

SurgRFO: Foundation Model Based Compositional Synthesis of Critical Retained Foreign Objects in Intraoperative Chest X-rays

Yuanyun Hu^{1,2,3*}, Yuli Wang^{2,3*}, Noemi Acevedo Rodriguez², Ronald Yang², Wen-Chi Hsu⁴, Siwei Luo⁵, Zihao Bai¹, Jing Wu⁶, Yuwei Dai², Shaoju Wu², Jonathon Lindquist², Justin Honce², Premal Trivedi², Zhicheng Jiao⁷, Ihab Kamel², Elliott Haut³, Pamela Johnson³, John Eng³, Cheng Ting Lin³, Nan Su¹, Bo Chen^{8*}, Sun Yu⁹, and Harrison Bai^{2*}

¹ Department of Electronic Engineering, Tsinghua University, China

² Department of Radiology, University of Colorado Denver Anschutz Medical Campus, USA

³ Department of Radiology and Radiological Science, Johns Hopkins University School of Medicine, Baltimore, USA

⁴ Department of Medical Imaging and Intervention, Chang Gung Memorial Hospital at Linkou, Taiwan

⁵ Department of Electronic Engineering, Southwest University of Science and Technology, China

⁶ Department of Radiology, Second Xiangya Hospital, Central South University, Changsha, Hunan, China

⁷ Department of Diagnostic Imaging, Brown University Health, Providence, USA

⁸ School of Clinical Medicine, Tsinghua University, China

⁹ Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, USA

cba04570@btch.edu.cn and harrison.bai@cuanschutz.edu

Abstract. Critical retained foreign objects (RFOs) on intraoperative chest radiographs are rare but high-risk events. Their scarcity limits robust automated detection model training and generalization. We introduce **SurgRFO**, a two-stage synthesis framework for generating realistic RFO-present intraoperative chest X-rays. In **Stage 1**, a Roentgen chest X-ray foundation model is fine-tuned on surgical-domain images to generate realistic RFO-free backgrounds that preserve anatomy, indwelling lines and tubes, and intraoperative imaging characteristics. In **Stage 2**, a lightweight generator trained on localized RFO patches from limited positive cases synthesizes diverse RFO instances, which are composited onto generated backgrounds using conditional Poisson fusion to improve photometric consistency. We evaluate SurgRFO through (i) a blinded clinician study assessing realism and clinical plausibility, and (ii) downstream detection experiments in which synthesized data are used to augment Faster R-CNN, YOLOv8, and RetinaNet. SurgRFO consistently improves sensitivity at low false-positive-per-image

* These authors contributed equally to this work.

(FPPI) operating points on internal and external test sets. Clinician ratings indicate that the synthesized images achieve realism comparable to real intraoperative images. Ablation analyses further examine fusion strategies and synthesis scale. Ethical safeguards for synthetic surgical data are also discussed. The code is publicly available on GitHub at <https://github.com/Yuanyun936/SurgRFO>.

Keywords: Image Synthesis · Critical Retained Foreign Objects · Data Augmentation · Surgical X-rays.

1 Introduction

Critical retained foreign objects (RFOs), such as surgical sponges and needle fragments, are rare but high-impact patient safety events. Unlike non-critical, intentionally placed objects, critical RFOs are unintended remnants that may cause severe complications, re-operations, and substantial medico-legal and financial burden. [18, 8, 17, 13, 16, 22, 23]. Clinically, intraoperative imaging aims to rule out RFOs before procedure completion, where missed detections can have serious consequences. However, work explicitly targeting critical RFOs remains limited. Most prior studies focus on standardized chest X-rays (CXRs) and predominantly address *non-critical* objects [30, 24, 21, 10]. Robust detection of critical RFOs in real-world intraoperative X-rays, therefore, remains an open and clinically significant challenge.

Critical RFO detection relies on intraoperative radiographs acquired under far less controlled conditions than standard CXRs, with cluttered scenes, overlapping instruments, variable anatomy exposure, and heterogeneous imaging settings. Two major challenges arise. First, visual cues are often subtle and confounded by surgical context. Sponges or needle fragments may be blood-soaked, partially occluded, or blended with surrounding tissue, producing weak signals that are difficult to recognize, even for experienced clinicians. [11, 14]. False negatives are therefore not uncommon, and prior studies suggest that intraoperative imaging may fail to detect a substantial fraction of RFO cases [2]. Second, positive cases are extremely rare and difficult to scale. Even at a large hospital, only low hundreds of cases may accumulate over decades, and to our knowledge, no public dataset of intraoperative radiographs dedicated to critical RFO detection exists. Together, these constraints create a fundamental limitation: accurate detection requires sensitivity to subtle local patterns, yet available data provide limited exposure to such events, challenging conventional supervised learning approaches.

Synthetic data augmentation is a natural strategy to mitigate data scarcity. Physics-based RFO synthesis, such as CT-projection pipelines with rendered 3D objects [31], provides controllable images with precise annotations. However, these approaches typically insert object models into relatively clean CT volumes and project them to X-ray images, leaving a substantial domain gap to real intraoperative radiographs acquired in cluttered surgical workflows. Recent advances in diffusion models (DMs), which generate images through iterative denoising of




Fig. 1. Overview of the **Intraoperative RFO Bench** and **SurgRFO** pipeline. (a) Multi-institutional intraoperative radiographs are curated and expert-annotated as real positive (critical RFO) and negative cases. (b) SurgRFO applies a two-stage synthesis framework, combining surgical X-ray background generation with localized RFO sampling and fusion to create annotated synthetic images. (c) Synthetic data are used for training augmentation and evaluated on held-out real cases for critical RFO detection and classification. (d) A radiologist-led Turing test assesses visual realism through confidence-based quality ratings and image discrimination tasks.

a learned reverse process [5], have enabled high-fidelity and diverse X-ray synthesis [1, 27, 29, 25, 31], supporting augmentation for imbalanced medical imaging tasks [9, 15, 7, 12, 6, 28]. However, fully end-to-end generation of critical RFO radiographs remains unreliable. The same rarity and subtle visual signatures that hinder detection also limit generative fidelity: standard pipelines struggle to simultaneously preserve complex surgical context and localized RFO cues, often leading to visually implausible samples for training detectors (see comparisons in Section 3). To address this gap, we propose **SurgRFO**, a two-stage synthesis framework for generating surgical radiographs with critical RFOs. Stage 1 learns a strong prior over cluttered surgical radiographs by training a latent diffusion model (LDM) as a generic surgical X-ray generator. Stage 2 trains a lightweight U-Net to model local critical RFO appearance, then conditionally fuses the sampled RFO content into Stage 1 backgrounds via a Poisson-based blending mechanism guided by a prior over plausible RFO locations and sizes.

We also release **Intraoperative RFO Bench**, to our knowledge the first publicly available real de-identified intraoperative chest radiograph dataset dedicated to *critical* RFOs, comprising 144 expert-annotated positive cases curated from 18 years of surgical radiographs from two health systems (Johns Hopkins Hospital (JHH) and University of Colorado (CU)). Augmenting training with SurgRFO yields consistent improvements across multiple detector families. In a radiologist-led Turing-style evaluation, synthetic images are frequently indistinguishable from real cases. Together, these results demonstrate the value of

releasing this benchmark and the effectiveness of the proposed framework for advancing critical RFO detection.

2 Method

2.1 Dataset Curation

Under institutional review board approval, we retrospectively curated surgical chest radiographs from the JHH, comprising 144 *critical* RFO-positive cases and 944 negative cases with no RFOs. An additional 20 *critical* RFO-positive cases were collected from CU for external evaluation. Positives were identified via structured report search and expert verification using radiology reports and operative notes. All images were de-identified and annotated by trained radiologists under a standardized protocol, including image-level labels and object-level annotations (bounding boxes or polygons), with each foreign object categorized as critical or non-critical. Patient-level splits into training, validation, and test sets were applied consistently across cohorts to prevent information leakage. Training images were preprocessed with center cropping, resolution normalization, and artifact filtering.

2.2 SurgRFO

Let $\Omega \subset \mathbb{R}^2$ denote the image lattice, and let $I \in \mathbb{R}^{H \times W}$ denote a radiograph defined on Ω . Two datasets are considered, $\mathcal{D}_{\text{bg}} = \{I_i\}_{i=1}^{N_{\text{bg}}}$ and $\mathcal{D}_{\text{rfo}} = \{(I_j, M_j, y_j)\}_{j=1}^{N_{\text{rfo}}}$. The set $\mathcal{D}_{\text{bg}}$ contains surgical-domain negative radiographs that provide RFO-free backgrounds. The set $\mathcal{D}_{\text{rfo}}$ contains radiographs with annotated critical RFOs. The mask $M_j \subset \Omega$ specifies the RFO support, and the label $y_j \in \mathcal{Y}$ specifies the semantic type.

SurgRFO is a two-stage synthesis framework that decouples global background modeling from local critical RFO modeling. Stage 1 learns a generic prior over surgical radiograph backgrounds. Stage 2 learns a local sampler for critical RFO appearance and composes sampled RFO evidence onto Stage 1 backgrounds through a controlled fusion mechanism.

Stage 1: Generic Surgical Radiograph Generator. A background generator is defined as a conditional LDM f_θ parameterized by θ . Background sampling is written as

$$\hat{I}_{\text{bg}} \sim p_\theta(I \mid t, z_1), \quad z_1 \sim \mathcal{N}(0, I), \quad (1)$$

where $\hat{I}_{\text{bg}}$ denotes an RFO-free surgical background, t denotes a domain anchor text token, and z_1 denotes the diffusion noise seed. Stage 1 is designed for domain-level background modeling rather than instance-specific textual control. A fixed domain anchor prompt is therefore used, and diversity is induced primarily by the stochasticity of z_1 . In implementation, we adapt RoentGen [1], an advanced CXRs LDM, to surgical X-ray domain leveraging $\mathcal{D}_{\text{bg}}$.

Stage 2: Local RFO Sampler and Conditional Fusion. A local RFO sampler is defined by a compact generator g_ϕ that models patch-level RFO appearance. For each annotated sample $(I, M, y) \in \mathcal{D}_{\text{rfo}}$, an RFO-centered canonical patch is extracted by a mask-driven cropping operator $\mathcal{C}_M$,

$$P = \mathcal{C}_M(I). \quad (2)$$

A plausible RFO patch is sampled from a noise seed z_2 as

$$\hat{P}_{\text{rfo}} \sim p_\phi(P \mid z_2), \quad z_2 \sim \mathcal{N}(0, I), \quad (3)$$

where p_ϕ denotes the implicit patch distribution defined by g_ϕ .

An anatomically plausible placement is specified by a location-scale variable $\ell = (x, y, s)$, where (x, y) denotes the insertion center on the background canvas and s denotes the insertion scale. An empirical placement prior $\hat{p}_{\text{anno}}(\ell)$ is estimated from annotation statistics, and a concrete placement parameter is drawn as $\hat{\ell} \sim \hat{p}_{\text{anno}}(\ell)$. A placement operator $\mathcal{T}_{\hat{\ell}}$ maps both the sampled patch and its support to the full image grid,

$$\hat{P}_{\hat{\ell}} = \mathcal{T}_{\hat{\ell}}(\hat{P}_{\text{rfo}}), \quad M_{\hat{\ell}} = \mathcal{T}_{\hat{\ell}}(S(\hat{P}_{\text{rfo}})), \quad (4)$$

where $S(\cdot)$ produces a coarse support mask in patch coordinates and $M_{\hat{\ell}}$ denotes the corresponding insertion support on the $H \times W$ canvas.

The final composite image I_{syn} is defined by an aligned mixed-cloning objective over the insertion region $\Omega_{\hat{\ell}} = \{\mathbf{u} \in \Omega : M_{\hat{\ell}}(\mathbf{u}) = 1\}$,

$$I_{\text{syn}} = \arg \min_I \sum_{\mathbf{u} \in \Omega_{\hat{\ell}}} \left\| \nabla I(\mathbf{u}) - \left(\beta \nabla \mathcal{A}_\alpha(\hat{P}_{\hat{\ell}})(\mathbf{u}) + (1 - \beta) \nabla \hat{I}_{\text{bg}}(\mathbf{u}) \right) \right\|_2^2, \quad (5)$$

where ∇ denotes the spatial gradient operator. The parameter $\beta \in [0, 1]$ controls the strength of mixed cloning by balancing patch-driven gradients and background-driven gradients. The operator $\mathcal{A}_\alpha$ performs local intensity alignment of the placed patch to the background within $M_{\hat{\ell}}$. The parameter $\alpha \in [0, 1]$ controls the degree of alignment by interpolating between the raw patch intensity and an intensity level matched to the local background statistics.

3 Experiments and Results

Experimental setup. SurgRFO is evaluated on the Intraoperative RFOs Bench, which is split into training, validation, and test sets with a 6:1:3 patient-level ratio. In Stage 1, we instantiate the generator from RoentGen under the Stable Diffusion (SD) latent diffusion formulation [19], and adapt it to the surgical background cohort using a fixed domain prompt and varying noise seeds for diversity. Stage 2 trains a compact U-Net diffusion sampler [3] on RFO-centered crops from the positive cohort and synthesizes RFO-positive composites via conditional fusion. We evaluate SurgRFO from: (i) synthesis quality, and (ii) the




Fig. 2. Visualization of synthesis results. Each row shows one representative case. Columns (a-b) show Stage-1 RFO-negative background generation using the original RoentGen and DreamBooth fine-tuning. Columns (c-d) compare Stage-2 RFO-positive synthesis using inpainting Stable Diffusion versus the proposed SurgRFO framework, with real positive cases shown in (e) for reference. Black insets provide zoomed views; red boxes mark target regions, and turquoise overlays indicate critical RFOs. SurgRFO maintains surgical radiograph fidelity while generating clinically plausible critical RFO appearances.

impact of synthetic augmentation on downstream critical RFO detection. We use FID [4] and MS-SSIM [26] for qualifying synthesis quality. Baselines include DreamBooth fine-tuning [20], and an inpainting SD [19]. RoentGen serves as the common initialization for both the baselines and SurgRFO. Downstream performance is evaluated on held-out real radiographs using detection metrics, including mean average precision (mAP), false negative rate (FNR), false omission rate (FOR), free-response receiver operating characteristic (FROC), and binary classification accuracy (ACC) for positive versus negative prediction.

3.1 Synthesis Quality

Table 1 evaluates Stage-1 RFO-negative synthesis on 5k images. Compared with the original RoentGen SD and a DreamBooth-adapted baseline, SurgRFO achieves consistently lower FID across checkpoints, with the best fidelity at 2k steps (FID 43.675). While DreamBooth improves domain alignment compared to the original RoentGen Model, it remains inferior to SurgRFO. Diversity improves with training and is best at 4k steps (lowest MS-SSIM), suggesting that Stage 1 captures surgical radiograph appearance with a selectable fidelity-diversity trade-off along training. Table 2 assesses final RFO-positive composites after Stage 2. SurgRFO improves fidelity over both the original RoentGen and inpainting baselines. Because Stage 2 edits only localized regions, MS-SSIM is dominated by preserved backgrounds and is reported as a secondary diagnostic, while FID remains the primary criterion for comparison. Visual comparisons

Table 1. Quantitative assessment of image fidelity and diversity for Stage-1 RFO-negative synthesis. FID and MS-SSIM for surgical background generation. **Bold** indicates the best, and underlined indicates the second best among adapted models across training steps.

Method	Steps	FID ↓	MS-SSIM ↓
Original RoentGen SD (no fine-tuning)	–	175.167	0.262 ± 0.130
DreamBooth SD (few-shot fine-tuning)	–	80.963	0.424 ± 0.085
Ours	1k	48.568	0.406 ± 0.071
	2k	43.675	0.474 ± 0.079
	3k	<u>46.100</u>	0.460 ± 0.052
	4k	57.896	0.362 ± 0.054
	5k	47.948	<u>0.393 ± 0.059</u>

Table 2. Quantitative assessment for Stage-2 RFO-positive synthesis. **Bold** indicates the best FID among compared methods.

Method	FID ↓	MS-SSIM ↓
Original RoentGen SD	197.837	0.262 ± 0.130
Inpainting SD	154.523	0.118 ± 0.045
Ours	69.178	0.480 ± 0.064

(Fig. 2) further demonstrate that end-to-end baselines struggle to maintain radiographic realism and render subtle RFO cues under limited positive data. In contrast, SurgRFO preserves global surgical fidelity and produces coherent, clinically plausible critical RFO appearances, as evidenced by the zoomed-in insets.

3.2 Downstream Detection with SurgRFO Augmentation

Internal evaluation. The downstream evaluation assesses whether SurgRFO improves critical RFO detection on real radiographs. As shown in Table 3, augmenting training with synthetic positives yields consistent gains across three detector families, suggesting that synthetic samples broaden effective training coverage rather than favoring a specific detector. Faster R-CNN demonstrates marked performance improvement with a substantial reduction in missed positive cases. RetinaNet shows monotonic gains as more synthetic samples are added. YOLOv8l performs poorly when trained on the base regime alone but achieves meaningful detection once synthetic samples are introduced, indicating particular benefit under limited positive supervision.

External evaluation. Generalization is further evaluated on an external real-world cohort from the CU. After SurgRFO-based augmentation, all models demonstrate clear improvement, with consistent reductions in false negatives and gains in overall accuracy. RetinaNet shows the largest benefit, with mAP@0.3 improving by +0.260, FNR decreasing by -0.750, FOR decreasing by -0.321, FROC increasing by +0.083, and ACC improving by +0.428, transitioning to meaningful retrieval performance. Faster R-CNN and YOLOv8l also exhibit reduced FNR (-0.250 each) and improved accuracy (+0.142 each), with YOLOv8l

Table 3. Downstream critical RFO detection with synthetic augmentation. Detectors are trained on the base real set with additional SurgRFO-generated positive samples at varying amounts. **Bold** marks the best result and underlined marks the second best within each detector block.

Model	Training Set	mAP@0.3 $\uparrow$	FNR $\downarrow$	FOR $\downarrow$	FROC $\uparrow$	ACC $\uparrow$
Faster R-CNN	Base	0.184	78.79%	8.72%	0.242	0.672
	Base+500	0.304	66.67%	7.48%	0.424	0.928
	Base+1000	0.390	60.61%	6.87%	<u>0.455</u>	<u>0.931</u>
	Base+2000	0.510	33.33%	3.93%	0.636	0.954
RetinaNet	Base	0.099	72.73%	8.11%	0.455	0.921
	Base+500	0.364	63.64%	7.17%	0.485	0.931
	Base+1000	0.455	54.55%	6.23%	0.667	0.938
	Base+2000	0.564	36.36%	4.26%	0.727	0.954
YOLOv8l	Base	0.000	100.00%	10.82%	–	0.892
	Base+500	0.290	69.70%	7.85%	0.303	0.918
	Base+1000	0.357	60.61%	6.87%	<u>0.394</u>	0.931
	Base+2000	<u>0.355</u>	63.64%	<u>7.24%</u>	0.394	<u>0.921</u>

Table 4. Radiologist-driven Turing tests for clinical realism.

(a) Blinded image-quality ratings on a 1–5 Likert scale. Values are mean±standard deviation. Welch’s test compares real and synthetic images per rater.				(b) Discrimination between real and synthetic images. Metrics include accuracy, F1 score, the fraction of synthetic images labeled as real, and confidence–accuracy AUC.				
	Real	Synthetic	Welch p		Acc	F1	Syn→Real	ConfAUC
Rater 1	2.86±1.11	2.47±1.02	0.166	Rater 1	0.750	0.762	0.30	0.558
Rater 2	2.64±1.22	2.31±0.74	0.220	Rater 2	0.750	0.750	0.25	0.495

achieving a notable FROC gain (+0.250). These findings indicate that synthesized positives provide effective supervision beyond the development cohort and enhance robustness under dataset shift.

4 Radiologist-driven Turing Test

To assess clinical plausibility beyond automated metrics, we conducted a blinded radiologist evaluation using a Turing-style protocol (Fig. 1 d) comprising two tasks: (1) blinded image quality rating without disclosure of provenance, and (2) explicit discriminability between real and synthetic images.

Turing Test 1: Blinded image quality rating. Radiologists independently scored randomized radiographs using a five-point Likert rubric. As shown in Table 4 (a), synthetic images received quality ratings not significantly different from real images under Welch’s test, indicating the synthetic cohort preserves clinically plausible radiographic appearance under blinded assessment.

Turing Test 2: Real versus synthetic discrimination with confidence. Radiologists classified each image as real or synthetic and assigned a five-point confidence score. As shown in Table 4 (b), both raters achieved moderate discrimination, indicating that some synthetic cues remain detectable. However, a substantial fraction of synthetic images were labeled as real, supporting their

visual plausibility. Confidence was only weakly correlated with correctness, suggesting limited calibration in this setting.

5 Conclusion

Critical RFO detection represents a clinically important yet data-scarce regime, where positive samples are rare and diagnostic cues are subtle and highly localized, challenging both conventional supervised learning and end-to-end synthesis. We introduced SurgRFO, a two-stage synthesis framework that separates surgical background modeling from localized RFO evidence generation. By combining a background diffusion generator with a lightweight local sampler through controlled fusion, SurgRFO produces high-quality RFO-positive radiographs for augmentation. We also released the Intraoperative RFO Bench, the first public benchmark dedicated to critical RFOs with expert annotations. On this benchmark, SurgRFO-augmented training yielded consistent improvements across multiple detector families, including under cross-institutional evaluation. A blinded radiologist study further supported the clinical plausibility of synthetic images, with comparable quality ratings to real radiographs and only moderate discriminability. These results demonstrate that structured synthesis can expand exposure to rare critical events and enhance detection sensitivity, while emphasizing the need for continued clinician-centered and workflow-aware validation prior to deployment.

Acknowledgments. No funding was received for this work.

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

References

1. Bluethgen, C., Chambon, P., Delbrouck, J.B., Van Der Sluijs, R., Połacin, M., Zambrano Chaves, J.M., Abraham, T.M., Purohit, S., Langlotz, C.P., Chaudhari, A.S.: A vision–language foundation model for the generation of realistic chest x-ray images. *Nature Biomedical Engineering* **9**(4), 494–506 (2025)
2. Cima, R.R., Kollengode, A., Garnatz, J., Storsveen, A., Weisbrod, C., Deschamps, C.: Incidence and characteristics of potential and actual retained foreign object events in surgical patients. *Journal of the American College of Surgeons* **207**(1), 80–87 (2008)
3. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
4. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)
5. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
6. Hsu, W.C., Wang, Y., Wu, Y.F., Chen, R., Afyouni, S., Liu, J., Vin, S., Shi, V., Imami, M., Chotiyanonta, J.S., et al.: Mri-based ovarian lesion classification via a foundation segmentation model and multimodal analysis: a multicenter study. *Radiology* **316**(2), e243412 (2025)
7. Huijben, E.M., Pluim, J.P., van Eijnatten, M.A.: Denoising diffusion probabilistic models for addressing data limitations in chest x-ray classification. *Informatics in Medicine Unlocked* **50**, 101575 (2024)
8. Kawakubo, M., Waki, H., Shirasaka, T., Kojima, T., Mikayama, R., Hamasaki, H., Akamine, H., Kato, T., Baba, S., Ushiro, S., et al.: A deep learning model based on fusion images of chest radiography and x-ray sponge images supports human visual characteristics of retained surgical items detection. *International Journal of Computer Assisted Radiology and Surgery* **18**(8), 1459–1467 (2023)
9. Khosravi, B., Li, F., Dapamede, T., Rouzrokh, P., Gamble, C.U., Trivedi, H.M., Wyles, C.C., Sellergren, A.B., Purkayastha, S., Erickson, B.J., et al.: Synthetically enhanced: unveiling synthetic data’s potential in medical imaging research. *EBioMedicine* **104** (2024)
10. Kufel, J., Bargiel-Łączek, K., Koźlik, M., Czogalik, Ł., Dudek, P., Magiera, M., Bartnikowska, W., Lis, A., Paszkiewicz, I., Kocot, S., et al.: Chest x-ray foreign objects detection using artificial intelligence. *Journal of Clinical Medicine* **12**(18), 5841 (2023)
11. Kumar, G.S., Ramani, S., Mahajan, A., Jain, N., Sequeira, R., Thakur, M.: Imaging of retained surgical items: A pictorial review including new innovations. *Indian Journal of Radiology and Imaging* **27**(03), 354–361 (2017)
12. Mahaulpatha, P., Abeywardane, T., George, T.: Ddpm based x-ray image synthesizer. *arXiv preprint arXiv:2401.01539* (2024)
13. Moffatt-Bruce, S.D., Cook, C.H., Steinberg, S.M., Stawicki, S.P.: Risk factors for retained surgical items: a meta-analysis and proposed risk stratification system. *Journal of surgical research* **190**(2), 429–436 (2014)
14. O’Connor, A.R., Coakley, F.V., Meng, M.V., Eberhardt, S.: Imaging of retained surgical sponges in the abdomen and pelvis. *American journal of roentgenology* **180**(2), 481–489 (2003)
15. Prakash, E., Valanarasu, J.M.J., Chen, Z., Reis, E.P., Johnston, A., Pareek, A., Bluethgen, C., Gatidis, S., Olsen, C., Chaudhari, A.S., et al.: Evaluating and improving the effectiveness of synthetic chest x-rays for medical image analysis. In:

- Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4413–4421 (2025)
16. Rajagopal, A., Martin, J.: Gossypiboma—“a surgeon’s legacy” report of a case and review of the literature. *Diseases of the colon & rectum* **45**(1), 119–120 (2002)
 17. Regenbogen, S.E., Greenberg, C.C., Resch, S.C., Kollengode, A., Cima, R.R., Zinner, M.J., Gawande, A.A.: Prevention of retained surgical sponges: a decision-analytic model predicting relative cost-effectiveness. *Surgery* **145**(5), 527–535 (2009)
 18. Rigamonti, D., Rigamonti, K.H., Rigamonti, A.S.: Retained foreign object signals a dangerous atmosphere in the operating room. *Cureus* **17**(3) (2025)
 19. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
 20. Ruiz, N., Li, Y., Jampani, V., Pritch, Y., Rubinstein, M., Aberman, K.: Dream-booth: Fine tuning text-to-image diffusion models for subject-driven generation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 22500–22510 (2023)
 21. Santosh, K., Roy, S., Allu, S.: Generic foreign object detection in chest x-rays. In: International Conference on Recent Trends in Image Processing and Pattern Recognition. pp. 93–104. Springer (2021)
 22. Stawicki, S.P., Evans, D., Cipolla, J., Seamon, M., Lukaszczyk, J., Prosciak, M., Torigian, D., Doraiswamy, V., Yazzie, N., Gunter Jr, O., et al.: Retained surgical foreign bodies: a comprehensive review of risks and preventive strategies. *Scandinavian Journal of Surgery* **98**(1), 8–17 (2009)
 23. Sun, H.S., Chen, S.L., Kuo, C.C., Wang, S.C., Kao, Y.L.: Gossypiboma—retained surgical sponge. *Journal of the Chinese Medical Association* **70**(11), 511–513 (2007)
 24. Viola, P., Jones, M.: Rapid object detection using a boosted cascade of simple features. In: Proceedings of the 2001 IEEE computer society conference on computer vision and pattern recognition. CVPR 2001. vol. 1, pp. I–I. Ieee (2001)
 25. Wang, Y., Shi, V.R., Zhou, L., Chin, R., Dai, Y., Hu, Y., Li, C.Y., Guan, H., Cheng, J., Sun, Y., et al.: Dataset and benchmark for enhancing critical retained foreign object detection. arXiv preprint arXiv:2507.06937 (2025)
 26. Wang, Z., Simoncelli, E.P., Bovik, A.C.: Multiscale structural similarity for image quality assessment. In: The thirty-seventh asilomar conference on signals, systems & computers, 2003. vol. 2, pp. 1398–1402. Ieee (2003)
 27. Weber, T., Ingrisch, M., Bischl, B., Rügamer, D.: Cascaded latent diffusion models for high-resolution chest x-ray synthesis. In: Pacific-Asia conference on knowledge discovery and data mining. pp. 180–191. Springer (2023)
 28. Wu, J., Wang, Y., Zhong, Z., Liao, W., Trayanova, N., Jiao, Z., Bai, H.X.: Vision-language foundation model for 3d medical imaging. *npj Artificial Intelligence* **1**(1), 17 (2025)
 29. Xie, C., Yoshii, Y., Kitahara, I.: Sv-drr: High-fidelity novel view x-ray synthesis using diffusion model. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 572–582. Springer (2025)
 30. Xue, Z., Candemir, S., Antani, S., Long, L.R., Jaeger, S., Demner-Fushman, D., Thoma, G.R.: Foreign object detection in chest x-rays. In: 2015 IEEE international conference on bioinformatics and biomedicine (BIBM). pp. 956–961. IEEE (2015)
 31. Zhou, L., Wang, Y., Shi, V., Feng, J., Zhao, L.m., Atagu, N., Madhu, P., Mehta, T., Trivedi, P., Lin, C.T., et al.: A dataset and benchmark for enhancing retained foreign object detection through physics-based image synthesis. In: MICCAI Workshop on Data Engineering in Medical Imaging. pp. 147–156. Springer (2025)