuring element.
 - 4: Apply a voxel-wise AND operation between the original mask (before erosion) and the dilated mask, recovering the shape of structures not removed by the binary erosion.
-

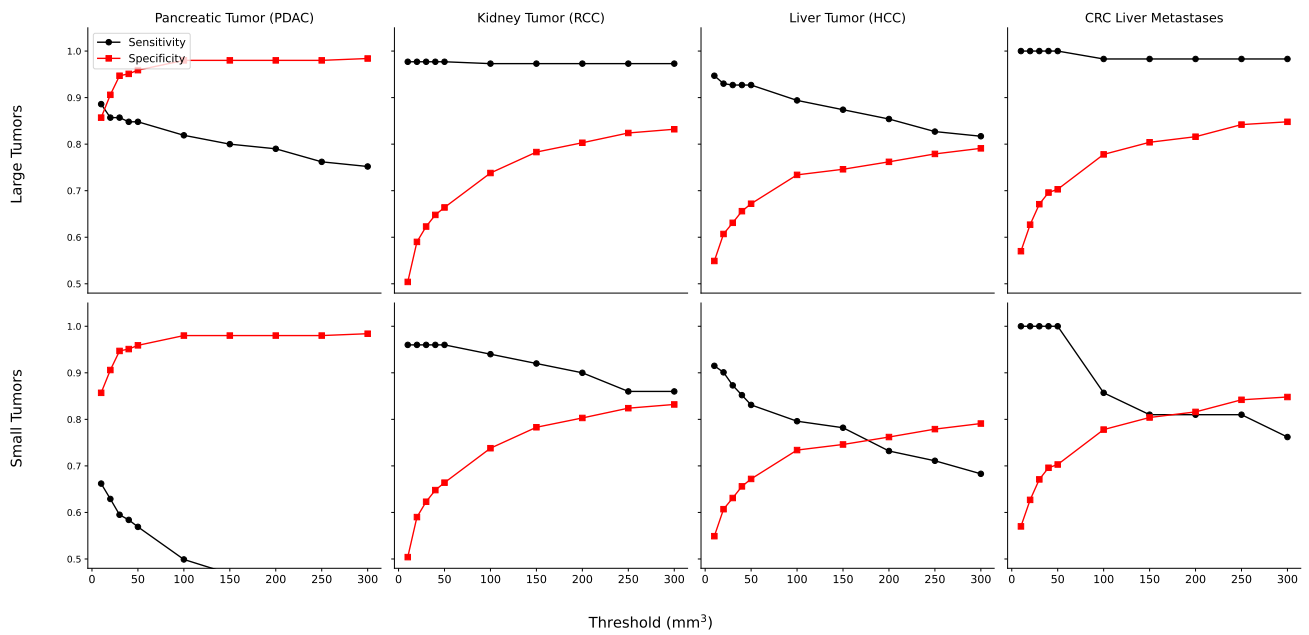


Figure 14. **Tumor detection sensitivity and specificity for RadGPT with diverse thresholds.** Evaluation performed on a private dataset from a hospital never observed during training, detailed in Table 2.

B.3. LLM Prompts

B.3.1. Style Adaptation

Or prompt for style adaptation is the following:

You are provided with a **structured radiology report** and n other radiology reports that have different writing styles compared to the structured report.

Task:

Please **paraphrase** the structured report to match the writing style of the other reports.

Important Guidelines:

1. **Do Not Alter Medical Information:** Do not change, add, or remove any medical details such as tumor measurements, types, or locations. You may remove HU values.
2. **Maintain Original Meaning:** Ensure that the rephrased report conveys the same information as the original structured report.
3. **Match Writing Style:** Adapt the language, tone, and structure to align with the provided example reports.
4. **Begin your report text with #start and finish it with #end.**
5. **Provide justification:** Go through all medical findings in your rephrased report (e.g., tumor size, no evidence of metastasis) and show where the information comes from in the structured report. Justification should come after #end.
6. **Pay attention to the Example Reports:** Your writing style must be consistent with the examples.
7. **Organization must match:** If the examples have an *Impressions* and *Results* section, you must add them. If the example reports talk about all abdominal organs in a single paragraph, you must do so too. You may skip sections you cannot fill due to lack of information, like patient history.
8. **Do not add new findings:** If the structured report does not mention the presence or absence of a medical condition (e.g., metastases), you must NOT include it in your rephrased report.
9. **Keep coherence:** Avoid going back and forth between medical findings or organs. For example, do not talk about the size of a pancreatic tumor, then mention the liver, and then go back to pancreatic findings. Keep the information about each organ together.
10. **Always include an impressions section with the most important findings.**

Example of Rephrasing:

Structured Report:

PDAC 1: Pancreatic body/tail. Hypoattenuating pancreas PDAC measuring 6.0 x 3.4 cm (centered on slice 356). Its mean HU value is 39.17 +/- 29.65, and its volume is 27.519 cm³.

Paraphrased Report:

#start

The patient has a pancreatic adenocarcinoma located in the body and tail of the pancreas, measuring 6.0 x 3.4 centimeters (image slice 356). The tumor is hypoattenuating and has a volume of 27.519 cm³.

#end

Justification:

- a. **Tumor Type:** Maintained as "pancreatic adenocarcinoma", originally "PDAC".
- b. **Location:** Preserved as "body and tail of the pancreas", originally "Pancreatic body/tail".
- c. **Measurements:** Kept as "6.0 x 3.4 centimeters", originally "measuring 6.0 x 3.4 cm".
- d. **Imaging Slice:** Retained as "image slice 356", originally "centered on slice 356".
- e. **Attenuation:** Maintained as "hypoattenuating", originally "Hypoattenuating pancreas PDAC".
- f. **Volume:** Kept as "27.519 cm³", originally "volume is 27.519 cm³".

Note: Removed mean HU value as per guidelines.

Example Reports (Target Style): {examples}

Structured Report to Paraphrase: {structured_report}

B.3.2. AI-Human Report Fusion

Our Report Fusion prompt is:

You are provided with a CT scan **structured radiology report** and notes written by a radiologist, about the same CT scan.

Your task is to identify any information in the notes that is not already included in the structured report and add it to the appropriate sections of the report. Please follow these guidelines:

1. **Do not remove** any existing information from the structured report. However, you may improve the report's details using **only** relevant information from the notes.
2. **Avoid adding any new findings** not already mentioned in either the notes or the structured report.
3. **Maintain the report's structure.** Carefully place new information in the correct sections inside "FINDINGS", considering which organ the information mentions. For instance, if the notes mention "cirrhosis," add it to the "**Liver**" section under "FINDINGS".
4. **Add new sections if necessary.** If the notes refer to an organ not covered in the structured report, create a new section for it. If the notes mention patient metadata (e.g., sex and age), you may add it to the beginning of the report.
5. **Update the IMPRESSION section if needed.** Besides the FINDINGS, include any critical information from the notes in the report's **IMPRESSION** section, summarizing or rephrasing it. Do not add new sections if the notes do not provide concrete information for them.
6. **Use consistent terminology.** If possible, make the terminology in the sentences you add to the report match the terminology in the original structured report.
7. **Begin your report text with #start and finish it with #end.**
8. **Provide justification.** Explain where in the report you added each piece of information from the notes. Also, explain why other information in the report was not removed or altered.
9. **Do not** write non-informative sentences such as "Patient metadata: Not available in the provided notes" or "Sex: Not specified."

The notes are as follows:

{clinical_info}

The current structured report is:

{structured_report}

B.3.3. Labeling/Report Evaluation

Our prompt is:

Instructions: Discover if the CT scan radiology report below indicates the presence of liver tumors, pancreas tumors, or kidney tumors. Output labels for each of these categories: **yes** to indicate tumor presence, **no** for tumor absence, and **U** for uncertain tumor presence.

Example: liver tumor presence=yes; kidney tumor presence=U; pancreas tumor presence=no.

Answer with only the labels, do not repeat this prompt.

Follow these rules for interpreting radiology reports:

1. 'Unremarkable' means that an organ has no tumor.
2. Multiple words can describe tumors. Check both the **findings** and **impressions** sections of the report (if present) to understand if an organ has tumors. Some words include: metastasis, tumor, tumor, mass, cyst, neoplasm, growth, cancer, index tumor in cancer patients, and tumors listed as oncologic findings.
3. Consider any tumor, hyperdensity, or hypodensity a tumor, unless the report explicitly states otherwise. Many conditions are not tumors and should not be interpreted as such unless a tumor is also reported along with the disease. Examples include:
 - **Liver conditions:** Hepatitis, Cirrhosis, Fatty Liver Disease (FLD), Liver Fibrosis, Hemochromatosis, Primary Biliary Cholangitis (PBC), Primary Sclerosing Cholangitis (PSC), Wilson's Disease, Liver Abscess, Alpha-1 Antitrypsin Deficiency (A1ATD), Steatosis, Granulomas, Cholestasis, Budd-Chiari Syndrome (BCS), Transplant, Gilbert's Syndrome, ulcers, wounds, infections, inflammations, and scars.
 - **Kidney conditions:** Stents, inflammation, postinflammatory calcification, transplant, Chronic Kidney Disease (CKD), Acute Kidney Injury (AKI), Glomerulonephritis, Nephrotic Syndrome, Polycystic Kidney Disease (PKD), Pyelonephritis, Hydronephrosis, Renal Artery Stenosis (RAS), Diabetic Nephropathy, Hypertensive Nephrosclerosis, Interstitial Nephritis, Renal Tubular Acidosis (RTA), Goodpasture Syndrome, and Alport Syndrome.
 - **Pancreas conditions:** Pancreatitis, Pancreatic Insufficiency, Cystic Fibrosis (CF), Diabetes Mellitus (DM), Exocrine Pancreatic Insufficiency (EPI), Pancreatectomy, and Pancreatic Pseudocyst.
4. Examples of specific tumor names include:
 - **Liver:** Hepatic Hemangioma (HH), Focal Nodular Hyperplasia (FNH), Bile Duct Adenoma, Simple Liver Cyst (SLC), Hepatocellular Carcinoma (HCC), Cholangiocarcinoma (CCA), Hepatic Adenoma (HA), Mucinous Cystic Neoplasm (MCN).
 - **Pancreas:** Serous Cystadenoma (SCA), Pancreatic Ductal Adenocarcinoma (PDAC), Mucinous Cystadenocarcinoma (MCC), Mucinous Cystadenoma (MCA), Intraductal Papillary Mucinous Neoplasm (IPMN), Solid Pseudopapillary Neoplasm (SPN), Pancreatic Neuroendocrine Tumor (PNET).
 - **Kidney:** Renal Oncocytoma (RO), Angiomyolipoma (AML), Simple Renal Cyst, Bosniak IIF Cystic Tumor, Renal Cell Carcinoma (RCC), Transitional Cell Carcinoma (TCC), Wilms Tumor, Cystic Nephroma (CN), Multilocular Cystic Renal Neoplasm of Low Malignant Potential (MCRNLMP), Hydronephrosis, Allograft.
5. Consider any benign (e.g., cyst) or malignant tumor as a tumor. Thus, any type of cyst is a tumor.
6. Organs never mentioned in the report have no tumors.
7. Do not assume a tumor is uncertain unless it is explicitly reported as uncertain. Many words can describe uncertainty, such as: ill-defined, too small to characterize, nonspecific, and uncertain. Reports may express uncertainty about tumor type (e.g., cyst or hemangioma) but still confirm it is a tumor—in this case, consider the tumor a tumor.
8. Organs with no tumor but other pathologies should be reported as **no**.

B.4. Organ size standards

Our standards for considering organs as large are based on widely accepted thresholds in radiological and anatomical studies. For the spleen, we consider volumes greater than 314.5 cm³ as large and over 430.8 cm³ as massive, based on thresholds provided by Taylor et al. [56]. For the kidneys, a volume exceeding 415.2 cm³ for men is considered large, with the threshold adjusted for individual kidneys (half of the total volume) [31]. Similarly, a liver volume exceeding 3000 cm³ is deemed large, which represents an upper limit for larger individuals, such as a 150 kg man, and highly depends on factors like weight and sex. For the pancreas, volumes above 83 cm³ are classified as large, as per imaging standards discussed by Kondoh et al. [33].

When size standards depend on variables like weight or sex, we apply thresholds suitable for larger individuals to ensure comprehensive assessments. This approach minimizes the risk of underestimating organ size variations in diverse populations.

C. Detailed Tumor Statistics

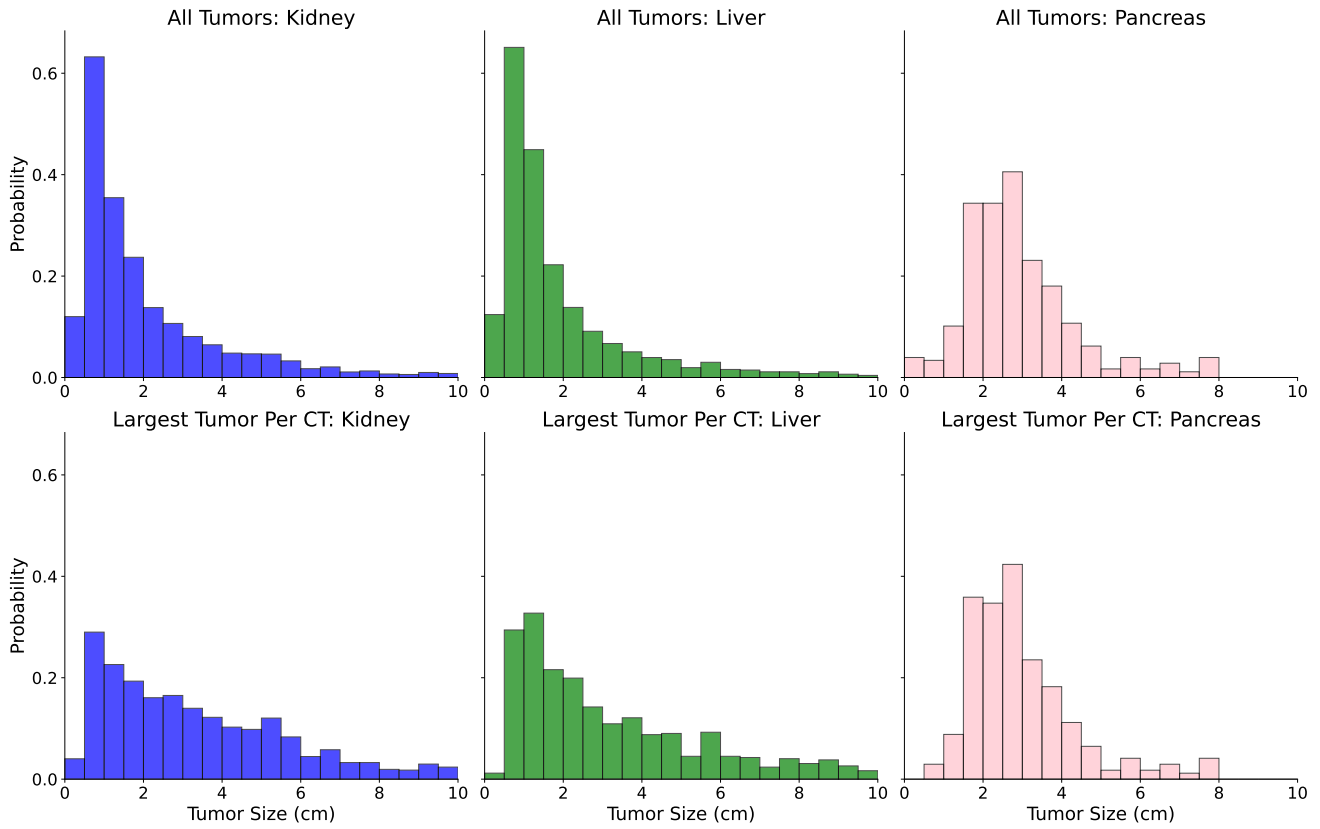


Figure 15. **Tumor size distribution in AbdomenAtlas 3.0. A large proportion of the CT scans, 35%, presents small tumors only (≤ 2 cm).** The figure’s top row shows histograms of all annotated tumors, while the bottom row focuses on the largest tumor in each organ. Notably, even considering only the largest tumor per organ, the proportion of small tumors (≤ 2 cm) is large in AbdomenAtlas 3.0: 35.59% for kidney, 38.25% for liver, and 23.68% for pancreas. These small tumor reports are vital for training vision-language AI models to detect early-stage cancers, where identifying subtle abnormalities is critical for early cancer detection and treatment.




Figure 16. **Tumor size probability distribution for liver tumors across all datasets in AbdomenAtlas 3.0.** Each subplot represents a dataset with at least three tumor occurrences. The x-axis shows tumor size (cm), and the y-axis represents the probability of tumors within each size range. The figure highlights the variability in tumor sizes annotated across datasets, and the significant presence of small tumors.




Figure 17. **Tumor size probability distribution for pancreas tumors across all datasets in AbdomenAtlas 3.0.** Each subplot represents a dataset with at least three tumor occurrences. The x-axis shows tumor size (cm), and the y-axis represents the probability of tumors within each size range. The figure highlights the variability in tumor sizes annotated across datasets, and the significant presence of small tumors.

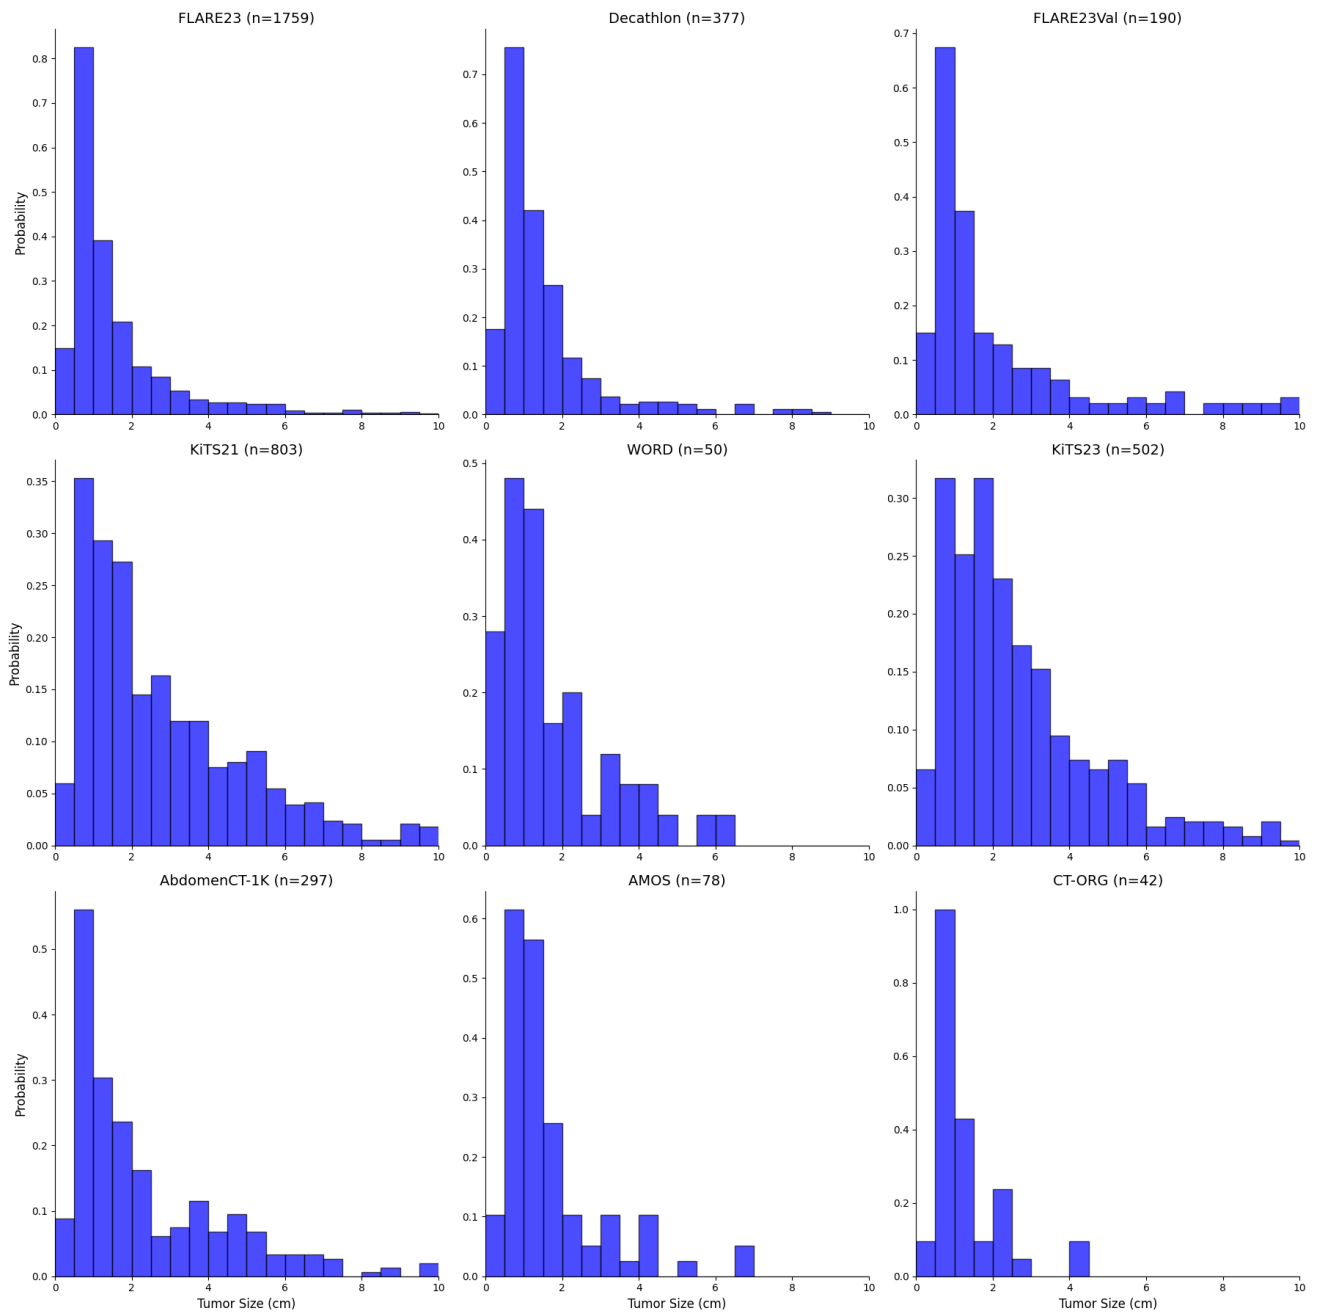


Figure 18. **Tumor size probability distribution for kidney tumors across all datasets in AbdomenAtlas 3.0.** Each subplot represents a dataset with at least three tumor occurrences. The x-axis shows tumor size (cm), and the y-axis represents the probability of tumors within each size range. The figure highlights the variability in tumor sizes annotated across datasets, and the significant presence of small tumors.




Figure 19. **Tumor size frequency histogram for liver tumors across all datasets in AbdomenAtlas 3.0.** Each subplot represents a dataset with at least three tumor occurrences. The x-axis shows tumor size (cm), and the y-axis represents the number of tumors within each size range. The figure highlights the variability in tumor sizes annotated across datasets, and the significant presence of small tumors.




Figure 20. **Tumor size frequency histogram for pancreas tumors across all datasets in AbdomenAtlas 3.0.** Each subplot represents a dataset with at least three tumor occurrences. The x-axis shows tumor size (cm), and the y-axis represents the number of tumors within each size range. The figure highlights the variability in tumor sizes annotated across datasets, and the significant presence of small tumors.




Figure 21. **Tumor size frequency histogram for kidney tumors across all datasets in AbdomenAtlas 3.0.** Each subplot represents a dataset with at least three tumor occurrences. The x-axis shows tumor size (cm), and the y-axis represents the number of tumors within each size range. The figure highlights the variability in tumor sizes annotated across datasets, and the significant presence of small tumors.