FluoroSAM: A Language-aligned Foundation Model for X-ray Image Segmentation

Benjamin D. Killeen^{*1}, Liam J. Wang¹, Han Zhang¹, Mehran Armand², Russell H. Taylor¹, Dave Dreizin³, Greg Osgood², and Mathias Unberath¹

> Johns Hopkins University, Baltimore, MD, USA *Corresponding author. ¹{killeen, wwang136, hzhang206, rht, unberath}@jhu.edu ²{marmand2, gosgood2}@jhmi.edu ³ ddreizin@umm.edu

Abstract. Automated X-ray image segmentation would accelerate research and development in diagnostic and interventional precision medicine. Prior efforts have contributed task-specific models capable of solving specific image analysis problems, but the utility of these models is restricted to their particular task domain, and expanding to broader use requires additional data, labels, and retraining efforts. Recently, foundation models (FMs) – machine learning models trained on large amounts of highly variable data thus enabling broad applicability – have emerged as promising tools for automated image analysis. Existing FMs for medical image analysis focus on scenarios and modalities where objects are clearly defined by visually apparent boundaries, such as surgical tool segmentation in endoscopy. X-ray imaging, by contrast, does not generally offer such clearly delineated boundaries or structure priors. During X-ray image formation, complex 3D structures are projected in transmission onto the imaging plane, resulting in overlapping features of varying opacity and shape. To pave the way toward an FM for comprehensive and automated analysis of arbitrary medical X-ray images, we develop FluoroSAM, a language-aligned variant of the Segment-Anything Model, trained from scratch on 1.6M synthetic X-ray images from a wide variety of human anatomies, X-ray projection geometries, energy spectra, and viewing angles. FluoroSAM is trained on data including masks for 128 organ types and 464 non-anatomical objects, such as tools and implants. In real X-ray images of cadaveric specimens, FluoroSAM is able to segment bony anatomical structures based on text-only prompting with 0.51 and 0.79 DICE with point-based refinement, outperforming competing SAM variants for all structures. FluoroSAM is also capable of zero-shot generalization to segmenting classes beyond the training set thanks to its language alignment, which we demonstrate for full lung segmentation on real chest X-rays.

Code, data, and model weights are available.¹

Keywords: fluoroscopy · radiology · image-guided surgery · medical AI

¹ https://github.com/arcadelab/fluorosam



Fig. 1: Overview of the FluoroSAM dataset. (a) Virtual C-arm views are sampled uniformly on the hemisphere. We project the segmentation maps for 128 organs and up to 15 out of 464 devices placed at random in the field of view. (b) Fluoroscopic images are transmissive by nature, with many overlapping masks but SAM is designed for nested masks that align with visually evident boundaries, (annotated here by hand). CLAHE applied for visualization.

1 Introduction

X-ray imaging is a workhorse imaging modality for diagnostic and interventional healthcare. There is enormous opportunity for quantitative, comprehensive, and automated segmentation of X-ray images to accelerate research and development in precision medicine [3, 4, 8, 14, 16, 25, 26, 29]. Prior efforts have contributed machine learning techniques for X-ray image analysis that perform well within a narrow scope, but fail to apply broadly to a large swath of possible uses. Extending these techniques to new applications requires labeled data in addition to personnel effort for retraining of models, that may or may not be available in sufficient quantities, thus inhibiting progress. Recently, foundation models (FMs) have emerged as a promising direction for overcoming these limitations [2, 5, 18, 20, 22]. FMs are characterized by scalable training strategies, often accomplished through self-supervision, that enable learning from very big and highly diverse data. Because of this, they exhibit strong generalizability and robust performance, resulting in a tendency to outperform more specialized models on downstream tasks [1]. However, due to the transmission imaging nature of X-ray, FMs – even the ones specialized for medical imaging [22] – often fail to generalize to this modality [9]. Further, the scarcity of X-ray data (beyond chest X-rays [3, 7, 10]) and accompanying labels inhibits development of an X-ray domain-specific model.

The Segment-Anything Model (SAM) is an attractive candidate for this task because of its flexibility and compatibility with human-in-the-loop workflows [18]. Given a user prompt, including points, bounding boxes, or text, SAM predicts a segmentation mask for semantically meaningful objects. Trained on a large scale dataset of over 1B masks, the original SAM is a powerful tool for automated and collaborative image segmentation, and it has been successfully fine-tuned on a variety of medical imaging modalities [22]. However, although models like MedSAM [22] are able to segment certain structures in chest Xray images, they suffer from an adherence to well-defined boundaries, which are typical of masks in other imaging modalities. At the same time, point-based prompting suffers from a high degree of ambiguity (see Fig. 1b) while box-based prompting requires expert user input, limiting the potential for automation. In the context of X-ray, text-based prompting is most desirable as it 1) allows for interaction-free segmentation given pre-defined prompts while 2) preserving the flexibility of zero-shot generalization and point prompt-based refinement; and 3) can unambiguously express the desired segmentation among many overlapping but highly disparate structures.

Recent advances in simulation [16, 26] and sim-to-real transfer [8, 16, 26] introduce the possibility of training FMs for X-ray image analysis in a fully supervised manner. This opens the door to a language-aligned FM in the X-ray domain, since perfect knowledge of simulated objects allows for automatic annotation of mask descriptions. Here, we use a physics-based simulation pipeline [26] to synthesize a large scale, full body dataset of 1.6M digitally reconstructed radiographs (DRRs), uniformly sampled over a wide variety of human anatomies, X-ray geometries, energy spectra, and viewing angles. Anatomical masks of 128 organs are projected from automatic segmentations of each CT image, using TotalSegmentator [28]. To encourage generalizability, we simulate alongside each patient model a subset of implants, surgical tools, and other devices commonly found in X-ray, associated with comprehensive text descriptions.

Based on this novel dataset, which to the best of our knowledge is the largest publicly available X-ray dataset, we present FluoroSAM, a language-aligned FM for automatic and interactive segmentation of objects and anatomical structures in medical X-ray images. During training, FluoroSAM uses organ names and device descriptions as the initial prompts, augmented using a large-language model to encourage flexbility. Subsequent point-based prompts allow for refinement of the initial segmentation. This approach allows for fully automated image segmentation by relying solely on text-based prompting, as well as for interactive workflows initiated by an initial prompt of "bone," "organ," or "device," for example and then followed by point prompts, allowing for flexibility in downstream applications. We evaluate FluoroSAM's zero-shot sim-to-real performance on real diagnostic and interventional X-ray images, including publicly available chest X-ray data and internally collected X-ray images of cadaveric torsos.

2 Methods

The FluoroSAM dataset is a large-scale, synthetic dataset of X-ray images of all human anatomies with a wide variety of medical devices and implants, including more than 63M masks with descriptions. Based on this dataset, we train FluoroSAM, a SAM-style FM trained with text and point prompting for organs and devices. Taking an innovative approach to text prompting in this space, FluoroSAM uses a large language model (LLM) to augment descriptions during training, ensuring flexibility and compatibility with diverse inputs.

2.1 A Large Scale Dataset for X-ray Image Analysis

First, we describe our DRR sampling strategy given a single human anatomical model defined by a CT image. We segment the CT using TotalSegmentator V2 [28], yielding 117 bone and soft tissue structures in the base model. We also segment the appendicular bones (11 classes) and body, which is used for view sampling, for a total of 128 organ masks. All segmentations are converted to meshes using marching cubes for downstream efficiency. Following this, we sample the detector size and source-to-detector distance D for a virtual C-arm uniformly in [180, 500] and [700, 1200]mm, respectively, based on commercially available systems. With a 10% probability, we instead sample a mini C-arm geometry, in the range of [100, 500] and [300, 400]mm, respectively. For each view, we sample a point of interest \mathbf{p} uniformly within the body mesh and a principle ray direction $\hat{\mathbf{r}}$ uniformly on the hemisphere, pointing in the anterior direction, omitting rays within 25° of the S/I axis. (These are not obtainable because the C-arm would collide with the table.) The X-ray source (or camera center) s is randomly offset from the point of interest to obtain the distribution of C-arm views shown in Fig. 1a.

For each view, then, we randomly place up to 15 surgical tools in the field of view, with a tendency toward fewer tools. The tools are chosen from among 464 tool mesh models available from [21] (111), collected from GrabCAD (296), or hand-modeled internally (57) and manually annotated with a text description, e.g. "cannulated 110mm screw with a 6.5mm diameter and 16mm thread." These are placed along the field of view with random location and orientation. We project digitally reconstructed radiographs (DRRs) from the patient CT and tool mesh models using a modified version of the DeepDRR simulator [26], which has been shown to support sim-to-real transfer for X-rays [8, 11, 15, 15]. Using this version, we simultaneously obtain realistic X-ray tansmission images and corresponding projected segmentations for each organ or tool present, enabling synchronous dataset generation of 448×448 images at a rate of ~ 4 images / s on an RTX 2080 Ti, using less than 4 GB of GPU memory. To increase throughput, which is rate-limited by the time required to obtain multi-organ segmentations in CTs, we parallelize simulation across 4 instances. Each annotation is associated with a base description, either from the TotalSegmentator class name or the original tool description.

At the time of submission the FluoroSAM dataset contains over 1.6M images synthesized from 1,699 CT scans. Because of compute constraints, this is only a portion of the CT scans available, including 747 scans from the New Mexico Decedent Image Database (410 head and neck, 37 torso, and 300 lower extremity scans) [6]; 108 scans of various regions from the TotalSegmentator dataset [28]; and 844 torso scans from internal collection at an emergency medicine department. We project up to 1000 images per scan, which are later filtered for tools that collide with the detector due to random orientations (roughly 5% of



Fig. 2: FluoroSAM's training pipeline, which uses an EfficientViT [20] image encoder, MedCLIP [27] text encoder, and GPT-3.5 for augmentation [2].



Fig. 3: When using only point prompts, SAM tends to over-segment fluoroscopic images. While an initial bounding box prompt helps, additional points simply confuse the network. MedSAM tends to segment an oval inside the prompt. FluoroSAM, by contrast, is able to propose a reasonable mask based only on text prompting and, crucially, incorporate new points to refine its prediction.

images). This results in a total of 1,672,969 X-ray images containing 63,625,805 masks, with 2 - 100 masks per image (\sim 38 on average), generated over approximately four weeks. This is comparable to the size of similar datasets for training foundation models, such as the MedSAM aggregated dataset of 1,570,263 imagemask pairs [22]. Finally, we split the FluoroSAM dataset according to an 90/5/5 split, so that no base CT scans overlap between training, validation, and testing.

2.2 A Language-aligned Foundation Model for X-ray Image Segmentation

Language alignment is crucial for the segment-anything task in X-ray. SAM and its variants predict three masks corresponding to the whole, part, or sub-part of an object, as in Fig. 1b, running backpropagation only for the mask with the lowest loss. X-ray images, on the other hand, often contain overlapping projections of many anatomical and non-anatomical objects. MedSAM [22] mitigates this ambiguity for other medical imaging domains by only allowing for bounding box prompts. For general X-ray imaging, however, this approach is undesirable because it (a) makes automatic or even non-expert prompting all but impossible, (b) precludes refinement of the initial mask using points, (c) still features significant ambiguity.

Thus, we train FluoroSAM with a text description as the first prompt, followed by points. We use the MedCLIP [27] text encoder with frozen weights to embed each description to 512, followed by a linear layer resizing it to the transformer dimension 256. To prevent overfitting on the MedCLIP embeddings of our original descriptions, we apply LLM-based augmentation to descriptions at training time, making up to 20 variations for each description by removing nonessential information or adding commands. For example, the class name "L4 vertebra" becomes (a) "Lumbar vertebrae 4," (b) "Fourth lumbar vertebrae," (c) "Identify the L4 vertebrae," (d) "Find the vertebrae labeled as L4," *etc.* We use the gpt-3.5-turbo from OpenAI to perform these augmentations [2], which are cached after the initial request. The full text of instructions and examples provided to the LLM for each prompt are provided in the supplement.

Because the effective resolution of our synthetic X-ray images is limited by the spatial resolution of the original CT images, we restrict our image size to 448^2 . Masks are predicted at 224^2 and upscaled via transposed convolution to the original image size, which is comparable to the original SAM output size. We train FluoroSAM from scratch using a b1 EfficientViT [20] image encoder and the frozen MedCLIP [27]. For subsequent prompts, we follow the strategy in the original SAM model, using the previous mask as an additional prompt, followed by up to 8 point embeddings. Following [8], we use strong domain randomization to enable sim-to-real transfer, including coarse dropout, inversion, blurring, Gaussian contrast adjustment [15], random windowing, and CLAHE histogram equalization. At inference time, we window the negative log transformed images to the 99th percentile and scale to [-1, 1]. No text augmentation is performed at inference. We use a batch size of 96 images with up to 16 masks per image, randomly chosen from among the available masks, and train for 6 epochs on 4 A100 GPUs with 80 GB of memory each. The initial learning rate was 0.001, reduced by a factor of 10 at epoch 3 and 5.

3 Evaluation

FluoroSAM Test Set. We first compare FluoroSAM's performance on a withheld test set of 47,190 synthetic X-rays. Given only a text prompt, FluoroSAM predicts a qualitatively reasonable mask and attains a 0.43 ± 0.26 DICE across all classes, with 0.85 ± 0.11 after further prompting. With two points, comparable to MedSAM and SAM's initial box prompt, it achieves superior performance, with a 0.68 ± 0.20 DICE compared to 0.60 ± 0.22 and 0.58 ± 0.19 , respectively. **Fluoroscopic Cadaver Study**. To explore FluoroSAM's performance on real X-ray images, we collect a dataset with 464 fluoroscopic images of the lower torso using the Brainlab Loop-X imaging device. Images for all cadaveric specimens were acquired from -30 to 30° in the cranial/caudal direction and -90 to 90° in the orbital angle, with random spacing. These images were registered automatically via optical tracking with the Brainlab Curve platform, which enabled Organ Segmentation on Cadaveric Lower Torso Images



Fig. 4: FluoroSAM incorporates prompts to refine its prediction for real X-ray images. MedSAM and SAM both achieve comparable performance to FluoroSAM with box prompts, but they do not improve with further prompting.

gold-standard 2D/3D registration with CBCT scans taken before and after fluoroscopic acquisitions. We then obtain 3D organ segmentations using TotalSegmentator [28] and project 2D masks onto each image as shown in Fig. 3 and 5a (GT). In some cases, TotalSegmentator struggled with soft tissue structures, possibly because of natural decomposition not present in its training set.

We find that FluoroSAM is an effective tool for segmenting structures in real fluoroscopic images. We observe qualitatively reasonable masks based on text prompts alone, with a DICE score of 0.39 on hard tissues and 0.26 overall. FluoroSAM is able to effectively incorporate new information from point prompts, achieving a DICE of 0.9 ± 0.15 on hard and 0.73 ± 0.15 on soft tissue, whereas SAM does not effectively learn from new prompts, as evident in Fig 3. Likewise MedSAM, although trained on diagnostic X-ray images, consistently predicts an oval-shaped mask inside the bounding box prompt, achieving a DICE of 0.51 ± 0.21 overall. This trend can be seen in Fig. 4, where the performance of FluoroSAM aligns with MedSAM and SAM when comparable information is available, but only FluoroSAM is able to incorporate new information. **Zero-shot Lung Segmentation on Chest X-rays**. FluoroSAM's language

alignment allows for segmentation of chest X-rays. FitoroSAM's language alignment allows for segmentation of objects not seen during training, such as whole long segmentation in Chest X-rays [3]. Using only text prompts, FluoroSAM proposes a reasonable segmentation of either lung with 0.52 ± 0.21 DICE, despite never training on the full lung class, only the individual lobes. Further refinement with point prompts enables 0.90 ± 0.04 DICE, as seen in Fig. 5b.

4 Discussion and Conclusion

The generalizability of FMs makes it tempting to apply them wherever suitable datasets can be found. Here, we make the case for a FM seemingly out of step with the current trend, *i.e.* to ingest as much data as possible, and instead focus on a single imaging modality. Yet, X-ray imaging of the full human body and from various viewing directions represents an enormous variety of images



Fig. 5: (a) FluoroSAM's failure modes on text-only prompts include fine details and soft tissue structures, but it exhibits promising zero- and few-shot performance on whole lung segmentation, not included during training.

with markedly different features than visible light or tomographic imaging. The potential applications of a model like FluoroSAM are significant, as exemplified by proliferation of specialized models for X-ray image analysis in chest X-ray diagnosis [3, 5, 24], dental exams [25], forensics [23], intelligent surgical systems [11–16], and AI-driven educational curricula [17, 19]. Within this broad spectrum of applications, the benefits of language alignment for X-ray imaging have so far been limited to diagnostic systems, where text descriptions are available as the byproduct of routine clinical workflows [5, 10]. The success of these systems demonstrates the potential for language-aligned models to analyze X-ray images, and we hope that by adding a large scale, full-body dataset, we will spur further innovation in this space.

Given the large scale data available, the choice of a SAM framework is sensible but not certain. On one hand, it allows for training strategies that focus on known objects, by only providing true positive masks during training. This avoids learning on false negatives in the data, such as implants not segmented by TotalSegmentator [28], but it does limit text-only prompts to small variations on the 128 organs and 464 tools seen during training. More complicated prompts like "Show me the bone fragment below that screw," necessitate first recognizing the desired object (a bone fragment rather than a screw), identifying the disambiguating object (the screw), and finally segmenting a bone fragment, which is not among the data available during training. We use LLM-based text augmentation to avoid overfitting on the finite number of descriptions available from the FluoroSAM dataset, but this approach would not support full language understanding. Rather, FluoroSAM allows the user to express the kind of desired object—a bone rather than an implant or soft tissue—which enables unambiguous point-based prompting.

5 Conclusion

We have introduced FluoroSAM, a foundation model trained from scratch on 63M masks for X-ray image segmentation. We have evaluated its performance on synthetic and real cadaveric X-ray images as well as its zero-shot generalizability for lung segmentation on chest X-ray images, of which it saw no examples during training. Thanks to its built-in language alignment, FluoroSAM is able to automatically segment anatomical and non-anatomical objects with reasonable accuracy, and its compatibility with point-based prompting allows for further refinement. Although out initial models show room for improvement, possibly due to lack of convergence at very small learning rates, we are continuing to train on the same and newly simulated data, and we look forward to making foundational-grade models available to the community in X-ray-enabled health-care research.

Acknowledgements

This work was supported by the Link Foundation Fellowship for Modeling, Training, and Simulation; the NIH under Grant No. R21EB028505, the NSF under Award No. 2239077, and Johns Hopkins University Internal Funds. Thank you to Lynne Jones, Wenhao Gu, Justin Opfermann, and Keshuai Xu for assistance with cadaver studies. An enormous thank you, also, to the Cloud & Research IT Services group in IT@JH for their technical and computational support.

Bibliography

[1] Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J.Q., Demszky, D., Donahue, C., Doumbouya, M., Durmus, E., Ermon, S., Etchemendy, J., Ethayarajh, K., Fei-Fei, L., Finn, C., Gale, T., Gillespie, L., Goel, K., Goodman, N., Grossman, S., Guha, N., Hashimoto, T., Henderson, P., Hewitt, J., Ho, D.E., Hong, J., Hsu, K., Huang, J., Icard, T., Jain, S., Jurafsky, D., Kalluri, P., Karamcheti, S., Keeling, G., Khani, F., Khattab, O., Koh, P.W., Krass, M., Krishna, R., Kuditipudi, R., Kumar, A., Ladhak, F., Lee, M., Lee, T., Leskovec, J., Levent, I., Li, X.L., Li, X., Ma, T., Malik, A., Manning, C.D., Mirchandani, S., Mitchell, E., Munyikwa, Z., Nair, S., Narayan, A., Narayanan, D., Newman, B., Nie, A., Niebles, J.C., Nilforoshan, H., Nyarko, J., Ogut, G., Orr, L., Papadimitriou, I., Park, J.S., Piech, C., Portelance, E., Potts, C., Raghunathan, A., Reich, R., Ren, H., Rong, F., Roohani, Y., Ruiz, C., Ryan, J., Ré, C., Sadigh, D., Sagawa, S., Santhanam, K., Shih, A., Srinivasan, K., Tamkin, A., Taori, R., Thomas, A.W., Tramèr, F., Wang, R.E., Wang, W., Wu, B., Wu, J., Wu, Y., Xie, S.M., Yasunaga, M., You, J., Zaharia, M., Zhang, M., Zhang, T., Zhang, X., Zhang, Y., Zheng, L., Zhou, K., Liang, P.: 10 B.D. Killeen et al.

On the Opportunities and Risks of Foundation Models. arXiv (Aug 2021). https://doi.org/10.48550/arXiv.2108.07258

- [2] Brown, T.B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D.M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., Amodei, D.: Language Models are Few-Shot Learners. arXiv (May 2020). https://doi.org/10.48550/arXiv.2005.14165
- [3] Candemir, S., Jaeger, S., Palaniappan, K., Musco, J.P., Singh, R.K., Xue, Z., Karargyris, A., Antani, S., Thoma, G., McDonald, C.J.: Lung segmentation in chest radiographs using anatomical atlases with nonrigid registration. IEEE Trans. Med. Imaging 33(2), 577–590 (Feb 2014). https://doi.org/10.1109/TMI.2013.2290491
- [4] Çallı, E., Sogancioglu, E., van Ginneken, B., van Leeuwen, K.G., Murphy, K.: Deep learning for chest X-ray analysis: A survey. Med. Image Anal. 72, 102125 (Aug 2021). https://doi.org/10.1016/j.media.2021.102125
- [5] Chen, Z., Varma, M., Delbrouck, J.B., Paschali, M., Blankemeier, L., Van Veen, D., Valanarasu, J.M.J., Youssef, A., Cohen, J.P., Reis, E.P., Tsai, E.B., Johnston, A., Olsen, C., Abraham, T.M., Gatidis, S., Chaudhari, A.S., Langlotz, C.: CheXagent: Towards a Foundation Model for Chest X-Ray Interpretation. arXiv (Jan 2024). https://doi.org/10.48550/arXiv.2401.12208
- [6] Edgar, H., Daneshvari Berry, S., Moes, E., Adolphi, N., Bridges, P., Nolte, K.: New mexico decedent image database. Office of the Medical Investigator, University of New Mexico (2020). https://doi.org/10.25827/5s8c-n515
- [7] Gaggion, N., Mosquera, C., Mansilla, L., Aineseder, M., Milone, D.H., Ferrante, E.: CheXmask: a large-scale dataset of anatomical segmentation masks for multi-center chest x-ray images. arXiv (Jul 2023). https://doi.org/10.48550/arXiv.2307.03293
- [8] Gao, C., Killeen, B.D., Hu, Y., Grupp, R.B., Taylor, R.H., Armand, M., Unberath, M.: Synthetic data accelerates the development of generalizable learning-based algorithms for X-ray image analysis. Nat. Mach. Intell. 5(3), 294–308 (Mar 2023). https://doi.org/10.1038/s42256-023-00629-1
- [9] He, S., Bao, R., Li, J., Stout, J., Bjornerud, A., Grant, P.E., Ou, Y.: Computer-Vision Benchmark Segment-Anything Model (SAM) in Medical Images: Accuracy in 12 Datasets. arXiv (Apr 2023). https://doi.org/10.48550/arXiv.2304.09324
- [10] Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghgoo, B., Ball, R., Shpanskaya, K., Seekins, J., Mong, D.A., Halabi, S.S., Sandberg, J.K., Jones, R., Larson, D.B., Langlotz, C.P., Patel, B.N., Lungren, M.P., Ng, A.Y.: CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison. arXiv (Jan 2019). https://doi.org/10.48550/arXiv.1901.07031
- [11] Kausch, L., Thomas, S., Kunze, H., Privalov, M., Vetter, S., Franke, J., Mahnken, A.H., Maier-Hein, L., Maier-Hein, K.: Toward automatic C-arm

positioning for standard projections in orthopedic surgery. Int. J. CARS 15(7), 1095–1105 (Jul 2020). https://doi.org/10.1007/s11548-020-02204-0

- Killeen, B.D., Chakraborty, S., Osgood, G., Unberath, M.: Toward perception-based anticipation of cortical breach during K-wire fixation of the pelvis. In: Proceedings Volume 12031, Medical Imaging 2022: Physics of Medical Imaging, vol. 12031, pp. 410–415. SPIE (Apr 2022). https://doi.org/10.1117/12.2612989
- [13] Killeen, B.D., Chaudhary, S., Osgood, G., Unberath, M.: Take a shot! natural language control of robotic x-ray systems for image-guided surgery. IPCAI (to appear) (Jun 2024)
- [14] Killeen, B.D., Cho, S.M., Armand, M., Taylor, R.H., Unberath, M.: In silico simulation: a key enabling technology for next-generation intelligent surgical systems. Prog. Biomed. Eng. 5(3), 032001 (May 2023). https://doi.org/10.1088/2516-1091/acd28b
- [15] Killeen, B.D., Gao, C., Oguine, K.J., Darcy, S., Armand, M., Taylor, R.H., Osgood, G., Unberath, M.: An autonomous X-ray image acquisition and interpretation system for assisting percutaneous pelvic fracture fixation. Int. J. CARS pp. 1–8 (May 2023). https://doi.org/10.1007/s11548-023-02941-y
- [16] Killeen, B.D., Zhang, H., Mangulabnan, J., Armand, M., Taylor, R.H., Osgood, G., Unberath, M.: Pelphix: Surgical Phase Recognition from X-Ray Images in Percutaneous Pelvic Fixation. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2023, pp. 133–143. Springer, Cham, Switzerland (Oct 2023). https://doi.org/10.1007/978-3-031-43996-4_13
- [17] Killeen, B.D., Zhang, H., Wang, L., Liu, Z., Kleinbeck, C., Rosen, M., Taylor, R.H., Osgood, G., Unberath, M.: Stand in surgeon's shoes: Virtual reality cross-training to enhance teamwork in surgery. IPCAI (to appear) (Jun 2024)
- [18] Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment Anything. arXiv (Apr 2023). https://doi.org/10.48550/arXiv.2304.02643
- [19] Kleinbeck, C., Zhang, H., Killeen, B.D., Roth, D., Unberath, M.: Neural digital twins: Reconstructing complex medical environments for spatial planning in virtual reality. IPCAI (to appear) (Jun 2024)
- [20] Liu, X., Peng, H., Zheng, N., Yang, Y., Hu, H., Yuan, Y.: Efficientvit: Memory efficient vision transformer with cascaded group attention. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14420–14430 (June 2023)
- [21] Luijten, G., Gsaxner, C., Li, J., Pepe, A., Ambigapathy, N., Kim, M., Chen, X., Kleesiek, J., Hölzle, F., Puladi, B., Egger, J.: 3D surgical instrument collection for computer vision and extended reality. Sci. Data 10(796), 1– 12 (Nov 2023). https://doi.org/10.1038/s41597-023-02684-0
- [22] Ma, J., He, Y., Li, F., Han, L., You, C., Wang, B.: Segment anything in medical images. Nat. Commun. 15(654), 1–9 (Jan 2024). https://doi.org/10.1038/s41467-024-44824-z

- 12 B.D. Killeen et al.
- [23] Milošević, D., Vodanović, M., Galić, I., Subašić, M.: Automated estimation of chronological age from panoramic dental X-ray images using deep learning. Expert Syst. Appl. 189, 116038 (Mar 2022). https://doi.org/10.1016/j.eswa.2021.116038
- [24] Mohammad-Rahimi, H., Nadimi, M., Ghalyanchi-Langeroudi, A., Taheri, M., Ghafouri-Fard, S.: Application of Machine Learning in Diagnosis of COVID-19 Through X-Ray and CT Images: A Scoping Review. Front. Cardiovasc. Med. 8, 638011 (Mar 2021). https://doi.org/10.3389/fcvm.2021.638011
- [25] Silva, G., Oliveira, L., Pithon, M.: Automatic segmenting teeth in X-ray images: Trends, a novel data set, benchmarking and future perspectives. Expert Syst. Appl. 107, 15–31 (Oct 2018). https://doi.org/10.1016/j.eswa.2018.04.001
- [26] Unberath, M., Zaech, J.N., Lee, S.C., Bier, B., Fotouhi, J., Armand, M., Navab, N.: DeepDRR – A Catalyst for Machine Learning in Fluoroscopy-Guided Procedures. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2018, pp. 98–106. Springer, Cham, Switzerland (Sep 2018). https://doi.org/10.1007/978-3-030-00937-3_12
- [27] Wang, Z., Wu, Z., Agarwal, D., Sun, J.: MedCLIP: Contrastive Learning from Unpaired Medical Images and Text. arXiv (Oct 2022). https://doi.org/10.48550/arXiv.2210.10163
- [28] Wasserthal, J., Breit, H.C., Meyer, M.T., Pradella, M., Hinck, D., Sauter, A.W., Heye, T., Boll, D.T., Cyriac, J., Yang, S., Bach, M., Segeroth, M.: TotalSegmentator: Robust Segmentation of 104 Anatomic Structures in CT Images. Radiology: Artificial Intelligence (Jul 2023), https://pubs.rsna.org/ doi/10.1148/ryai.230024
- [29] Zhao, W., Shen, L., Islam, M.T., Qin, W., Zhang, Z., Liang, X., Zhang, G., Xu, S., Li, X.: Artificial intelligence in image-guided radiotherapy: a review of treatment target localization. Quantitative Imaging in Medicine and Surgery 11(12), 4881 (Dec 2021). https://doi.org/10.21037/qims-21-199

Text Augmentation Details.

The full text of instructions for text augmentation prompting are below:

- You are a data scientist trying to create a consistent but varied set of descriptions for training a textaligned deep neural network for detecting and segmenting objects in X-ray images. In subsequent messages, I will give you a comprehensive description of an anatomical or surgical object, such as an organ or an instrument. Your job is to generate a set of descriptions that are consistent with the original description but varied enough to provide a diverse training set. There are several ways to do this, each of which you should use, either alone or in combination, to generate a diverse set of descriptions . You should also use your creativity to come up with other methods for generating diverse descriptions. Here are some examples of the methods you should use:
- Remove some or all of the nonessential information about the object.
- Add some nonessential information that would not confuse the network.
- Stay close to the original description, but change some of the nonessential information.
- Convert the description to a command, such as "Find the [object]," "Identify the [object]," or any other similar commands.
- Convert the description to a question, such as "What is the [object]?," "Where is the [object]?," or any other similar questions.
- When responding to each query, list 5-20 variations in plain text, separated by newlines. Do not explain your decisions. Do not provide any additional text, other than the varied descriptions. Do not number the descriptions or use bullet points to separate them. Putting them on different lines is fine.
- Finally, it is okay if you feel you do not have enough information, I'm sure you will do your best as the expert surgical data annotator that you are.

Here are some examples to get you started:

```
2 Anonymous
```

```
input:
cannulated 110mm screw with a 6.5mm diameter and 16mm
   thread
output:
cannulated screw
Show me the cannulated screw.
screw
The one-ten screw.
Where is the screw?
Where is the long screw?
Find the screw with the 16mm thread.
Find the screw.
The orthopaedic screw.
input:
orthopedic drill with a K-wire adaptor
orthopedic drill
drill
the drill
Show me the drill.
drill with K-wire adaptor
Orthopedic drill with adapter
Drill for orthopedic procedure
. . .
```



Fig. 1: On synthetic X-ray images from a wide variety of views onto the full body, FluoroSAM outperforms SAM and MedSAM, and it shows the ability to improve based on additional information.