

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 MODELING THE DENSITY OF PIXEL-LEVEL SELF-SUPERVISED EMBEDDINGS FOR UNSUPERVISED PATHOLOGY SEGMENTATION IN MEDICAL CT

Anonymous authors

Paper under double-blind review

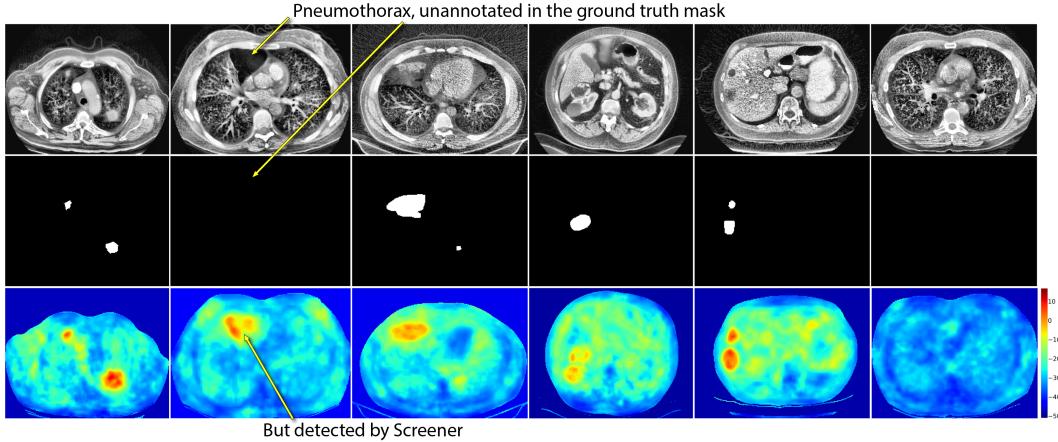


Figure 1: Examples of CT image slices (the first row), the ground truth pathology masks (the second row) and the anomaly maps predicted by our *fully self-supervised* Screener model (the third row).

ABSTRACT

Accurate detection of all pathological findings in 3D medical images remains a significant challenge, as supervised models are limited to detecting only the few pathology classes annotated in existing datasets. To address this, we frame pathology detection as an unsupervised visual anomaly segmentation (UVAS) problem, leveraging the inherent rarity of pathological patterns compared to healthy ones. We enhance the existing density-based UVAS framework with two key innovations: (1) dense self-supervised learning for feature extraction, eliminating the need for supervised pretraining, and (2) learned, masking-invariant dense features as conditioning variables, replacing hand-crafted positional encodings. Trained on over 30,000 unlabeled 3D CT volumes, our fully self-supervised model, Screener, outperforms existing UVAS methods on four large-scale test datasets comprising 1,820 scans with diverse pathologies. Furthermore, in a *low-shot* supervised fine-tuning setting, Screener surpasses existing self-supervised pretraining methods, establishing it as a state-of-the-art foundation for pathology segmentation. The code and pretrained models will be made publicly available.

1 INTRODUCTION

Accurate identification, localization, and classification of *all* pathological findings in 3D medical images remain a significant challenge in medical computer vision. While supervised models have shown promise, their utility is limited by the scarcity of labeled datasets, which often contain annotations for only a few pathologies. For example, Figure 1 shows 2D slices of 3D computed tomography (CT) images (first row) from public datasets providing annotations of lung cancer, pneumonia, kidney tumors, or liver tumors, while annotations of other pathologies, e.g., pneumothorax, are missing. This restricts supervised models to narrow, task-specific applications.

054 However, unlabeled CT datasets are abundant but often remain unused. Leveraging these datasets,
 055 we aim to develop an unsupervised model capable of distinguishing pathological regions from nor-
 056 mal ones. Our core assumption is that pathological patterns are statistically rarer than healthy pat-
 057 terns in CT images. This frames pathology segmentation as an unsupervised visual anomaly seg-
 058 mentation (UVAS) problem.

059 Although existing UVAS methods have been extensively explored for natural images, their adapta-
 060 tion to medical imaging is challenging. One obstacle is that uncurated CT datasets include many
 061 patients with pathologies, and there is no automatic way to filter them out to ensure a training set
 062 composed entirely of normal (healthy) images — a common requirement for synthetic-based (Zavr-
 063 tanik et al., 2021; Marimont & Tarroni, 2023) and reconstruction-based (Baur et al., 2021; Schlegl
 064 et al., 2019) UVAS methods. Density-based approaches (Gudovskiy et al., 2022; Zhou et al., 2024)
 065 are better suited, as they model image patterns probabilistically and assume abnormal patterns are
 066 rare rather than absent. To model the density of image patterns, existing methods encode them into
 067 feature maps using an ImageNet-pretrained encoder. Therefore, their performance on medical im-
 068 ages degrades due to a domain shift. Supervised medical encoders like STU-Net (Huang et al., 2023)
 069 might seem viable, but our experiments show they also underperform, likely because their features
 070 are too specific and lack discriminative information for pathology segmentation.

071 To address these challenges, we propose using dense self-supervised learning (SSL) (O. Pinheiro
 072 et al., 2020) to pretrain more discriminative feature maps of CT images and employ them in the
 073 density-based UVAS framework. Thus, our model learns the distribution of dense SSL embeddings
 074 and assigns high anomaly scores to image regions where embeddings fall into low-density regions.

075 Inspired by dense SSL, we also generalize the idea of conditioning in density-based UVAS methods.
 076 Existing works (Gudovskiy et al., 2022; Zhou et al., 2024) use hand-crafted conditioning variables
 077 such as pixel-wise sinusoidal positional embeddings. We replace them by *learned* pixel-wise context-
 078 ual embeddings capturing global characteristics of individual image regions, e.g. their anatomical
 079 position, patient’s age, etc. At the same time, we eliminate local information about presence of
 080 pathologies from the learned conditioning variables by enforcing their invariance to image masking.

081 We train our model, *Screener*, on 30,000 unlabeled CT volumes and evaluate it on 1,820 scans in
 082 two settings. First, as a fully unsupervised model, it achieves remarkable results (Figure 1), signifi-
 083 cantly outperforming existing UVAS methods. Second, after fine-tuning for downstream pathology
 084 segmentation tasks, *Screener* rivals other state-of-the-art pretrained models.

085 Our key contributions are four-fold:

- 087 • **Dense self-supervised features for density-based UVAS.** We demonstrate that dense self-
 088 supervised representations can be successfully used and even preferred over supervised fea-
 089 ture extractors in density-based UVAS methods. This enables a novel fully self-supervised
 090 UVAS framework for domains with limited labeled data.
- 091 • **Learned conditioning variables.** We propose novel self-supervised conditioning variables
 092 for density-based UVAS, simplifying the conditional distributions and enabling a simple
 093 Gaussian density model to perform on par with normalizing flows.
- 094 • **State-of-the-art UVAS results in CT.** This work presents the first large-scale evaluation
 095 of UVAS methods for CT images, showing state-of-the-art performance on unsupervised
 096 semantic segmentation of pathologies in diverse anatomical regions, including lung cancer,
 097 pneumonia, liver and kidney tumors.
- 098 • **State-of-the-art pretraining for pathology segmentation.** We introduce a novel pretrain-
 099 ing method that distills *Screener* into a UNet, enabling supervised fine-tuning and matching
 100 the performance of state-of-the-art self-supervised pretraining methods.

102 2 BACKGROUND & NOTATION

104 2.1 DENSITY-BASED UVAS

105 The core idea of density-based UVAS methods is to assign high anomaly scores to image regions
 106 containing *statistically rare* patterns. To implement this idea, they involve two models, which we

108 call a *descriptor model* and a *density model*. The descriptor model encodes image patterns into
 109 vector representations, while the density model learns their distribution and assigns anomaly scores.
 110

111 The descriptor model $f_{\theta^{\text{desc}}}$ is usually a pretrained fully-convolutional neural network. For a 3D
 112 image $\mathbf{x} \in \mathbb{R}^{H \times W \times S}$, it produces feature maps $\mathbf{y} \in \mathbb{R}^{h \times w \times s \times d^{\text{desc}}}$ consisting of vectors $\mathbf{y}[p] \in$
 113 $\mathbb{R}^{d^{\text{desc}}}$, which we call *descriptors* of positions $p \in P = \{1, \dots, h\} \times \{1, \dots, w\} \times \{1, \dots, s\}$.

114 The density model $q_{\theta^{\text{dens}}}(y)$ estimates the descriptors' marginal density $q_Y(y)$ (here, Y denotes the
 115 descriptor of a random position in a random image). For an abnormal pattern at position p , the
 116 descriptor $\mathbf{y}[p]$ is expected to lie in a low-density region, resulting in a low $q_{\theta^{\text{dens}}}(\mathbf{y}[p])$. Conversely,
 117 normal patterns correspond to high density values. During inference, the negative log-density values,
 118 $-\log q_{\theta^{\text{dens}}}(\mathbf{y}[p])$, are used as anomaly segmentation scores.

119 This framework can be extended with a conditioning mechanism. For each position p , an auxiliary
 120 variable $\mathbf{c}[p] \in \mathbb{R}^{d^{\text{cond}}}$, called a *condition*, is introduced. Instead of modeling the marginal density
 121 $q_Y(y)$, the conditional density $q_{Y|C}(y | c)$ is learned for each condition c , where (Y, C) represents
 122 the descriptor and condition at a random position in a random image. At inference, the negative
 123 log-conditional densities, $-\log q_{\theta^{\text{dens}}}(\mathbf{y}[p] | \mathbf{c}[p])$, serve as anomaly scores. State-of-the-art meth-
 124 ods (Gudovskiy et al., 2022; Zhou et al., 2024) follow this conditional framework using sinusoidal
 125 positional encodings as conditions.

126

127 2.2 DENSE JOINT EMBEDDING SSL

128

129 Joint embedding SSL models learn meaningful image embeddings by generating positive pairs —
 130 augmented views of the same image (e.g., random crops). They optimize embeddings to capture
 131 mutual information between views, making them both discriminative (distinguishing images) and
 132 augmentation-invariant (predictable across views). Contrastive methods, e.g., SimCLR (Chen et al.,
 133 2020), explicitly push apart embeddings of different images, while non-contrastive methods, e.g.,
 134 VICReg (Bardes et al., 2021), avoid embeddings' collapse through regularization. Details on Sim-
 135 CLR and VICReg are in the Appendix B.

136

137 *Dense* SSL methods extend this idea to learn image feature maps consisting of pixel-wise embed-
 138 dings that encode information about different spatial positions in the image. To this end, they define
 139 positive pairs at the pixel level: two embeddings are positive if they correspond to the same absolute
 140 position in the original image, but are predicted from different augmented crops (see the upper part
 141 of Figure 2 for illustration). Thus, dense SSL enforces feature maps to be equivariant w.r.t. crops,
 142 while encouraging dissimilarity between embeddings from different positions. DenseCL (Wang
 143 et al., 2021) and VADER (O. Pinheiro et al., 2020) use contrastive losses, while VICRegL (Bardes
 et al., 2022) adopts a VICReg objective.

144

145

3 METHOD

146

147

Novelty statement. Our method, illustrated in Figure 2, enhances the density-based UVAS frame-
 148 work with two key innovations. First, instead of relying on generic backbones, we *pretrain our*
 149 *descriptor model via dense SSL* which enables domain-specific, high-resolution, customizable and
 150 more discriminative descriptors (Section 3.1). Second, we introduce novel *masking-invariant con-
 151 ditioning variables, also learned via dense SSL* (Section 3.2), and largely simplifying further con-
 152 ditional density modeling (Section 3.3). Beyond these contributions, we distill the overall UVAS
 153 inference pipeline to a single UNet architecture, which makes it suitable for further supervised fine-
 154 tuning. This allows us to reinterpret our framework as a *novel self-supervised pretraining method*.

155

156

157

158

159

160

161

3.1 DESCRIPTOR MODEL

The success of our method relies on high-quality descriptors that are discriminative of pathology yet
 robust to irrelevant normal variations. Dense SSL provides a principled way to achieve this balance:
 voxel-level objectives encourage spatial discrimination, while augmentation invariance eliminates
 low-level details, leading to a smoother, more semantically structured embedding space in which
 similar normal patterns map to high-density areas.

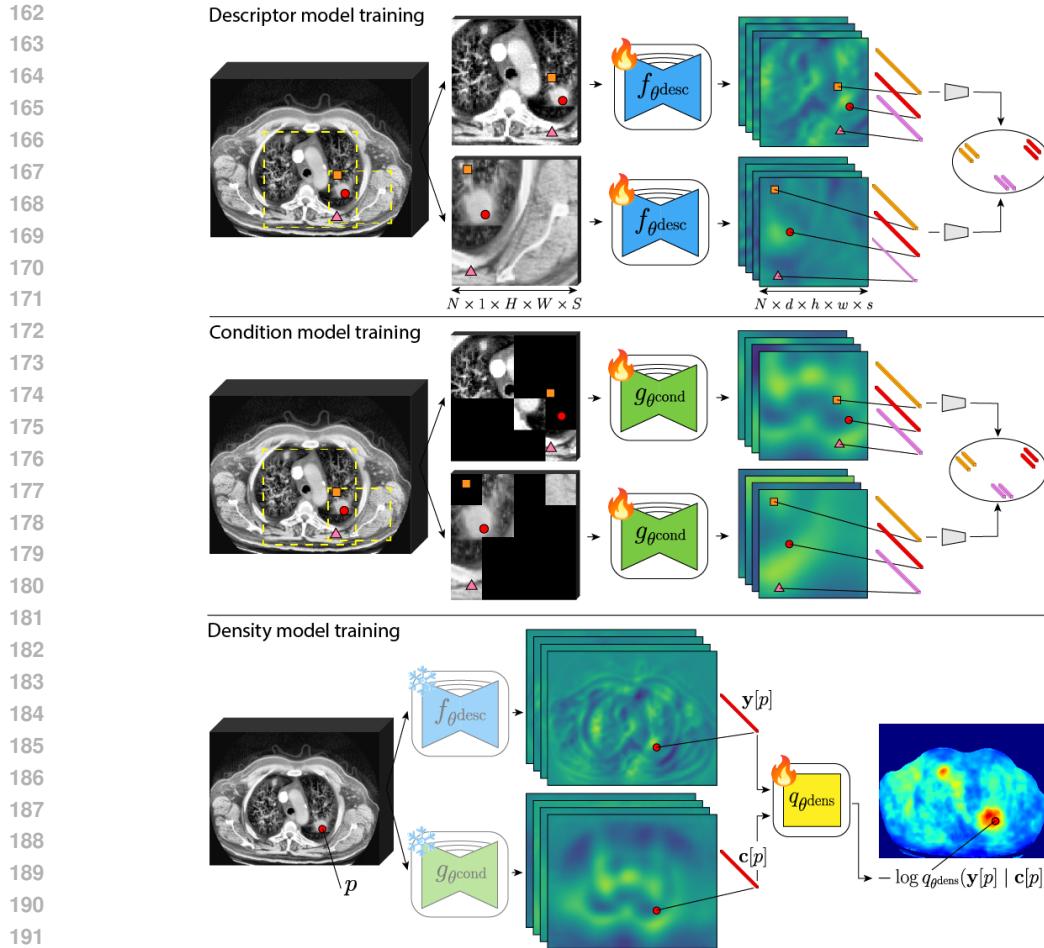


Figure 2: Illustration of Screener. First, we pretrain a *descriptor model* to produce discriminative feature maps, equivariant w.r.t. image crops and rescaling and invariant w.r.t. color jitter. Second, we train a *condition model* in the same way as the descriptor model, but also enforcing invariance to image masking. Third, a density model learns the conditional distribution $p_{Y|C}(y|c)$ of feature vectors $Y = y[p]$ and $C = c[p]$ extracted by the descriptor and condition models from random image at random position p . To obtain anomaly maps we apply the density model in a pixel-wise manner, which can be efficiently implemented using $1 \times 1 \times 1$ convolutions.

Our descriptor model design follows domain-driven, minimalistic principles, differing from the prior dense SSL literature (Wang et al., 2021; Bardes et al., 2022). We adopt a UNet-like architecture, which has proven a strong dense feature extractor in 3D medical imaging. Full resolution output enables precise localization of small pathologies. Each training batch includes embeddings from nearby voxels, forcing distinction of even spatially adjacent locations. We omit auxiliary global objectives or multi-scale feature pyramids — our approach is simple and principled, relying solely on dense self-supervision at full resolution.

The training process is illustrated in the upper part of Figure 2. From a random CT volume \mathbf{x} , we extract two overlapping, randomly sized 3D crops, resize them to $H \times W \times S$, and apply augmentations such as color jitter. The augmented crops, denoted $\mathbf{x}^{(1)}$ and $\mathbf{x}^{(2)}$, are passed through the descriptor model to produce feature maps $\mathbf{y}^{(1)}$ and $\mathbf{y}^{(2)}$. From the overlapping region of the two crops, we randomly select n positions. For each position p , we compute its coordinates $p^{(1)}$ and $p^{(2)}$ relative to the augmented views, and extract descriptors $y^{(1)} = \mathbf{y}^{(1)}[p^{(1)}]$ and $y^{(2)} = \mathbf{y}^{(2)}[p^{(2)}]$. These descriptors form a positive pair, as they correspond to the same position in the original volume but are predicted from different augmentations. Repeating this process for m different seed

216 CT volumes yields a batch of $N = n \cdot m$ positive pairs, denoted $\{(y_i^{(1)}, y_i^{(2)})\}_{i=1}^N$. These em-
 217 beddings are then optimized using standard dense SSL objectives, such as InfoNCE (Chen et al.,
 218 2020) or VICReg (Bardes et al., 2021), described in Appendix B. We refer to the resulting models
 219 as *DenseInfoNCE* and *DenseVICReg*, respectively.
 220

221 3.2 CONDITION MODEL

223 In medical imaging, the statistical plausibility of a local pattern often depends on its broader context,
 224 such as anatomical location or patient characteristics. This motivates modeling the conditional dis-
 225 tribution of descriptors, given relevant contextual variables. Conditioning offers two key advantages:
 226 it simplifies density estimation, as conditional distributions are usually less complex than marginal,
 227 and it may lead to more semantically meaningful anomalies, defined as deviations from what is ex-
 228 pected in a specific context. For example, a pattern normal in one anatomical region or patient group
 229 (e.g., a calcification in an elderly lung) might be abnormal in another (e.g., a calcification in breast).
 230

231 Conditioning variables can be global (e.g., patient metadata like age or sex) or voxel-wise, en-
 232 abling region-specific conditioning. General-domain UVAS methods (Gudovskiy et al., 2022; Zhou
 233 et al., 2024) utilize sinusoidal positional encodings of absolute spatial coordinates relative to the
 234 image origin. However, since medical scans may not be anatomically aligned, vanilla positional
 235 encodings lack consistent anatomical or patient-specific relevance. Anatomical Positional Embed-
 236 dings (APE) (Goncharov et al., 2024) offer an alternative by encoding pixels’ anatomical locations
 237 (though previously used for retrieval, not UVAS conditioning). However, it is domain-specific and
 238 may not capture all patient-level or fine-grained contextual nuances.
 239

240 To address the aforementioned limitations, we propose a domain-agnostic self-supervised method
 241 for *learning* conditions. Our key idea is to train a *condition model* $g_{\theta^{\text{cond}}}$ to predict voxel-wise
 242 embeddings that are consistent across different masked image views. For instance, as illustrated in
 243 Figure 2, the model learns to predict the same condition embedding for a location even if a pathology
 244 is visible in one masked view but not another. Consequently, the learned condition feature maps
 245 become *invariant to the presence / absence of anomalies*. At the same time, we encourage intra-
 246 subject, i.e. spatial, and inter-subject discriminativeness and expect feature maps to capture voxel-
 247 level features such as anatomical location and tissue type, and patient-level characteristics such as
 248 age or sex, which are robustly inferable from the global image structure. **We empirically probe**
 249 **pathology-related ignorance and anatomical awareness of our condition model in Appendix D.** The
 250 architecture and training procedure for the condition model $g_{\theta^{\text{cond}}}$ are exactly the same as those for
 251 the descriptor model, with the sole difference: random masking as an additional augmentation.
 252

253 3.3 DENSITY MODEL

254 The conditional density model $q_{\theta^{\text{dens}}}(y \mid c)$ can be viewed as a predictive model, which tries to
 255 predict descriptors based on the corresponding conditions. In this interpretation, anomaly scores
 256 $\{-\log q_{\theta^{\text{dens}}}(y[p] \mid c[p])\}_{p \in P}$ are position-wise prediction errors. Also note, that marginal density
 257 model $q_{\theta^{\text{dens}}}(y)$ is a special case of conditional model with a constant condition $c[p] = \text{const.}$
 258

259 During training, we sample a batch of m random crops, $\{x_i\}_{i=1}^m$, each of size $H \times W \times S$, from
 260 different CT images. For each crop, the pretrained descriptor and condition models produce the de-
 261 scriptor maps, $\{y_i\}_{i=1}^m$, and condition maps, $\{c_i\}_{i=1}^m$, and negative log-likelihood loss is optimized:
 262

$$\min_{\theta^{\text{dens}}} \frac{1}{m \cdot |P|} \sum_{i=1}^m \sum_{p \in P} -\log q_{\theta^{\text{dens}}}(y_i[p] \mid c_i[p]).$$

263 At inference, we divide an input CT image into M overlapping patches, $\{x_i\}_{i=1}^M$, each of size
 264 $H \times W \times S$. For each patch, we apply the descriptor, condition, and density models to compute the
 265 anomaly map, $\{-\log q_{\theta^{\text{dens}}}(y_i[p] \mid c_i[p])\}_{p \in P}$. These patch-wise anomaly maps are then aggregated
 266 into a single anomaly map aligned with the entire input volume. During aggregation, we average the
 267 predictions in patches’ overlapping regions.
 268

269 We explore two parameterizations for the density model $q_{\theta^{\text{dens}}}(y \mid c)$: Gaussian, as a straightforward
 270 baseline, and normalizing flows, similar to Gudovskiy et al. (2022); Zhou et al. (2024), as an expres-
 271 sive generative model enabling tractable density estimation. These parameterizations and the details
 272 of their implementation in the context of UVAS framework are further described in Appendix E.
 273

270 3.4 DISTILLATION AND SUPERVISED FINE-TUNING
271272 Although unsupervised Screener shows impressive results, supervised fine-tuning is the most practical
273 way to further improve its performance. The density-based UVAS pipeline, consisting of three
274 separate models, is not amenable to end-to-end optimization. To enable fine-tuning, we distill the
275 knowledge from the pretrained Screener into a single UNet architecture. This step can be viewed as
276 a novel self-supervised pretraining method for pathology segmentation tasks.277 During distillation, we sample random image crops, pass them through the pretrained modular
278 Screener to obtain ground truth anomaly score maps (negative log-density values). We then train
279 a regression UNet model (last conv has one output channel without activation) to predict these score
280 maps directly from the input image crops using a simple MSE loss. For supervised fine-tuning on
281 binary segmentation tasks, we randomly reinitialize the UNet’s last conv layer and append a sig-
282 moid activation. Then we fine-tune the model on task-specific labeled data using a combination of
283 voxel-wise binary cross-entropy and Dice losses.284 4 EXPERIMENTS
285

286 Our experiments can be divided into three main parts:

287

- **Unsupervised setting.** We show that our *unsupervised* Screener significantly outperforms
288 other UVAS methods on real-word medical CT datasets (Section 4.1).
- **Fine-tuning setting.** We demonstrate that Screener can serve as a state-of-the-art self-
289 supervised *pretraining* method. To this end, we fine-tune the distilled Screener (as de-
290 scribed in Section 3.4) for pathology segmentation tasks and compare it with supervised
291 model trained from scratch, as well as other fine-tuned pretrained models (Section 4.2).
- **Ablation study.** We explore how different choices of descriptor, condition and density
292 models in our method affect the UVAS results (Section 4.3).

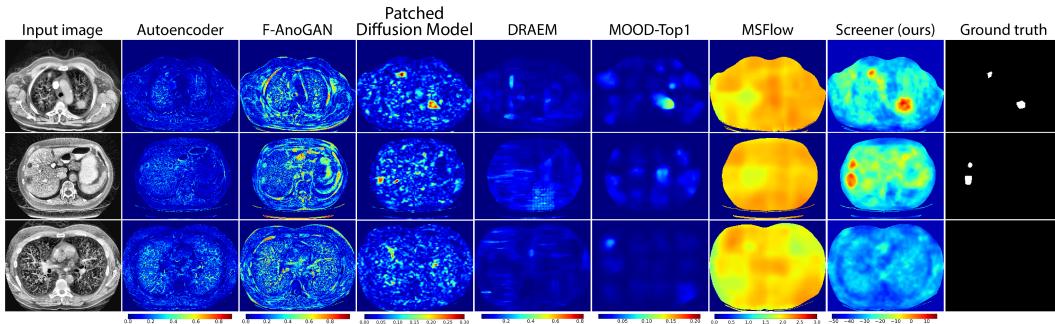
293 **Datasets.** We train Screener and other unsupervised models on three CT datasets: NLST (Team,
294 2011), AMOS (Ji et al., 2022), and AbdomenAtlas (Qu et al., 2024). These large-scale datasets
295 include diverse patients with potential pathologies, but their annotations are not available for data
296 filtering or training. For evaluation we use four datasets: LIDC (Armato III et al., 2011), MIDRC-
297 RICORD-1a (Tsai et al., 2020), KiTS (Heller et al., 2019) and LiTS (Bilic et al., 2023). These
298 datasets provide annotation masks only for certain pathologies. Any other pathologies present in
299 these datasets are not labeled. Summary table about the datasets is provided in Appendix F.

300 4.1 UNSUPERVISED SETTING

301 **Evaluation protocol.** We compare Screener with baseline UVAS models using voxel-level AU-
302 ROC and Dice score. [The segmentation threshold is selected to maximize Dice score on a subset](#)
303 [of 10 pathological cases, and then kept fixed for the entire test set.](#) Note that Dice scores are sig-
304 nificantly underestimated due to incomplete ground truth masks: while UVAS models aim to detect
305 *all* anomalies, the datasets provide annotations only for specific target pathologies. Detections cor-
306 responding to other unlabeled pathologies (see Figure 1 and Appendix A) are therefore mistakenly
307 counted as false positives. Voxel-level AUROC is a standard UVAS metric because its estimation
308 is more robust to the ground truth incompleteness issue. We estimate AUROC across all dataset
309 voxels by sampling 1000 pathological voxels (contributing to true positive rate) and 1000 out-of-
310 mask “normal” voxels (for false positive rate) per test image. The sampled “normal” voxels are
311 overwhelmingly normal, ensuring accurate AUROC estimation despite incomplete annotations.312 **Results.** Quantitative results are presented in Table 1. Qualitative results are shown in Figure 3.
313 Screener significantly outperforms the UVAS baselines. Autoencoder, f-AnoGAN and Patched
314 Diffusion Model tend to overfit to pathologies in the training data, and fail to reconstruct fine-
315 grained normal details (see also Appendix I). Synthetics-based DRAEM and MOOD-Top1 struggle
316 to generalize to the appearance of real medical pathologies. The density-based MSFlow, relying on
317 ImageNet-pretrained features, proves ineffective at discriminating pathologies from normal regions
318 in CT images.

324 Table 1: Comparison of Screener and the existing UVAS methods in unsupervised setting.
325

326 Model	327 Voxel-level AUROC						328 Dice score ¹		
	LIDC	MIDRC	KiTS	LiTS	LIDC	MIDRC	KiTS	LiTS	
329 Autoencoder (Baur et al., 2021)	0.71	0.65	0.66	0.68	0.00 ± 0.00	0.09 ± 0.07	0.01 ± 0.02	0.01 ± 0.01	
330 f-AnoGAN (Schlegl et al., 2019)	0.82	0.66	0.67	0.67	0.00 ± 0.00	0.09 ± 0.07	0.01 ± 0.02	0.01 ± 0.01	
331 Patched Diffusion Model (Behrendt et al., 2024)	0.87	0.76	0.76	0.80	0.01 ± 0.03	0.14 ± 0.08	0.02 ± 0.03	0.02 ± 0.04	
DRAEM (Zavrtanik et al., 2021)	0.63	0.72	0.82	0.83	0.00 ± 0.00	0.11 ± 0.08	0.03 ± 0.06	0.02 ± 0.04	
MOOD-Top1 (Marimont & Tarroni, 2023)	0.79	0.79	0.77	0.80	0.00 ± 0.01	0.13 ± 0.10	0.02 ± 0.07	0.06 ± 0.12	
MSFlow (Zhou et al., 2024)	0.71	0.67	0.63	0.63	0.00 ± 0.01	0.08 ± 0.06	0.01 ± 0.01	0.00 ± 0.01	
Screener (ours)	0.96	0.87	0.90	0.93	0.05 ± 0.13	0.30 ± 0.18	0.06 ± 0.09	0.10 ± 0.12	

345 Figure 3: Qualitative comparison of anomaly maps produced by baseline UVAS methods and un-
346 supervised Screener. First column contains CT slices, columns 2 to 7 are the baseline methods’
347 predictions, column 8 is the Screener’s prediction. Last column depicts the ground truth mask.
348

351 4.2 FINE-TUNING SETTING

353 **Evaluation protocol.** We fine-tune and test pretrained models on the evaluation datasets via 3-fold
354 cross-validation. For each training fold, we use only 25 labeled cases, to amplify pretraining benefits
355 and to conserve computational resources. [Full-data fine-tuning results can be found in Appendix J](#).
356 We assess the models using Dice score. We use a Wilcoxon signed-rank test to compare all the
357 fine-tuned models with the nnUNet (Isensee et al., 2021) trained from scratch.
358359 **Results.** Results in Table 2 demonstrate that Screener-based pretraining consistently improves
360 downstream segmentation performance across all test datasets in low-data fine-tuning setup, with
361 significant gains on LIDC (a 1.5-fold Dice increase) and LiTS. Screener is competitive with super-
362 vised pretraining (Huang et al., 2023) and state-of-the-art self-supervised VoCo (Wu et al., 2024),
363 and outperform other SSL models (Zhou et al., 2021; Tang et al., 2022; Valanarasu et al., 2023).
364 However, when fine-tuned on full datasets, Screener does not yield any gains compared to training
365 from scratch (see Appendix J), which is consistent with the broader observation that many SSL
366 methods provide their strongest advantages when training data are scarce.
367

368 4.3 ABLATION STUDY

369 Table 3 presents the ablation study of our proposed condition model. We compare our condition
370 model with two baselines: vanilla sinusoidal positional encodings and APE (Goncharov et al., 2024),
371 detailed in Appendix C. We evaluate condition models in combination with the fixed DenseVICReg
372 descriptor model and two different density models — Gaussian and normalizing flow — described
373 in Appendix E. When we use expressive normalizing flow density model, all conditioning strategies
374 yield results comparable to each other and to the unconditional model. However, in experiments
375 with simple Gaussian density models, we see that the results significantly improve as the condition-
376377 ¹Note that Dice scores are often underestimated in the unsupervised setting, as ground truth masks cover
378 only certain target pathologies, while UVAS models intentionally detect *all* pathologies. Many true positives
379 are thus mistakenly counted as false positives (see Figure 1 and Appendix A for examples).

378 Table 2: Dice scores of Screener and other self-supervised pretrained models after fine-tuning.
379 We highlight statistically significant improvements (green) or declines (red) relative to nnUNet
380 trained from random initialization. Note that **Grounding DINO** (Liu et al., 2023) and **LVM-
381 Med** (MH Nguyen et al., 2023) operate on 2D slices without access to full 3D image context, and
382 therefore show inferior performance to other baselines which use 3D convolutions or 3D attention.

Model	LIDC	MIDRC	KiTS	LiTS
nnUNet (random init.) (Isensee et al., 2021)	0.21	0.61	0.41	0.45
nnUNet (supervised pretrain.) (Huang et al., 2023)	0.29 \uparrow 40% ($p < 0.01$)	0.62 \uparrow 2% ($p = 0.51$)	0.46 \uparrow 10% ($p < 0.01$)	0.48 \uparrow 7% ($p < 0.01$)
Grounding DINO (Liu et al., 2023)	0.09 \uparrow 54% ($p < 0.01$)	0.58 \downarrow 5% ($p < 0.01$)	0.19 \downarrow 55% ($p < 0.01$)	0.33 \downarrow 26% ($p < 0.01$)
LVM-Med (MH Nguyen et al., 2023)	0.04 \uparrow 80% ($p < 0.01$)	0.55 \downarrow 8% ($p < 0.01$)	0.16 \downarrow 62% ($p < 0.01$)	0.33 \downarrow 26% ($p < 0.01$)
Model Genesis (Zhou et al., 2021)	0.21 \uparrow 1% ($p = 0.76$)	0.59 \downarrow 2% ($p = 0.05$)	0.34 \downarrow 18% ($p < 0.01$)	0.39 \downarrow 12% ($p = 0.01$)
SwinUNETR (Tang et al., 2022)	0.16 \downarrow 24% ($p < 0.01$)	0.55 \downarrow 9% ($p < 0.01$)	0.19 \downarrow 53% ($p < 0.01$)	0.39 \downarrow 13% ($p < 0.01$)
DAE (Valanarasu et al., 2023)	0.15 \downarrow 26% ($p < 0.01$)	0.58 \downarrow 4% ($p < 0.01$)	0.26 \downarrow 38% ($p < 0.01$)	0.36 \downarrow 20% ($p < 0.01$)
VoCo (Wu et al., 2024)	0.20 \downarrow 2% ($p = 0.79$)	0.61 \uparrow 1% ($p = 0.89$)	0.49 \uparrow 17% ($p < 0.01$)	0.49 \uparrow 10% ($p < 0.01$)
MAE (Wald et al., 2025)	0.24 \uparrow 18% ($p < 0.01$)	0.62 \uparrow 3% ($p = 0.01$)	0.35 \downarrow 17% ($p = 0.04$)	0.49 \uparrow 10% ($p < 0.01$)
DenseVICReg	0.22 \uparrow 7% ($p = 0.15$)	0.58 \downarrow 4% ($p < 0.01$)	0.31 \downarrow 26% ($p < 0.01$)	0.44 \downarrow 2% ($p = 0.92$)
Screener (ours)	0.31 \uparrow 49% ($p < 0.01$)	0.62 \uparrow 3% ($p = 0.45$)	0.43 \uparrow 4% ($p = 0.17$)	0.48 \uparrow 7% ($p < 0.01$)

394 ing variables becomes more informative. Remarkably, our proposed masking-invariant condition
395 model allows Gaussian model to achieve very strong anomaly segmentation results competing with
396 complex flow-based models.

398 Table 3: Ablation study of the effect of conditional model for gaussian and flow-based density
399 models. None in Condition model column means that results are given for a marginal model.

Descriptor model	Condition model	Density model	Voxel-level AUROC				Dice score			
			LIDC	MIDRC	KiTS	LiTS	LIDC	MIDRC	KiTS	LiTS
DenseVICReg, $d^{\text{desc}} = 32$	None	Gaussian	0.81	0.81	0.61	0.71	0.00 ± 0.00	0.17 ± 0.13	0.00 ± 0.01	0.00 ± 0.01
— “ —	Sin-cos pos.	— “ —	0.82	0.80	0.74	0.77	0.00 ± 0.00	0.14 ± 0.11	0.01 ± 0.02	0.01 ± 0.02
— “ —	APE	— “ —	0.88	0.80	0.78	0.86	0.00 ± 0.03	0.14 ± 0.10	0.01 ± 0.01	0.01 ± 0.03
— “ —	Masking-invariant	— “ —	0.96	0.84	0.87	0.90	0.04 ± 0.08	0.21 \pm 0.13	0.03 \pm 0.05	0.13 \pm 0.19
— “ —	None	Norm. flow	0.96	0.89	0.88	0.93	0.05 ± 0.12	0.31 \pm 0.18	0.04 ± 0.06	0.09 ± 0.12
— “ —	Sin-cos pos.	— “ —	0.96	0.89	0.90	0.94	0.05 ± 0.13	0.30 ± 0.18	0.06 ± 0.09	0.10 \pm 0.12
— “ —	APE	— “ —	0.96	0.88	0.94	0.94	0.04 ± 0.11	0.28 ± 0.18	0.05 ± 0.08	0.09 ± 0.13
— “ —	Masking-invariant	— “ —	0.96	0.87	0.90	0.93	0.05 ± 0.13	0.28 ± 0.18	0.07 ± 0.11	0.10 ± 0.13

408 We also ablate different choices of descriptor model in Table 4. We compare DenseInfoNCE and
409 DenseVICReg and conclude that dense VICReg objective works slightly better. We also compare
410 two DenseVICReg models with different descriptors’ dimensionality $d^{\text{desc}} = 32$ or $d^{\text{desc}} = 128$ and
411 conclude that increasing dimensionality does not improve the results. To demonstrate the superior-
412 ity of our domain-specific self-supervised descriptor model over supervised feature extractors, we
413 compare them it with ImageNet-pretrained ResNet50 (Zhou et al., 2024) and STU-Net (Huang et al.,
414 2023) — a UNet pretrained in a supervised manner on anatomical structure segmentation tasks.

416 Table 4: Ablation study of the effect of descriptor model. In these experiments we do not use con-
417 ditioning and use normalizing flow as a marginal density model. We include MSFlow (Zhou et al.,
418 2024) to demonstrate that ImageNet-pretrained descriptor model is inappropriate for 3D medical CT
419 images.

Descriptor model	Condition model	Density model	Voxel-level AUROC				Dice score			
			LIDC	MIDRC	KiTS	LiTS	LIDC	MIDRC	KiTS	LiTS
ImageNet	Sin-cos pos.	MSFlow	0.70	0.66	0.64	0.64	0.00 ± 0.01	0.08 ± 0.06	0.01 ± 0.01	0.00 ± 0.01
STU-Net (Huang et al., 2023)	None	Norm. flow	0.52	0.44	0.52	0.64	0.00 ± 0.00	0.02 ± 0.03	0.01 ± 0.02	0.01 ± 0.01
DenseInfoNCE, $d^{\text{desc}} = 32$	None	Norm. flow	0.96	0.87	0.87	0.91	0.04 ± 0.11	0.28 ± 0.18	0.04 \pm 0.06	0.05 ± 0.09
DenseVICReg, $d^{\text{desc}} = 32$	None	Norm. flow	0.96	0.89	0.88	0.93	0.05 ± 0.12	0.31 \pm 0.18	0.04 \pm 0.06	0.09 \pm 0.12
DenseVICReg, $d^{\text{desc}} = 128$	None	Norm. flow	0.96	0.90	0.87	0.93	0.04 ± 0.09	0.31 \pm 0.18	0.03 ± 0.06	0.08 ± 0.12

5 RELATED WORK

428 **Reconstruction-based UVAS.** Reconstruction-based methods train a generative model to recon-
429 struct the original image from its compressed representation (Baur et al., 2021; Schlegl et al., 2019)
430 or from its corrupted, e.g., noised (Behrendt et al., 2024), version. If training set is anomaly-free

432 these models struggle to reconstruct anomalies in the test set and absolute differences between the
 433 original and reconstructed pixel values can be used as anomaly maps. However, when training
 434 dataset contains real anomalies, reconstruction-based models can learn to reconstruct anomalies as
 435 well as normal regions, diminishing their ability to differentiate. Another limitation is that measuring
 436 reconstruction errors in raw pixel space can be problematic: some abnormal pixels can accidentally
 437 have small reconstruction errors, while some normal fine-grained details, which are inherently
 438 difficult to reconstruct precisely, might yield high reconstruction errors.

439
 440 **Synthetics-based UVAS.** These methods rely on generating synthetic image anomalies and training
 441 a supervised model to segment them. Anomalies can be simulated by corrupting random image
 442 regions with noise, replacing them with random patterns from a specialized set (Zavrtanik et al.,
 443 2021), or using parts of other training images (Marimont & Tarroni, 2023). While these models are
 444 straightforward to implement and train, they overfit to synthetic anomalies and struggle to generalize
 445 effectively to real-world anomalies.

446
 447 **Density-based UVAS.** We explain the idea of density-based UVAS in Section 2.1. Some methods
 448 (Roth et al., 2022) use non-parametric density models based on memory banks. More scalable
 449 flow-based methods (Yu et al., 2021; Gudovskiy et al., 2022; Zhou et al., 2024), leverage normalizing
 450 flows. In our experiments, we included MSFlow (Zhou et al., 2024), as it was among the top-5
 451 performing methods on MVTecAD (Bergmann et al., 2021) at the time.

452
 453 **Medical UVAS.** Recognized methods are either reconstruction-based (Baur et al., 2021; Schlegl
 454 et al., 2019; Pinaya et al., 2022; Behrendt et al., 2024) or synthetics-based (Marimont & Tarroni,
 455 2023). f-AnoGAN (Schlegl et al., 2019) trains generator g and discriminator d , to generate anomaly-
 456 free images $x \sim g(z)$ from latent variables z . Then, it trains encoder f to map anomaly-free
 457 images x to the latent space, s.t. they can be reconstructed via frozen generator $\hat{x} = g(f(x)) \approx x$.
 458 Patched Diffusion Model (Behrendt et al., 2024) cuts out image patches and trains a diffusion
 459 model to reconstruct them based on the surrounding context. At inference, an image is split into a
 460 grid of patches and Diffusion model reconstructs each patch from its noised version based on the
 461 remaining clean patches. MOOD-Top1 (Marimont & Tarroni, 2023) is a straightforward synthetics-
 462 based method showing top-1 performance on MOOD (Zimmerer et al., 2022).

463
 464 **Medical self-supervised pretraining.** Methods like Model Genesis (Zhou et al., 2021) and Swi-
 465 nUNETR (Tang et al., 2022) utilize combinations of contrastive learning, masked image modeling,
 466 and various pretext tasks re-implemented for 3D CT volumes. DAE (Valanarasu et al., 2023) pre-
 467 train a model to reconstruct original images from their disrupted versions created by local masking
 468 across channel embeddings and low-level perturbations like noise and downsampling. Volume
 469 Contrast (VoCo) (Wu et al., 2024) employs a contrastive approach to implicitly encode contextual
 470 position priors, treating different image regions as distinct "classes" and predicting which region a
 471 random sub-volume belongs to by contrasting its representation against base crops. To our knowl-
 472 edge, Screener is the first work to propose and demonstrate the effectiveness of using unsupervised
 473 anomaly segmentation as a pretraining strategy for downstream pathology segmentation tasks.

474 6 CONCLUSION

475
 476 Our work addresses the critical challenge of detecting all pathological findings in 3D CT images,
 477 a task hindered by limited labeled data. Assuming the inherent rarity of pathological patterns,
 478 we frame this as a UVAS problem. We propose Screener, a novel density-based UVAS frame-
 479 work with dense SSL, ensuring discriminative and robust domain-specific descriptors, and learned,
 480 masking-invariant conditioning variables that simplify density modeling. Evaluated on four large-
 481 scale datasets, the fully unsupervised Screener achieved state-of-the-art performance, effectively
 482 localizing diverse pathologies. Furthermore, when distilled and fine-tuned, Screener demonstrated
 483 strong performance on supervised segmentation tasks, establishing its value as a novel pretraining
 484 method. Screener represents a significant step towards comprehensive and scalable pathology de-
 485 tection, serving as a powerful unsupervised screening tool and a robust foundation for supervised
 486 applications.

486 **Limitations & future work.** Despite its promising performance, Screener has several limitations
 487 that warrant future investigation. Its reliance on the rarity assumption may lead to false negative
 488 errors for common or widespread pathologies, while statistical anomalies that lack clinical sig-
 489 nificance, e.g. artifacts, could result in false positives (though we analyze robustness to artifacts,
 490 low-dose and contrast agent in Appendix H). Comprehensive evaluation of UVAS methods remains
 491 challenging due to the lack of ground truth annotations for all potential pathologies. Currently val-
 492 idated on CT, the generalizability of our approach to other medical imaging modalities requires
 493 further exploration. Future work will also explore scaling laws to investigate how performance
 494 scales with model size and training data, potentially unlocking further improvements.

495 **REFERENCES**

496 Samuel G Armato III, Geoffrey McLennan, Luc Bidaut, Michael F McNitt-Gray, Charles R Meyer,
 497 Anthony P Reeves, Binsheng Zhao, Denise R Aberle, Claudia I Henschke, Eric A Hoffman,
 498 et al. The lung image database consortium (lidc) and image database resource initiative (idri): a
 499 completed reference database of lung nodules on ct scans. *Medical physics*, 38(2):915–931, 2011.
 500

501

502 Adrien Bardes, Jean Ponce, and Yann LeCun. Vicreg: Variance-invariance-covariance regularization
 503 for self-supervised learning. *arXiv preprint arXiv:2105.04906*, 2021.

504

505 Adrien Bardes, Jean Ponce, and Yann LeCun. Vicregl: Self-supervised learning of local visual
 506 features. *Advances in Neural Information Processing Systems*, 35:8799–8810, 2022.

507

508 Christoph Baur, Stefan Denner, Benedikt Wiestler, Nassir Navab, and Shadi Albarqouni. Autoen-
 509 coders for unsupervised anomaly segmentation in brain mr images: a comparative study. *Medical*
 510 *Image Analysis*, 69:101952, 2021.

511 Finn Behrendt, Debayan Bhattacharya, Julia Krüger, Roland Opfer, and Alexander Schlaefer.
 512 Patched diffusion models for unsupervised anomaly detection in brain mri. In *Medical Imag-*
 513 *ing with Deep Learning*, pp. 1019–1032. PMLR, 2024.

514

515 Paul Bergmann, Kilian Batzner, Michael Fauser, David Sattlegger, and Carsten Steger. The mvtec
 516 anomaly detection dataset: a comprehensive real-world dataset for unsupervised anomaly detec-
 517 tion. *International Journal of Computer Vision*, 129(4):1038–1059, 2021.

518

519 Patrick Bilic, Patrick Christ, Hongwei Bran Li, Eugene Vorontsov, Avi Ben-Cohen, Georgios
 520 Kaassis, Adi Szeskin, Colin Jacobs, Gabriel Efrain Humpire Mamani, Gabriel Chartrand, et al.
 521 The liver tumor segmentation benchmark (lits). *Medical Image Analysis*, 84:102680, 2023.

522

523 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for
 524 contrastive learning of visual representations. In *International conference on machine learning*,
 525 pp. 1597–1607. PMLR, 2020.

526

527 Mikhail Goncharov, Valentin Samokhin, Eugenia Soboleva, Roman Sokolov, Boris Shirokikh,
 528 Mikhail Belyaev, Anvar Kurmukov, and Ivan Oseledets. Anatomical positional embeddings.
 529 *arXiv preprint arXiv:2409.10291*, 2024.

530

531 Denis Gudovskiy, Shun Ishizaka, and Kazuki Kozuka. Cflow-ad: Real-time unsupervised anomaly
 532 detection with localization via conditional normalizing flows. In *Proceedings of the IEEE/CVF*
 533 *winter conference on applications of computer vision*, pp. 98–107, 2022.

534

535 Nicholas Heller, Niranjan Sathianathan, Arveen Kalapara, Edward Walczak, Keenan Moore,
 536 Heather Kaluzniak, Joel Rosenberg, Paul Blake, Zachary Rengel, Makinna Oestreich, et al. The
 537 kits19 challenge data: 300 kidney tumor cases with clinical context, ct semantic segmentations,
 538 and surgical outcomes. *arXiv preprint arXiv:1904.00445*, 2019.

539

540 Ziyan Huang, Haoyu Wang, Zhongying Deng, Jin Ye, Yanzhou Su, Hui Sun, Junjun He, Yun Gu,
 541 Lixu Gu, Shaoting Zhang, et al. Stu-net: Scalable and transferable medical image segmentation
 542 models empowered by large-scale supervised pre-training. *arXiv preprint arXiv:2304.06716*,
 543 2023.

540 Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnunet: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021.

541

542

543 Yuanfeng Ji, Haotian Bai, Chongjian Ge, Jie Yang, Ye Zhu, Ruimao Zhang, Zhen Li, Lingyan Zhanng, Wanling Ma, Xiang Wan, et al. Amos: A large-scale abdominal multi-organ benchmark for versatile medical image segmentation. *Advances in neural information processing systems*, 35:36722–36732, 2022.

544

545

546

547

548 Durk P Kingma and Prafulla Dhariwal. Glow: Generative flow with invertible 1x1 convolutions. *Advances in neural information processing systems*, 31, 2018.

549

550 Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *corr abs/2303.05499* (2023). *arXiv preprint arXiv:2303.05499*, 2023.

551

552

553

554 Sergio Naval Marimont and Giacomo Tarroni. Achieving state-of-the-art performance in the medical outof-distribution (mood) challenge using plausible synthetic anomalies. *arXiv preprint arXiv:2308.01412*, 2023.

555

556

557 Duy MH Nguyen, Hoang Nguyen, Nghiêm Diep, Tan Ngoc Pham, Tri Cao, Binh Nguyen, Paul Swoboda, Nhat Ho, Shadi Albarqouni, Pengtao Xie, et al. Lvm-med: Learning large-scale self-supervised vision models for medical imaging via second-order graph matching. *Advances in Neural Information Processing Systems*, 36:27922–27950, 2023.

558

559

560

561 Pedro O. Pinheiro, Amjad Almahairi, Ryan Benmalek, Florian Golemo, and Aaron C Courville. Unsupervised learning of dense visual representations. *Advances in Neural Information Processing Systems*, 33:4489–4500, 2020.

562

563

564

565 Walter HL Pinaya, Petru-Daniel Tudosi, Robert Gray, Geraint Rees, Parashkev Nachev, Sébastien Ourselin, and M Jorge Cardoso. Unsupervised brain imaging 3d anomaly detection and segmentation with transformers. *Medical Image Analysis*, 79:102475, 2022.

566

567

568 Stephen M Pizer, E Philip Amburn, John D Austin, Robert Cromartie, Ari Geselowitz, Trey Greer, Bart ter Haar Romeny, John B Zimmerman, and Karel Zuiderveld. Adaptive histogram equalization and its variations. *Computer vision, graphics, and image processing*, 39(3):355–368, 1987.

569

570

571 Chongyu Qu, Tiezheng Zhang, Hualin Qiao, Yucheng Tang, Alan L Yuille, Zongwei Zhou, et al. Abdomenatlas-8k: Annotating 8,000 ct volumes for multi-organ segmentation in three weeks. *Advances in Neural Information Processing Systems*, 36, 2024.

572

573

574

575 Karsten Roth, Latha Pemula, Joaquin Zepeda, Bernhard Schölkopf, Thomas Brox, and Peter Gehler. Towards total recall in industrial anomaly detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14318–14328, 2022.

576

577

578 Thomas Schlegl, Philipp Seeböck, Sebastian M Waldstein, Georg Langs, and Ursula Schmidt-Erfurth. f-anogan: Fast unsupervised anomaly detection with generative adversarial networks. *Medical image analysis*, 54:30–44, 2019.

579

580

581 Yucheng Tang, Dong Yang, Wenqi Li, Holger R Roth, Bennett Landman, Daguang Xu, Vishwesh Nath, and Ali Hatamizadeh. Self-supervised pre-training of swin transformers for 3d medical image analysis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 20730–20740, 2022.

582

583

584

585 National Lung Screening Trial Research Team. The national lung screening trial: overview and study design. *Radiology*, 258(1):243–253, 2011.

586

587

588 Emily Tsai, Scott Simpson, Matthew P. Lungren, Michelle Hershman, Leonid Roshkovan, Errol Colak, Bradley J. Erickson, George Shih, Anouk Stein, Jayashree Kalpathy-Cramer, Jody Shen, Mona A.F. Hafez, Susan John, Prabhakar Rajiah, Brian P. Pogatchnik, John Thomas Mongan, Emre Altinmakas, Erik Ranschaert, Felipe Campos Kitamura, Laurens Topff, Linda Moy, Jeffrey P. Kanne, and Carol C. Wu. Medical imaging data resource center - rsna international covid radiology database release 1a - chest ct covid+ (midrc-ricord-1a). *The Cancer Imaging Archive*, 2020.

589

590

591

592

593

594 Jeya Maria Jose Valanarasu, Yucheng Tang, Dong Yang, Ziyue Xu, Can Zhao, Wenqi Li, Vishal M
 595 Patel, Bennett Landman, Daguang Xu, Yufan He, et al. Disruptive autoencoders: Leveraging
 596 low-level features for 3d medical image pre-training. *arXiv preprint arXiv:2307.16896*, 2023.
 597

598 Tassilo Wald, Constantin Ulrich, Stanislav Lukyanenko, Andrei Goncharov, Alberto Paderno, Max-
 599 imilian Miller, Leander Maerkisch, Paul Jaeger, and Klaus Maier-Hein. Revisiting mae pre-
 600 training for 3d medical image segmentation. In *Proceedings of the Computer Vision and Pattern
 601 Recognition Conference*, pp. 5186–5196, 2025.

602 Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning
 603 for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF conference on computer
 604 vision and pattern recognition*, pp. 3024–3033, 2021.
 605

606 Linshan Wu, Jiaxin Zhuang, and Hao Chen. Voco: A simple-yet-effective volume contrastive learn-
 607 ing framework for 3d medical image analysis. In *Proceedings of the IEEE/CVF Conference on
 608 Computer Vision and Pattern Recognition*, pp. 22873–22882, 2024.

609 Jiawei Yu, Ye Zheng, Xiang Wang, Wei Li, Yushuang Wu, Rui Zhao, and Liwei Wu. Fast-
 610 flow: Unsupervised anomaly detection and localization via 2d normalizing flows. *arXiv preprint
 611 arXiv:2111.07677*, 2021.
 612

613 Vitjan Zavrtanik, Matej Kristan, and Danijel Skočaj. Draem-a discriminatively trained reconstruc-
 614 tion embedding for surface anomaly detection. In *Proceedings of the IEEE/CVF International
 615 Conference on Computer Vision*, pp. 8330–8339, 2021.

616 Yixuan Zhou, Xing Xu, Jingkuan Song, Fumin Shen, and Heng Tao Shen. Msflow: Multiscale flow-
 617 based framework for unsupervised anomaly detection. *IEEE Transactions on Neural Networks
 618 and Learning Systems*, 2024.
 619

620 Zongwei Zhou, Vatsal Sodha, Jiaxuan Pang, Michael B Gotway, and Jianming Liang. Models
 621 genesis. *Medical image analysis*, 67:101840, 2021.
 622

623 David Zimmerer, Jens Petersen, Gregor Köhler, Paul Jäger, Peter Full, Klaus Maier-Hein, Tobias
 624 Roß, Tim Adler, Annika Reinke, and Lena Maier-Hein. Medical out-of-distribution analysis
 625 challenge 2022. In *25th International Conference on Medical Image Computing and Computer
 626 Assisted Intervention (MICCAI 2022)*. Zenodo, 2022.
 627

628 A DICE SCORES UNDERESTIMATION IN UNSUPERVISED SETTING

630 B SELF-SUPERVISED LEARNING

632 **InfoNCE.** As in SimCLR Chen et al. (2020), batch of positive pairs $\{(y_i^{(1)}, y_i^{(2)})\}_{i=1}^N$ is passed
 633 through a trainable MLP-projector $g_{\theta\text{proj}}$ and L2-normalized: $z_i^{(k)} = g_{\theta\text{proj}}(y_i^{(k)}) / \|g_{\theta\text{proj}}(y_i^{(k)})\| \in \mathbb{R}^d$,
 634 where $k = 1, 2$ and $i = 1, \dots, N$. Then, the objective is to maximize similarity in positive pairs
 635 while minimizing similarity in negative pairs. To this end, InfoNCE loss is written as:
 636

$$637 \min_{\theta} \sum_{i=1}^N \sum_{k \in \{1,2\}} -\log \frac{\exp(\langle z_i^{(1)}, z_i^{(2)} \rangle / \tau)}{\exp(\langle z_i^{(1)}, z_i^{(2)} \rangle / \tau) + \sum_{j \neq i} \sum_{l \in \{1,2\}} \exp(\langle z_i^{(k)}, z_j^{(l)} \rangle / \tau)}. \quad (1)$$

641 **VICReg.** VICReg objective consists of three terms:

$$642 \min_{\theta} \alpha \cdot \mathcal{L}^{\text{inv}} + \beta \cdot \mathcal{L}^{\text{var}} + \gamma \cdot \mathcal{L}^{\text{cov}}. \quad (2)$$

644 The first term enforces embeddings to be invariant to augmentations:
 645

$$646 \mathcal{L}^{\text{inv}} = \frac{1}{N \cdot D} \sum_{i=1}^N \|z_i^{(1)} - z_i^{(2)}\|^2. \quad (3)$$

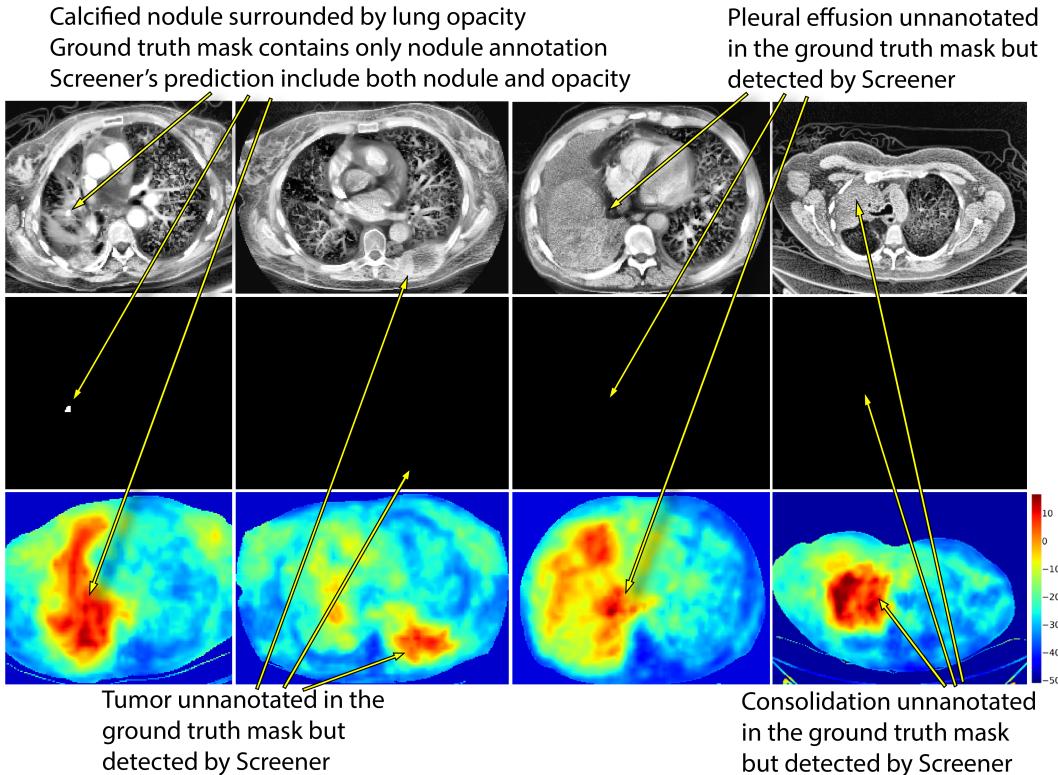


Figure 4: Examples of Screener's **true positive predictions** (third row) counted as "false positives" due to incompleteness of the ground truth masks (second row), leading to Dice score underestimation.

The second term ensures that individual embeddings' dimensions have a least unit variance:

$$\mathcal{L}^{\text{var}} = \sum_{k \in \{1,2\}} \frac{1}{D} \sum_{i=1}^D \max \left(0, 1 - \sqrt{C_{i,i}^{(k)} + \varepsilon} \right). \quad (4)$$

The third term encourages different embeddings' dimensions to be uncorrelated, increasing the total information content of the embeddings:

$$\mathcal{L}^{\text{cov}} = \sum_{k \in \{1,2\}} \frac{1}{D} \sum_{i \neq j} \left(C_{i,j}^{(k)} \right)^2. \quad (5)$$

In VICReg, embeddings $\{z_i^{(k)}\}$ are not L2-normalized and obtained through a trainable MLP-expander which increases the dimensionality up to 8192.

C BASELINE CONDITION MODELS

Sin-cos positional encodings. The existing density-based UVAS methods Gudovskiy et al. (2022); Zhou et al. (2024) for natural images use standard sin-cos positional encodings for conditioning. We also employ them as an option for condition model in our framework. However, let us clarify what we mean by sin-cos positional embeddings in CT images. Note that we never apply descriptor, condition or density models to the whole CT images due to memory constraints. Instead, at all the training stages and at the inference stage of our framework we always apply them to image crops of size $H \times W \times S$, as described in Sections 3.1 and 3.3. When we say that we apply sin-cos positional embeddings condition model to an image crop, we mean that compute sin-cos encodings of absolute positions of its pixels w.r.t. to the whole CT image.

702 **Anatomical positional embeddings.** To implement the idea of learning the conditional distribution
 703 of image patterns at each certain anatomical region, we need a condition model producing
 704 conditions $c[p]$ that encode which anatomical region is present in the image at every position p .
 705 Supervised model for organs’ semantic segmentation would be an ideal condition model for this
 706 purpose. However, to our best knowledge, there is no supervised models that are able to segment
 707 all organs in CT images. That is why, we decided to try the self-supervised APE Goncharov et al.
 708 (2024) model which produces continuous embeddings of anatomical position of CT image pixels.
 709

710 D DESCRIPTOR AND CONDITION MODELS PROBING ANALYSIS

711 Our intuition — introduced in Section 3, and empirically validated in Table 3 — is that an effective
 712 conditioning variables should exclude information about presence / absence of pathologies, while
 713 capturing as much contextual information as possible (e.g., anatomical region, tissue type, patient-
 714 level characteristics etc.). This combination simplifies conditional density modeling and sharpens
 715 anomaly detection.

716 To make this intuition explicit, we conducted an additional probing experiment comparing the
 717 amount of pathology-related information retained by the descriptor model versus the condition
 718 model. We trained a lightweight segmentation head on top of frozen descriptor vs. condition em-
 719 beddings and evaluated Dice scores on four pathology segmentation tasks. As seen in Table 5, our
 720 condition model probing yields substantially lower Dice scores (p -value $< 10^{-10}$ for all datasets).
 721 For lung cancer and liver tumors, scores are essentially zero, confirming that pathology-specific in-
 722 formation is effectively removed. For pneumonia and kidney tumors, non-zero Dice arises largely
 723 because these pathologies occupy substantial portions of the corresponding organs; the segmentation
 724 head tends to predict the whole organ, producing non-zero Dice scores.
 725

726 **Table 5:** Probing of frozen descriptor model (DenseVICReg) and condition model (masking-
 727 invariant DenseVICReg) feature maps on supervised pathology segmentation tasks (results are Dice
 728 scores).

730 Model	731 LIDC	732 MIDRC	733 KiTS	734 LiTS
Descriptor model probe	0.19 ± 0.23	0.40 ± 0.19	0.17 ± 0.18	0.25 ± 0.23
Condition model probe	0.00 ± 0.00	0.27 ± 0.16	0.06 ± 0.10	0.01 ± 0.03

735 To demonstrate that our condition model indeed captures anatomical information better than base-
 736 lines, we also probe it on *liver* segmentation task on LiTS. As seen in Table 6, condition model
 737 embeddings separate liver vs. non-liver voxels substantially better than APE and sinusoidal posi-
 738 tional encodings.

739 **Table 6:** Probing of our masking-invariant condition model vs. APE vs. sin-cos positional embed-
 740 dings on supervised liver segmentation task (results are Dice scores).

742 Model	743 LiTS (liver)
744 Condition model probe	0.82 ± 0.14
745 APE probe	0.62 ± 0.13
746 Sin-cos pos. embed. probe	0.36 ± 0.14

748 E DENSITY MODELS

751 Below, we describe simple Gaussian density model and more expressive learnable Normalizing Flow
 752 model.

753 **Gaussian** marginal density model is written as

$$755 -\log q_{\theta^{\text{dens}}}(y) = \frac{1}{2}(y - \mu)^\top \Sigma^{-1}(y - \mu) + \frac{1}{2} \log \det \Sigma + \text{const}, \quad (6)$$

756 where the trainable parameters θ^{dens} are mean vector μ and diagonal covariance matrix Σ .
 757

758 Conditional Gaussian density model is written as

$$759 -\log q_{\theta^{\text{dens}}}(y | c) = \frac{1}{2}(y - \mu_{\theta^{\text{dens}}}(c))^{\top} (\Sigma_{\theta^{\text{dens}}}(c))^{-1} (y - \mu_{\theta^{\text{dens}}}(c)) + \frac{1}{2} \log \det \Sigma_{\theta^{\text{dens}}}(c) + \text{const}, \quad (7)$$

761 where $\mu_{\theta^{\text{dens}}}$ and $\Sigma_{\theta^{\text{dens}}}$ are MLP nets which take condition $c \in \mathbb{R}^{d^{\text{cond}}}$ as input and predict a
 762 conditional mean vector $\mu_{\theta^{\text{dens}}}(c) \in \mathbb{R}^{d^{\text{desc}}}$ and a vector of conditional variances which is used to construct
 763 the diagonal covariance matrix $\Sigma_{\theta^{\text{dens}}}(c) \in \mathbb{R}^{d^{\text{desc}} \times d^{\text{desc}}}$.
 764

765 As described in Section 3.3, at both training and inference stages, we need to obtain dense negative
 766 log-density maps. Dense prediction by MLP nets $\mu_{\theta^{\text{dens}}}(c)$ and $\Sigma_{\theta^{\text{dens}}}(c)$ can be implemented using
 767 convolutional layers with kernel size $1 \times 1 \times 1$. In practice, we increase this kernel size to $3 \times 3 \times 3$,
 768 which can be equivalently formulated as conditioning on locally aggregated conditions.

769 **Normalizing flow** model of descriptors' marginal distribution is written as:
 770

$$771 -\log p_{\theta^{\text{dens}}}(y) = \frac{1}{2} \|f_{\theta^{\text{dens}}}(y)\|^2 - \log \left| \det \frac{\partial f_{\theta^{\text{dens}}}(y)}{\partial y} \right| + \text{const}, \quad (8)$$

773 where neural net f_{θ} must be invertible and has a tractable Jacobian determinant.

774 Conditional normalizing flow model of descriptors' conditional distribution is given by:
 775

$$776 -\log p_{\theta^{\text{dens}}}(y | c) = \frac{1}{2} \|f_{\theta^{\text{dens}}}(y, c)\|^2 - \log \left| \det \frac{\partial f_{\theta^{\text{dens}}}(y, c)}{\partial y} \right| + \text{const}, \quad (9)$$

778 where neural net $f_{\theta} : \mathbb{R}^{d^{\text{desc}}} \times \mathbb{R}^{d^{\text{cond}}} \rightarrow \mathbb{R}^{d^{\text{desc}}}$ must be invertible w.r.t. the first argument, and the
 779 second term should be tractable.
 780

781 We construct f_{θ} by stacking Glow layers Kingma & Dhariwal (2018): act-norms, invertible linear
 782 transforms and affine coupling layers. Note that at both training and inference stages we apply f_{θ}
 783 to descriptor maps $y \in \mathbb{R}^{h \times w \times s \times d^{\text{desc}}}$ in a pixel-wise manner to obtain dense negative log-density
 784 maps. In conditional model, we apply conditioning in affine coupling layers similar to Gudovskiy
 785 et al. (2022) and also in each act-norm layer by predicting maps of rescaling parameters based on
 786 condition maps.
 787

F DATASETS

790 We utilized several publicly available datasets for training and evaluation summarized in Table 7.
 791 For training, we used the NLST Team (2011), AMOS Ji et al. (2022), and AbdomenAtlas Qu et al.
 792 (2024) datasets. NLST data access is controlled by the National Cancer Institute Data Access Com-
 793 mittee and is available for research use. AMOS is released under a Creative Commons Attribution-
 794 NonCommercial-ShareAlike 4.0 International License (CC BY-NC-SA 4.0). AbdomenAtlas is li-
 795 censed under CC BY-NC-SA 4.0 and intended for academic, research, and educational purposes.
 796 For evaluation, we used the LIDC-IDRI (LIDC) Armato III et al. (2011), MIDRC-RICORD-1a
 797 (MIDRC) Tsai et al. (2020), KiTS Heller et al. (2019), and LiTS Bilic et al. (2023) datasets. LIDC-
 798 IDRI is available through The Cancer Imaging Archive (TCIA) and is typically used under terms
 799 permitting research and education. MIDRC-RICORD-1a is also available through TCIA under sim-
 800 ilar terms, permitting non-commercial use for research and education. The KiTS dataset (version
 801 2021) is available under a CC BY-NC-SA 4.0 license, primarily for non-commercial research and
 802 educational purposes. The LiTS dataset is available for research purposes, often under a Creative
 803 Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (CC BY-NC-ND
 804 4.0) or similar terms, as specified by its organizers. We have used all datasets in accordance with
 805 their specified licenses and terms of use.
 806

G IMPLEMENTATION DETAILS

808 For our Screener model, we preprocess CT volumes by cropping them to dense foreground vox-
 809 ells (thresholded by -500HU), resizing to $1.5 \times 1.5 \times 2.25 \text{ mm}^3$ voxel spacing, clipping intensi-
 810 ties to $[-1000, 300]\text{HU}$ and rescaling them to $[0, 1]$ range. As an important final step we apply

810 Table 7: Summary information on the datasets that we use for training and testing of all models.
811
812

813 Dataset	814 # 3D images	815 Annotated 816 pathology
815 NLST (Team, 2011)	816 25,652	817 –
816 AMOS (Ji et al., 2022)	817 2,123	818 –
817 AbdomenAtlas (Qu et al., 2024)	818 4,607	819 –
818 LIDC (Armato III et al., 2011)	819 1,017	820 lung cancer
819 MIDRC (Tsai et al., 2020)	820 115	821 pneumonia
820 KiTS (Heller et al., 2019)	821 298	822 kidney tumors
821 LiTS (Bilic et al., 2023)	822 117	823 liver tumors

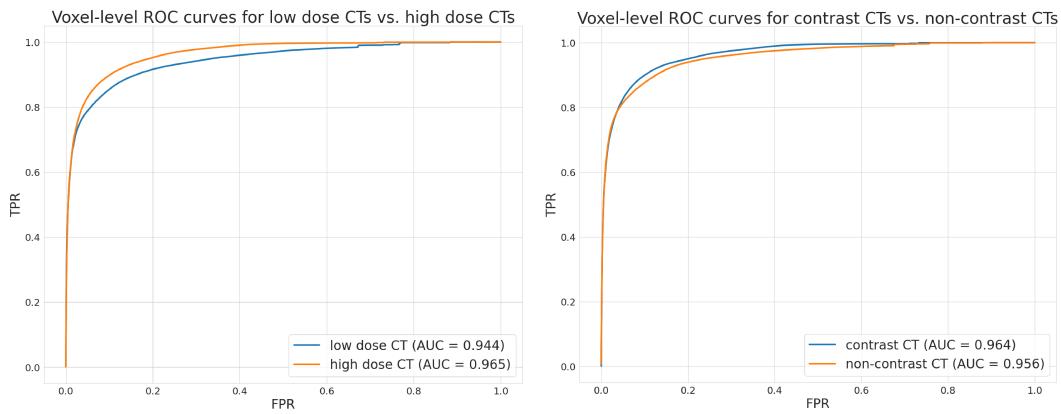
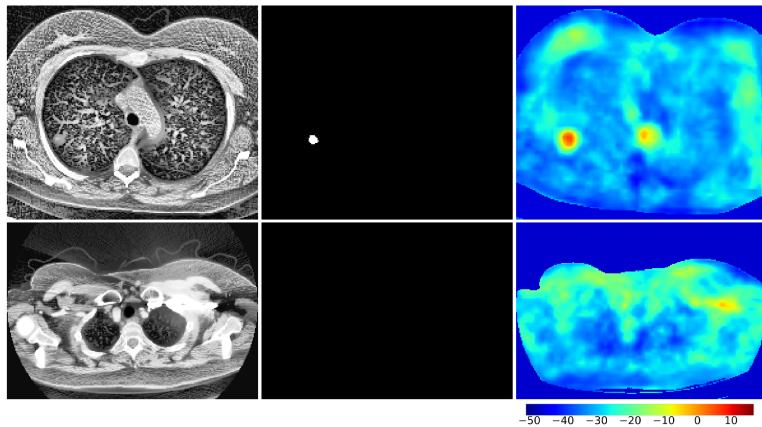
823 CLAHE Pizer et al. (1987). CLAHE ensures that color jitter augmentations preserve information
824 about presence of pathologies during descriptor model training (otherwise, the quality of our method
825 degrades largely).

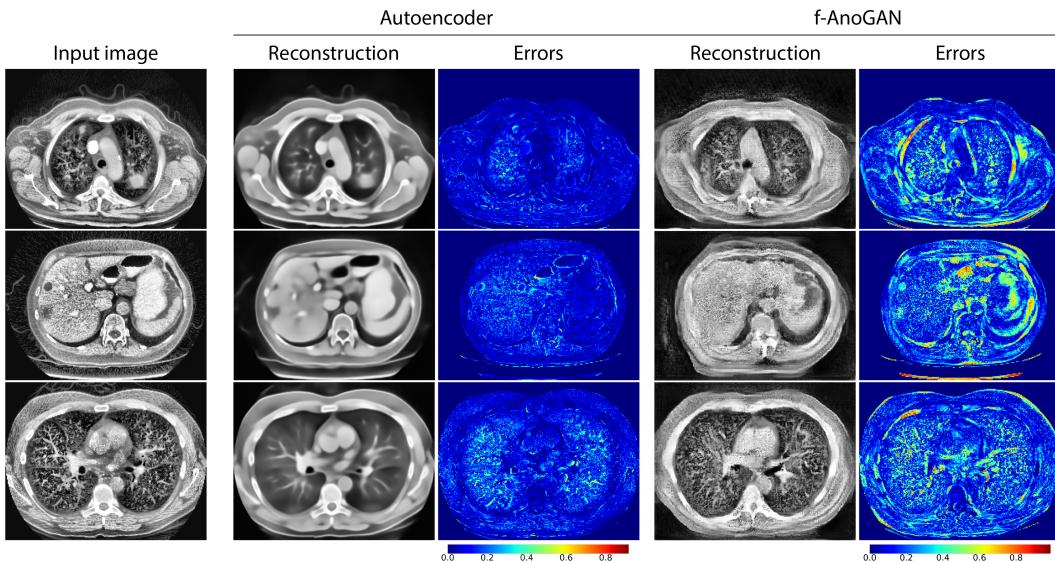
826 We train both the descriptor model and the condition model for 300k batches of $m = 8$ pairs of
827 overlapping patches with $N = 8192$ positive pairs of voxels. The training takes about 3 days on
828 a single NVIDIA RTX H100-80GB GPU. We use AdamW optimizer, warm-up learning rate from
829 0.0 to 0.0003 during first 10K batches, and then reduce it to zero till the end of the training. Weight
830 decay is set to 10^{-6} and gradient clipping to 1.0 norm. Patch size is set to $H \times W \times S = 96 \times 96 \times 64$.
831

832 During density model training, we apply average pooling operations with the $3 \times 3 \times 2$ stride to
833 feature maps produced by the descriptor model and the condition model, following Gudovskiy et al.
834 (2022); Zhou et al. (2024). Thus $h \times w \times s = 32 \times 32 \times 32$. We inject Gaussian noise with
835 0.1 standard deviation both to the descriptors and conditions in order to stabilize the training. We
836 train the density model for 500k batches each containing $m = 4$ patches. This training stage again
837 takes about 3 days on a single NVIDIA RTX H100-80GB GPU. We use the same optimizer and the
838 learning rate scheduler as for the descriptor and condition models.

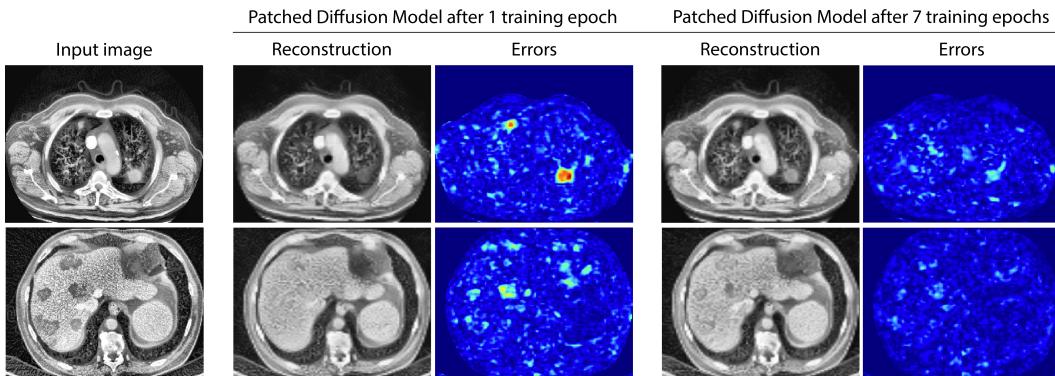
839 The modular Screener model has 133M parameters, patch-based inference for a whole CT volume on
840 NVIDIA RTX H100 GPU requires 4 Gb of GPU memory and takes about 5-10 seconds depending
841 on the number of slices. The distilled Screener has 350M parameters, its patch-based inference
842 requires 5 Gb of GPU memory and takes 0.5-1.0 seconds. We did not observe any difference in
843 quality metrics for the distilled model compared to the modular model.

844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
H ROBUSTNESS ANALYSIS880
881
882
883
884
Figure 5: Comparison of Screener's voxel-level AUROCs on high-dose vs. low-dose and on contrast
vs. non-contrast images from LIDC dataset.885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
Figure 6: Examples of Screener performance on low-dose CT and artifacts. First row: Screener
900 successfully segments lung cancer in low-dose CT. Second row: Screener assigns high anomaly
901 scores to artifact.
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917

918 I ANALYSIS OF RECONSTRUCTION-BASED MODELS
919

938
939 Figure 7: Reconstructions and anomaly maps predicted by Autoencoder Baur et al. (2021) (second
940 and third columns) and f-AnoGAN Schlegl et al. (2019) (last two columns). Autoencoder overfits to
941 reconstruct pathologies and thus fails to detect them. Also Autoencoder produces blurry generations,
942 leading to inaccurate reconstructions and high anomaly scores on fine details (e.g., vessels in the
943 lungs). f-AnoGAN avoids generating pathologies, but the reconstruction quality still is insufficient,
944 resulting in false positive errors. GANs are known to be unstable and sensitive to hyperparameters,
945 necessitating careful tuning and experimentation to achieve optimal results.
946
947



948
949
950 Figure 8: Reconstructions and anomaly maps predicted by Patched Diffusion Model Behrendt et al.
951 (2024) at different epochs. Note that at the beginning of the training (after 1 epoch) it reconstructs
952 healthy regions better than pathologies. However, after 7 epochs, it begins to reconstruct pathologies
953 as well.
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

972 **J FULL-DATA FINE-TUNING**
973974 In the fine-tuning experiments (Table 2, Section 4.2), in order to amplify the effect of pretraining
975 we trained all the models on random data subsets consisting of 25 labeled images. Here, in Table 8,
976 we also provide the results of fine-tuning on full training datasets. As anticipated, the benefits of
977 pretraining become substantially smaller. Noticeably, only supervised pretraining, VoCo Wu et al.
978 (2024), and MAE Wald et al. (2025) show statistically significant improvements on *some* datasets. In
979 contrast, Screener-based pretraining does not yield gains when abundant labeled data are available.
980981 **Table 8: Full-data** fine-tuning results for Screener and other self-supervised pretrained models. We
982 highlight statistically significant improvements (green) or declines (red) relative to nnUNet trained
983 from random initialization.
984

985 Model	986 LIDC	987 MIDRC	988 KitTS	989 LiTS
986 nnUNet (random init.) (Isensee et al., 2021)	986 0.56	986 0.64	986 0.78	986 0.59
987 nnUNet (supervised pretrain.) (Huang et al., 2023)	987 0.58 \uparrow 4% ($p = 0.11$)	987 0.65 \uparrow 2% ($p = 0.12$)	987 0.80 \uparrow 2% ($p < 0.01$)	987 0.62 \uparrow 4% ($p = 0.06$)
988 Grounding DINO (Liu et al., 2023)	988 0.43 \downarrow 23% ($p < 0.01$)	988 0.62 \downarrow 3% ($p < 0.01$)	988 0.63 \downarrow 19% ($p < 0.01$)	988 0.49 \downarrow 18% ($p < 0.01$)
989 LVM-Med (MH Nguyen et al., 2023)	989 0.41 \downarrow 25% ($p < 0.01$)	989 0.61 \downarrow 5% ($p < 0.01$)	989 0.62 \downarrow 21% ($p < 0.01$)	989 0.50 \downarrow 15% ($p < 0.01$)
990 Model Genesis (Zhou et al., 2021)	990 0.57 \uparrow 2% ($p < 0.01$)	990 0.62 \downarrow 2% ($p < 0.01$)	990 0.75 \downarrow 4% ($p < 0.01$)	990 0.53 \downarrow 10% ($p < 0.01$)
991 SwinUNETR (Tang et al., 2022)	991 0.52 \downarrow 6% ($p < 0.01$)	991 0.61 \downarrow 4% ($p < 0.01$)	991 0.70 \downarrow 10% ($p < 0.01$)	991 0.50 \downarrow 16% ($p < 0.01$)
992 DAE (Valanarasu et al., 2023)	992 0.54 \downarrow 3% ($p = 0.08$)	992 0.63 \downarrow 2% ($p < 0.01$)	992 0.75 \downarrow 4% ($p < 0.01$)	992 0.54 \downarrow 9% ($p < 0.01$)
993 VoCo (Wu et al., 2024)	993 0.57 \downarrow 3% ($p = 0.01$)	993 0.64 \uparrow 0% ($p = 0.93$)	993 0.79 \uparrow 2% ($p = 0.05$)	993 0.60 \uparrow 1% ($p = 0.37$)
994 MAE (Wald et al., 2025)	994 0.55 \downarrow 2% ($p = 0.21$)	994 0.66 \uparrow 3% ($p < 0.01$)	994 0.82 \uparrow 5% ($p < 0.01$)	994 0.59 \uparrow 0% ($p = 0.68$)
995 DenseVICReg	995 0.56 \uparrow 0% ($p = 0.83$)	995 0.63 \downarrow 2% ($p < 0.01$)	995 0.77 \downarrow 1% ($p = 0.06$)	995 0.58 \downarrow 3% ($p = 0.19$)
996 Screener (ours)	996 0.55 \uparrow 0% ($p = 0.31$)	996 0.65 \uparrow 1% ($p = 0.46$)	996 0.76 \downarrow 2% ($p < 0.01$)	996 0.58 \downarrow 2% ($p = 0.22$)

997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025