

Trade-Offs in Fine-Tuned Diffusion Models between Accuracy and Interpretability

Mischa Dombrowski¹, Hadrien Reynaud², Johanna P. Müller¹, Matthew Baugh²,
Bernhard Kainz^{1,2}

¹Friedrich-Alexander-Universität Erlangen-Nürnberg, DE

²Imperial College London, UK
mischa.dombrowski@fau.de

Abstract

Recent advancements in diffusion models have significantly impacted the trajectory of generative machine learning research, with many adopting the strategy of fine-tuning pre-trained models using domain-specific text-to-image datasets. Notably, this method has been readily employed for medical applications, such as X-ray image synthesis, leveraging the plethora of associated radiology reports. Yet, a prevailing concern is the lack of assurance on whether these models genuinely comprehend their generated content. With the evolution of text-conditional image generation, these models have grown potent enough to facilitate object localization scrutiny. Our research underscores this advancement in the critical realm of medical imaging, emphasizing the crucial role of interpretability. We further unravel a consequential trade-off between image fidelity – as gauged by conventional metrics – and model interpretability in generative diffusion models. Specifically, the adoption of learnable text encoders when fine-tuning results in diminished interpretability. Our in-depth exploration uncovers the underlying factors responsible for this divergence. Consequently, we present a set of design principles for the development of truly *interpretable* generative models. Code is available at <https://github.com/MischaD/chest-distillation>.

Introduction

Automatic detection and localisation of diseases in medical images have great potential to be applied at scale because of the considerable amount of available data that connects images to radiology reports. Recently, multi-modal models have come into focus because of their ability to capture both, textual and visual information and combine them to improve the performance of the model (Li et al. 2022) and its interpretability (Huang et al. 2021). Additionally, (Boecking et al. 2022) showed how vision-language models can greatly benefit from prompt engineering because of the meaningful representations learned by large language models such as (Devlin et al. 2018; Yan et al. 2022; Liu et al. 2020). Furthermore, these cross-modality abilities can also provide interpretable outputs, for example, by using contrastive learning approaches (Zhang et al. 2022; Huang et al. 2021; Gupta et al. 2020) to perform phrase grounding, which associates

certain tokens or words of the input prompt with regions in the image. Due to recent advances in diffusion models, there has been an increased focus on generative approaches to solve common problems, such as categorical data imbalances and counterfactual image generation. Additionally, since their recent popularization by Ho, Jain, and Abbeel (2020), diffusion models have propelled the performance of generative models (Dhariwal and Nichol 2021; Saharia et al. 2021, 2022). These advances have also been adopted in the medical domain, where most studies thus far have focused on improving generative capabilities, such as generating large corpora of MRI scans (Pinaya et al. 2022b) or 4D data (Kim and Ye 2022). Current approaches have reached a level of performance where the models can generate videos of such high quality that clinicians cannot visually distinguish them from real videos (Reynaud et al. 2023). These successes also led to advances in discriminative tasks such as anomaly detection (Pinaya et al. 2022a; Wolleb et al. 2022). One common approach is to use pre-trained diffusion models (Rombach et al. 2022) and fine-tune them to generate, for example, Chest X-ray (CXR) images. The resulting image quality is better compared to training diffusion models from scratch, as demonstrated in (Chambon et al. 2022). These approaches have in common that they do not interpret the results given by the diffusion models, which have recently demonstrated, particularly, interpretable latent spaces (Dombrowski et al. 2023; Tang et al. 2022; Hertz et al. 2022). Interpreting the latent space of diffusion models is of great importance because generative models have to properly represent what the tokens mean in order to generate correctly matching images from them. Even if we validate the generated samples on pre-trained classifiers (Chambon et al. 2022), the results can be deceiving. Many models have not been evaluated on robustness towards unrealistic artefacts introduced by generative models. Furthermore, if classifiers are already able to clearly label images as belonging to a class, it is unclear whether this sample is useful for applications such as data augmentation.

In this paper, we show that the state-of-the-art methods for fine-tuning diffusion models on radiology reports produce models which are no longer interpretable. We do this by analyzing the influence of jointly learning language embedding and image generation. Our experiments show that diffusion models trained for high image quality have no se-

mantic understanding of their conditional input, but instead, learn to produce images based on confounders. Our hypothesis is that jointly learning the text and image embedding is too difficult and that adapting language models to medical tasks should be approached with more care, including a thorough analysis of the interpretability capabilities of the diffusion models. To show this, we evaluate the phrase grounding potential of the diffusion model on the state-of-the-art approach to fine-tune diffusion models. Then we empirically analyze the effect of keeping the textual encoder frozen. This results in statistically inferior generative performance, but retains the interpretability of the model to such an extent that we beat discriminative phrase grounding approaches on two out of eight diseases in a selected clinical downstream scenario. Furthermore, we statistically evaluate the significance of our findings. By demonstrating the effect of losing accuracy to gain interpretability, this paper is the first work that unveils this important trade-off within the domain of generative medical imaging models.

Method

Pretrained foundation model We begin with Stable Diffusion v2 (SDv2) (Rombach et al. 2022) as a pre-trained foundation model. Using pre-trained models as a starting point is a common approach to improve training time and, furthermore, (Chambon et al. 2022) have shown that it produces better results in terms of image quality compared to starting from a randomly initialized model. SDv2 is a latent diffusion model, which means that the diffusion is learned on a reduced image size $Z = 64$. Following (Chambon et al. 2022), we keep the model used to compute the latent dimension fixed. Consequently, we can precompute the latent representation of our input images to speed up the learning time and to fit the entire dataset into system memory. Since SDv2 is a text-to-image diffusion model, it was trained to learn the function $f(y)$ that generates images conditioned on textual input y . At training time, this prompt could be “a photo of a swiss mountain dog”. In order to inject the conditional input into the diffusion model, the text is first tokenized using a tokenizer $T(y)$ and then embedded using a clip-like model (Radford et al. 2021), in this case, a “ViT-H-14” (Dosovitskiy et al. 2020). The tokenizer maps words to integers as input for the textual encoder and has a fixed vocabulary size determined at training time. Short input prompts are padded to a fixed length of T_{max} , which is typically set to 77 at training and inference time. Since the language embedder was trained on normal images and captions, many words from the medical domain do not appear in this vocabulary. The tokenizer solves this issue by splitting unknown words into multiple tokens. “Consolidation” for example exists, whereas the word “Cardiomegaly” is split into four tokens. To formalize this, we define $\tau_i := \forall \tau : \tau \in T(y_i)$ as the set of tokens corresponding to input word i . In the next step, all tokens are encoded using a pre-trained language encoder. It is common to use a learnable language encoder to simultaneously train the language encoder and the diffusion model, as this results in higher image quality (Chambon et al. 2022). **Classifier-free guidance** For image synthesis, we use classifier-free guidance, which has been shown to produce

better results in terms of image fidelity in the medical domain (Chambon et al. 2022). Intuitively, this means that the reverse diffusion step is performed twice, once for unconditional denoising and once for conditional denoising. Then the difference is computed and amplified to push the denoising step further into the direction of the conditioning and improve the quality of the results. To improve the performance of conditional denoising, we do unconditional guidance training, a common technique that drops the conditional input for healthy samples. We executed this in 30% of the “No Finding” cases.

Interpretability Following Zhang and Zhu (2018) we claim that interpreting mid-layer representations of deep neural networks helps to interpret the predictions of the model. In the case of text-conditioned diffusion models this means that we analyze the interpretability by looking at their attention layers following the approach proposed by Dombrowski et al. (2023). Internally, every input token is converted to a learned latent embedding. This embedding is inserted into the model at every level of the denoising U-Net. The denoising U-Net itself consists of multiple channels and pixels with a reduced width and height. Attention is computed at every layer as the dot-product of the latent representation of the image with the embedding of the token:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right) \cdot V \quad (1)$$

where K and V are the learned projection representations of the input token, Q is the learned projection of the input image, and d is a normalization constant. This means that the layers produce output conditioned on input to localize the relevant pixel-wise features. In the next step we can save the attention maps of multiple reverse diffusion steps of all cross-attention layers of the U-Net architecture and compute the mean over it. Attention maps in downsampled layers are resized to the size $Z \times Z$. The resulting attention map has the shape $\mathbb{R}^{T_{max} \times Z \times Z}$. Due to the softmax operation, attention is defined as a probability and the values over T_{max} sum up to one. To get a class prediction, all we have to do, is to compute the location of all tokens τ_i corresponding to this class and compute the average. For in-domain data, this is straightforward as all the words have a single token representation due to them being part of the training dataset of the model. For unknown words, the tokenizer splits the words into several known tokens. Therefore, we aggregate the attention over all tokens corresponding to this word. The pipeline to analyze this type of interpretability is summarized in Fig. 1.

Datasets

In-domain data To fine-tune a diffusion model on in-domain data we take MS-COCO (Lin et al. 2014), a giant dataset comprised of natural images in their usual environment, and fine-tune on the captions of these images. We split the official training set into 98169 train and 10702 validation images. We use the official validation set as our testset as we do not have ground truth segmentation labels available for the test set. If multiple captions are available for a single

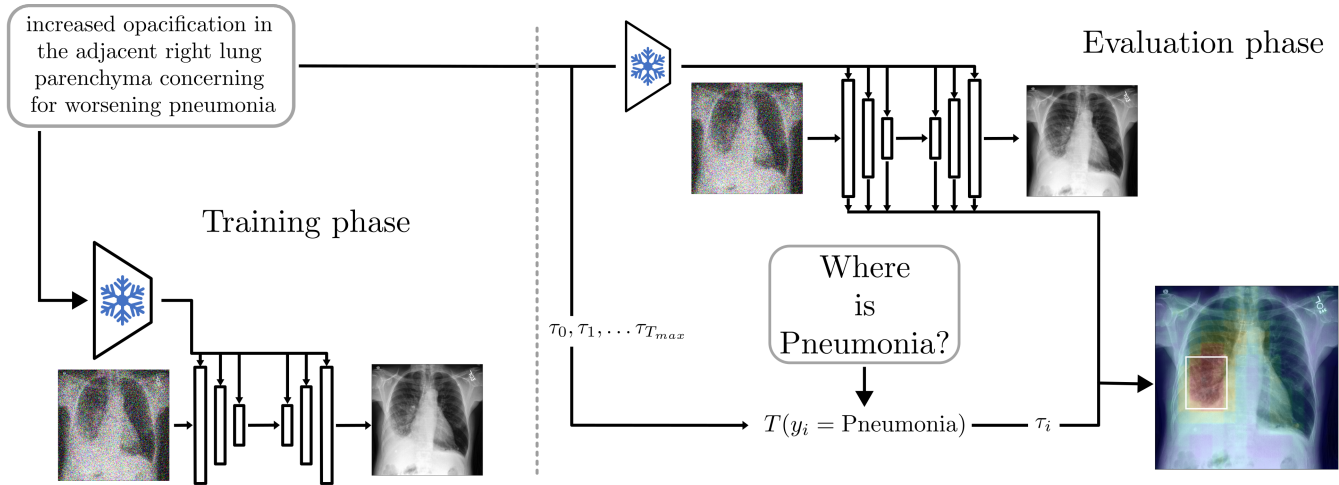


Figure 1: Visual summary of our findings. We investigate different ways of finetuning multi-modal diffusion models. We show that by keeping the language encoder frozen, visualized by the blue snowflake, we can train diffusion models that are interpretable by evaluating the localization abilities of the attention map.

image, we use all captions for training but randomly sample one of them for the training steps. During validation we first filter the dataset such that only images remain where the object class is part of the caption to make the task of token-based localization easier. For example, for the class of birds “two seagulls standing on a handrail” would be removed but “two birds standing on a hand rail” would be left in. Since these captions are short descriptions of the content of the image this step introduces a bias towards only evaluating our model on images where the object is prominent.

Out-of-domain data To fine-tune the models on data from a different domain, we use posterior-anterior (PA) and anterior-posterior (AP) views of MIMIC-CXR (Johnson et al. 2019), a large corpus of paired text and image CXR data. For evaluating localization results we use the phrases and bounding boxes of MS-CXR, a subset of MIMIC-CXR with revised bounding boxes and impressions curated by board-certified radiologists based on the original reports (Boecking et al. 2022) and therefore serve as a higher quality textual description of the image content. Evaluation of multi-label images is done by splitting the image into two instances, one for each disease class. To evaluate segmentation results, we do the following preprocessing step: If the image has multiple bounding boxes describing the same disease, we keep them together and provide the results for the union of multiple bounding boxes. For phrase grounding benchmarks, these are kept separate if the phrases for the bounding boxes are different so that we can compare our method to other methods. To evaluate the statistical significance of our findings we report the average results over five runs. Following (Chambon et al. 2022), we only considered samples with impressions that have over 7 characters for training. MS-CXR and the p19 subset of MIMIC are left out as a test set which results in a total of 162651 samples for the training set. The p19 subset is used to test the performance of our proposed image generation method. We sample 5000

images from this subset, reducing the number of “No Finding” samples to avoid reporting our generative metrics on a subset consisting of predominantly healthy images.

Evaluation Metrics

To evaluate feature localization accuracy we use AUC-ROC, Top-1 accuracy and contrast to noise ratio (CNR) (Boecking et al. 2022) as they allow quantifying the performance without tuning thresholds. Top-1 in this context means the accuracy of the highest predicted pixel being within the boundary of the ground truth annotation. Intuitively this score is supposed to caption the attention of the human observer who looks at the output and focuses on the highest prediction value first. In practice, this value is weak against outliers, *e.g.* if annotations in the corner of the image are detected as outliers. To assess the generative performance we report the Fréchet inception distance (FID) of 5000 images using the standard InceptionV3 and additionally a domain-specific DenseNet121 from (Cohen et al. 2022) (FID_{XRV}) for the CXR data. The images used for comparison are first resampled from the p19 test set to limit the number of “No Finding” input conditions (see Appendix). The sampling settings for the diffusion model follow (Chambon et al. 2022). To assess the diversity of our samples, we use MS-SSIM (Wang, Simoncelli, and Bovik 2003) on pairs of 4 images created with the same prompt, averaged over a set of 100 prompts.

Baseline Method

As the baseline method, we use Stable Diffusion v2, a text-to-image model trained for image generation of natural images of size 512×512 (Rombach et al. 2022). We perform all our fine-tuning experiments closely following the recommendations from (Chambon et al. 2022) ($Finetune_L$), by choosing a learning rate of 5×10^{-5} , a batch size of 256 and training for 60000 steps. When sampling, we use a classifier-free guidance scale of 4 over 75 sampling steps

using the PNDM sampler (Liu et al. 2022). When keeping the language encoder frozen (Finetune_F), we observe that the models perform better after only 30000 training steps, due to the simplified training objective and fewer parameters. We put ablations on this in the Appendix. Our approach differs from (Chambon et al. 2022) because we do not exclude AP views from our evaluation. The reason behind this is that the MS-CXR test set contains both AP and PA views. Furthermore, we refrain from limiting the number of healthy samples in our training set. All results are calculated after applying resizing and center-cropping to get to an image size of 512×512 . We split our fine-tuning over 16 80GB A100 GPUs which takes roughly 240 GPU hours with the frozen, and 580 with the learnable language encoder.

Results

In-Domain Fine-Tuning

First, we evaluate our method with respect to segmentation results on MS-COCO data. We argue that this data is much closer to the training dataset of SDv2 as this model was trained to be a general-purpose text-to-image model. Therefore, we first analyze whether the finetuning process itself destroys the ability of the generative model to be interpretable, especially after finetuning them on smaller compute clusters. The original model (Rombach et al. 2022) was trained on a batch size of 2048, which enables the model to learn a joint representation of the text and the image and therefore enables interpretation of the latent space (Dombrowski et al. 2023). In Tab. 2 we show, that finetuning on in-domain data remains possible even on a single Nvidia A100 GPU and a batch size of 16.

Trade-off of learnable vs frozen encoder Interestingly, we can see that the FID improves over time for both the learnable and the frozen encoder. However, in the case of the learnable encoder, the task of jointly learning language embedding and conditional image generation was too hard and therefore, the FID got worse first. We can also see this in terms of interpretability measured in AUC-ROC. The initial value is way worse than the pretrained baseline and the frozen model. However, eventually the interpretability and synthetic quality get better. The Frozen finetuning works better, with an improved FID score compared to the baseline even for a low number of steps. Further finetuning the model leads to the best FID score and a small improvement in terms of AUC-ROC. Interestingly, in terms of segmentation, the pretrained SDv2 model is the best out of all of them. Partially, this can be explained by the training dataset. SDv2 was trained on a subset of the LAION-5B dataset (Schuhmann et al. 2022), which is a huge agglomeration of natural images with text descriptions scraped from the internet. It is likely that a huge chunk of MS-COCO is already part of the training data. Additionally, the huge variability in the training dataset of SDv2 helps the model to understand concepts better and the result is marginally higher interpretability, even if we compare it to a model finetuned to a more specific dataset.

Localization Results

Next, we want to assess the interpretability of these models when training them on data of new domains. We do this by looking at localization metrics computed on the impressions from MS-CXR (Boecking et al. 2022). Choosing tokens τ_i to attend to is not straightforward for textual conditioning on medical data. Hence, in order to compare with phrase-grounding approaches, if the name of the disease, or some altered version of it appears in the impression, we use the attention maps of this token as the prediction. If this does not occur, we compute the attention for the tokens of all the words, excluding tokens indicating the start of a string, the end of a string, and padding. Manually choosing tokens in this case, *e.g.*, the token for “heart” to localize Cardiomegaly could potentially boost the performance even further, but we avoid this for the sake of generalizability. Since we observed limited localization results for the learnable approach, we do not compute the absolute value of CNR in order to not wrongly overestimate the performance.

Quantitative results are shown in Tab. 1. SDv2 generally has a bad localization performance, although some classes had better scores than expected. One possible explanation is that some diseases have correlations with certain regions, like the regions with medical equipment, lungs, or text labels, that SDv2 often puts higher attention to.

Finetune_L is not able to perform localization consistently. Its localization performance only reaches 54% AUC-ROC which is even worse than the SDv2 baseline and 10% Top-1 accuracy, which is slightly better. Overall, this indicates that this model has not acquired any understanding of the disease-defining features. Interestingly, the localization for “Edema”, “Consolidation” and “Pneumothorax” has a decent Top-1 Accuracy and the CNR results for “Pneumothorax” shows the best AUC-ROC out of all methods. However, the large standard deviation indicates that the understanding of this disease is not robust and might be related to random artefacts.

Trade-off of learnable vs frozen encoder Keeping the language encoder frozen during fine-tuning, on the other hand, shows excellent results in terms of localization. It achieves 79% AUC-ROC, 45% Top-1 accuracy, and 0.920 CNR across all diseases. The attention maps are therefore satisfactory indicators for the location of the disease, and we infer that the model has learned to localize those conditions. Fig. 2 confirms our qualitative observations by showing a considerable gap in the interpretability of the two different diffusion models. Finetune_L mistakenly relates τ_i to image features, such as ribs, spine, or bones, that are unrelated to the disease class. Our proposed approach, on the other hand, shows a good localization performance, including multi-instance input samples.

Phrase grounding benchmarks To substantiate the performance of our method, we can compare it with phrase grounding benchmarks. These methods were trained using contrastive methods and were specifically designed with the task of localization in mind. Our method, on the other hand, provides these localization maps without additional effort, and they can therefore be used for localization in a zero-shot manner. As it can be seen in Tab. 3, our best method

Category	SDv2			Finetune _L			Finetune _F (Ours)		
	AUC ↑	Top-1 ↑	CNR ↑	AUC ↑	Top-1 ↑	CNR ↑	AUC ↑	Top-1 ↑	CNR ↑
Atelectasis	58.5	0.0	0.23	45±4.9	2±1.6	-0.15±.13	86±0.9	80±3.3	1.25 ±.06
Cardiomegaly	53.1	8.4	0.14	45±1.6	2±0.8	-0.16±.05	73±2.3	40±6.8	0.67 ±.07
Consolidation	59.3	2.6	0.21	53±4.8	14±4.3	0.05 ±.11	86±1.6	48±7.5	1.22 ±.10
Edema	80.2	15.2	0.37	54±1.6	29±8.2	0.09 ±.02	89±0.1	70±3.8	1.36 ±.01
Lung Op.	64.7	1.2	0.87	53±3.8	7±4.6	0.05 ±.09	83±0.2	38±1.9	1.06 ±.01
Pl. Effusion	48.8	1.0	0.34	40±4.7	3±1.6	-0.27±.12	83±1.1	66±3.8	1.05 ±.06
Pneumonia	61.3	0.5	0.64	49±2.4	3±1.9	-0.05±.05	85±0.9	59±4.5	1.16 ±.04
Pneumothorax	71.0	10.2	0.01	81±6.9	27±26.5	0.97 ±.33	74±2.7	20±3.4	0.68 ±.08
Average	60.6	5.7	0.321	54±3.7	10±7.4	0.133±.12	79±0.5	45±1.5	0.920±0.03

Table 1: AUC-ROC and Top-1 accuracy for the phrase grounding benchmark of different chest diseases on MS-CXR. We compare our approach compared to SDv2 (Rombach et al. 2022) and Finetune_L (Chambon et al. 2022). Mean and standard deviation over three different trainings are given. The results show that training with a frozen language encoder significantly improves the localization performance throughout all metrics for seven out of eight disease classes.

Method	Steps	Frozen	FID ↓	AUC-ROC ↑
SDv2	0	-	26.1	91.8
Finetune	10000	✗	29.5	79.0
	150000	✗	25.0	90.1
	10000	✓	25.8	91.3
	150000	✓	19.4	91.5

Table 2: MS-COCO finetuning experiment: We compare the generative and object segmentation performance of the two suggested finetuning models. Keeping the language encoder frozen leads to better interpretability and faster convergence.

outperforms the discriminative approaches in terms of localization in two out of eight disease classes, which is remarkable, given that our model is *only* generative by nature. The learnable method, once again, is the worst method out of all generative approaches, with slightly worse results than our SDv2 baseline.

Generative Results

Next, we evaluate the generative quality of the different models. Table 4 shows the quantitative results of the finetuning. The fidelity scores of the models follow the observations reported in (Chambon et al. 2022) that the results are drastically improved by making the textual encoder learnable during fine-tuning, despite the evidence that the model has a worse understanding of what individual tokens mean. To investigate the correctness of the conditional generation, we evaluate the prediction AUC-ROC of a pre-trained classifier, similar to the approach chosen by (Chambon et al. 2022). Even though no ground-truth labels are available for the generated images, it indicates whether the generated artefacts resemble the ones used by trained classifiers. For that, we use the label predictions of the previously used DenseNet121 on our generated images and compare them with the predictions of the same label for the generated healthy images. As a baseline, we take the classification

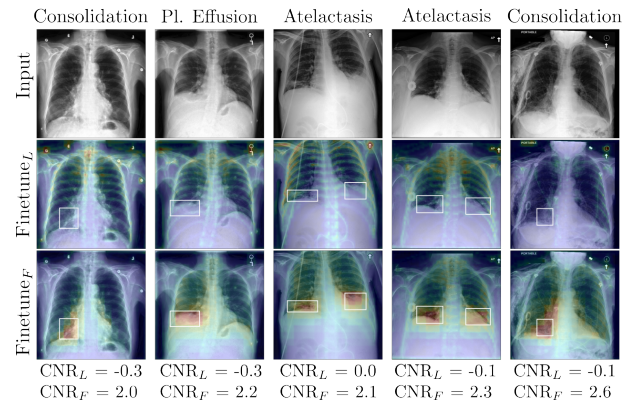


Figure 2: Phrase grounding examples from MS-CXR. White labels show which bounding boxes are ground truth. Pixels with higher importance are shown in red.

accuracy on 5000 real images from the p19 subset. Table 5 shows the results. We can conclude that the class conditional generation follows the same trend as the unconditional generation. The best model is the one that jointly trains language embedding and image generation. However, the difference between both models is very small compared to the FID scores. In fact, Finetune_F even outperforms the other method in six out of eight classes. The only exception are “Pneumothorax”, which was also the only disease class that had good AUC-ROC performance, and “Edema”, which was the only class with good Top-1 accuracy, as shown Tab. 1.

Lack of interpretability To investigate where this lack of interpretability comes from we look for every word that occurs more than 25 times in the test set, how good the localization is and report the three best words and values for all three models in Tab. 6. Interestingly, we can see that Finetune_L has good values for some of the words in the dataset but these words are never consistent with the diseases. The only exception is the end of string token which has a good localization performance throughout all diseases.

Weakly Supervised	Method	Atel.	Card.	Cons.	L-Op.	Edem.	Pnem.	Pnth.	P-Ef.	Avg.
		ConVIRT (Zhang et al. 2022)	0.86	0.64	1.25	0.78	0.68	1.03	0.28	1.02
GLoRIA (Huang et al. 2021)	0.98	0.53	1.38	1.05	0.66	1.18	0.47	1.20	0.930	
BioViL-L (Boecking et al. 2022)	1.17	0.95	1.45	1.19	0.96	1.19	0.74	1.50	1.142	
Zero-shot	SDv2 (Rombach et al. 2022)	0.23	0.14	0.21	0.37	0.87	0.34	0.64	0.01	0.321
	Finetune _L (Chambon et al. 2022)	-0.15	-0.16	0.05	0.05	0.09	-0.05	0.97	-0.27	0.133
	Finetune _F (Ours)	1.25	0.67	1.22	1.06	1.36	1.16	0.68	1.05	0.920

Table 3: Phrase grounding results for contrast-to-noise ratio (CNR) on MS-CXR of our approach compared to state-of-the-art weakly supervised phrase grounding results. Values for ConVIRT, GLoRIA, and BioViL-L are taken from (Boecking et al. 2022). Despite being generative in nature, interpreting the attention layers of the diffusion model and using them for phrase grounding reaches a similar performance as weakly supervised approaches.

Method	FID ↓		FID _{XRV} ↓		MS-SSIM ↓	
	Mimic	MS-CXR	Mimic	MS-CXR	Mimic	MS-CXR
SDv2 (Rombach et al. 2022)	237.6	236.2	104.9	109.4	12.9	11.1
Finetune _L (Chambon et al. 2022)	61.9	60.5	7.7	7.3	10.3	11.4
Finetune _F (Ours)	75.5	75.7	10.1	10.0	10.1	8.2

Table 4: Quantitative comparison of fidelity in terms of general FID and domain-specific FID_{XRV} score and diversity of image generation in terms of MS-SSIM (in %). We compare the results obtained from using impressions from Mimic or from MS-CXR as our prompts for generation.

Due to the language encoding, which is causal, tokens in the future are conditioned on tokens before that. Therefore, the end of string token has learned to relate the text to the location in which the important information is located. Overall the end-of-string (EOS) token of Finetune_L has a CNR of 0.68 which is still a bit worse than EOS for Finetune_F with 0.72 and far from the performance of Finetune_L conditioned on the disease describing words. For phrase grounding this method of localization is not feasible as it fails if parts of the phrase describe different regions of the image.

Influence of Complexity of Text Input

Finally, we want to understand the influence of the text data on the interpretability. Given the results of Tab. 1 we hypothesize that the learnable models struggle with the large variability and the complexity of the impressions of the radiology report in the training datasets. Therefore, we experiment with the variability of the textual input phrases. Our goal is to maximize the interpretability of the model by directly injecting the class condition into the text prompts assuming that we can boost phrase grounding performance. The highest variability is obtained by taking the impressions directly (compare Tab. 1). The lowest variability is obtained by taking the class, converting it to a string, and taking this as input. Additionally, to get a text dataset with medium variability, we take ChatGPT version 3.5 and ask it to generate 100 radiography impressions of CXR images with a certain disease (We show a few examples in the Appendix). Then we randomly sample from this set during training. To enable a fair comparison we only report the results on the subset of the MS-CXR test dataset that has the label of the image in its impression. The results are given in Tab. 7.

Trade-off of learnable vs frozen encoder The learnable models once again completely fail in terms of interpretability. We achieve the highest phrase grounding performance when training the model on the impressions directly. Training on labels surprisingly does not improve the localization results. We assume this is partial because of the domain gap in the text input when we try our localization method on impressions that not only consist of the class label itself. The synthetic impressions achieved the lowest result validating our observation that a high variability of the impressions is key for the model to be interpretable.

Discussion

Our experiments reveal that there are three factors that heavily impact the interpretability of diffusion models. First is the domain shift. Keeping this shift low helps in producing interpretable models. Secondly, keeping the variability high or using strong input signals helps. Finally, we revealed that keeping the language encoder frozen heavily impacts the interpretability of diffusion models, which in turn results in worse generative quality in terms of FID. This gap is smaller for conditional image generation with Finetune_F being slightly better at generation for six out of eight classes. The only two exceptions are “Edema” and “Pneumothorax”, the only two classes of Finetune_L that showed the slightest signs of interpretability. This is evidence that Finetune_L did not properly learn to align textual and spatial information for the majority of input prompts. Furthermore, we believe this is evidence that focusing on designing interpretable diffusion models could also boost their generative ability and that their interpretability is a reasonable indicator of generative performance. Explaining performance differences in

Method	Source	Atel.	Card.	Cons.	L-Op.	Edem.	Pnem.	Pnth.	Effu.	Avg.
p19 Test	Real	83.0	85.5	87.7	80.2	91.4	75.6	84.4	88.7	84.6
Finetune _L (Chambon et al. 2022)	Synthetic	75.8	78.7	77.8	75.1	85.9	66.2	79.6	79.6	77.3
Finetune _F		76.8	79.0	77.9	77.5	84.1	66.7	70.1	84.5	77.1

Table 5: AUC-ROC of a pre-trained classifier (Cohen et al. 2022) on real samples (p19) compared to the accuracy on synthetic images. The metrics are computed as the average over the scores from all diseases. Results for Pleural Effusion are reported using the prediction for Effusion since the classifier does not distinguish between different types of effusions.

Category	SDv2 (Rombach et al. 2022)			Finetune _L (Chambon et al. 2022)			Finetune _F		
	Word	Occ.	CNR	Word	Occ.	CNR	Word	Occ.	CNR
Atelactasis	EOS	61	0.37	EOS	61	0.80	atelectasis	50	1.11
	atelectasis	50	0.25	SOS	61	-0.17	EOS	61	0.85
	SOS	61	-0.17	atelectasis	50	-0.23	SOS	61	-0.85
Cardiomegaly	EOS	333	0.24	EOS	333	0.37	enlarged	222	1.16
	cardiac	220	0.21	SOS	333	0.19	silhouette	219	0.92
	silhouette	219	0.16	moderate	30	-0.09	moderate	30	0.84
Consolidation	lung	56	0.82	EOS	117	0.99	lung	56	1.41
	lobe	32	0.61	patchy	60	0.11	consolidation	93	1.34
	patchy	60	0.58	left	44	0.06	left	44	1.26
Edema	pulmonary	39	1.20	EOS	46	0.90	pulmonary	39	1.25
	EOS	46	0.73	pulmonary	39	0.18	edema	42	1.20
	edema	42	0.53	edema	42	0.13	EOS	46	0.96
Lung opacity	lung	35	0.75	EOS	82	0.97	lung	35	1.33
	ground-glass	32	0.60	the	43	0.10	EOS	82	1.06
	patchy	33	0.55	right	32	0.09	opacity	28	1.00
Effusion	SOS	96	0.11	EOS	96	0.52	left	31	0.84
	EOS	96	0.03	SOS	96	0.03	pleural	88	0.74
	left	31	-0.12	Small	39	-0.20	effusion	51	0.71
Pneumonia	lung	46	0.72	EOS	182	0.89	lung	46	1.33
	lobe	49	0.68	right	58	-0.01	consolidation	31	1.26
	EOS	182	0.45	opacities	30	-0.01	lobe	49	1.26
Pneumothorax	small	70	0.73	small	70	0.95	small	70	0.70
	pneumothorax	226	0.58	left	85	0.70	apical	119	0.66
	apical	119	0.54	apical	119	0.69	pneumothorax	226	0.46

Table 6: Number of occurrences and CNR for the three words with the highest localization accuracy of each model on MS-CXR. SOS and EOS are start and end of string respectively. The EOS string token shows consistently high localization performance for both fine-tuned models. However, the frozen model is the only model that learns the connection between the token describing the disease and its localization.

Textual input	Variability	CNR	
		Learnable	Frozen
Impressions	high	0.15	0.95
ChatGPT	middle	0.22	0.42
Class labels	low	0.29	0.80

Table 7: CNR results for different models in relation to the variability of the textual input prompts. Maximizing the variability of the prompts also maximizes the interpretability.

medical imaging is of utmost importance as the assessment requires trained experts, whose time is often limited. We worked around this by using FID, which has its own limitation due to it being a trained metric and requiring too many samples to be accurate, and conditional generation, which is limited as there is no ground-truth data available. Localization of descriptive features on the other hand is more straightforward, as there is a ground truth available. Given that the results shown in our study match with the results of conditional image generation, we believe that it is a valuable metric that should be considered for text-to-image models.

It would be interesting to perform a large-scale study that analyses the quality of images judged by domain experts in relation to FID and CNR to see which metric better captures generative quality.

Conclusion

In this paper, we show evidence that the state-of-the-art way of fine-tuning diffusion models to medical tasks results in models that have extraordinary image quality but completely loose interpretability. Although machine learning models with limited interpretability may be suitable for certain industrial applications and entertainment purposes, representation models intended for deployment in medical environments are likely to face significant scrutiny regarding their interpretability in the future. To alleviate this issue, we perform rigorous experiments that analyze and unveil important relations in terms of training and interpretability of the model.

Acknowledgements

This work was supported by HPC resources provided by the Erlangen National High Performance Computing Center (NHR@FAU) of FAU Erlangen-Nürnberg under the NHR projects b143dc and b180dc. NHR funding is provided by federal and Bavarian state authorities. NHR@FAU hardware is partially funded by the German Research Foundation (DFG) – 440719683. Key support was received by the ERC - project MIA-NORMAL 101083647, DFG DC-AIDE 512819079, and by the state of Bavaria. H. Reynaud was supported by Ultromics Ltd. and the UKRI Centre for Doctoral Training in Artificial Intelligence for Healthcare (EP/S023283/1).

References

- Boecking, B.; Usuyama, N.; Bannur, S.; Castro, D. C.; Schwaighofer, A.; Hyland, S.; Wetscherek, M.; Naumann, T.; Nori, A.; Alvarez-Valle, J.; et al. 2022. Making the most of text semantics to improve biomedical vision–language processing. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVI*, 1–21. Springer.
- Chambon, P.; Bluethgen, C.; Delbrouck, J.-B.; Van der Sluijs, R.; Polacin, M.; Chaves, J. M. Z.; Abraham, T. M.; Purohit, S.; Langlotz, C. P.; and Chaudhari, A. 2022. RoentGen: Vision-Language Foundation Model for Chest X-ray Generation. *arXiv preprint arXiv:2211.12737*.
- Cohen, J. P.; Viviano, J. D.; Bertin, P.; Morrison, P.; Torabian, P.; Guarrera, M.; Lungren, M. P.; Chaudhari, A.; Brooks, R.; Hashir, M.; and Bertrand, H. 2022. TorchXRrayVision: A library of chest X-ray datasets and models. In *Medical Imaging with Deep Learning*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems*, 34: 8780–8794.
- Dombrowski, M.; Reynaud, H.; Baugh, M.; and Kainz, B. 2023. Foreground-Background Separation through Concept Distillation from Generative Image Foundation Models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 988–998.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. 2020. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.
- Gupta, T.; Vahdat, A.; Chechik, G.; Yang, X.; Kautz, J.; and Hoiem, D. 2020. Contrastive learning for weakly supervised phrase grounding. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III*, 752–768. Springer.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2022. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33: 6840–6851.
- Huang, S.-C.; Shen, L.; Lungren, M. P.; and Yeung, S. 2021. Gloria: A multimodal global-local representation learning framework for label-efficient medical image recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3942–3951.
- Johnson, A. E.; Pollard, T. J.; Berkowitz, S. J.; Greenbaum, N. R.; Lungren, M. P.; Deng, C.-y.; Mark, R. G.; and Horng, S. 2019. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data*, 6(1): 317.
- Kim, B.; and Ye, J. C. 2022. Diffusion deformable model for 4D temporal medical image generation. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part I*, 539–548. Springer.
- Li, L. H.; Zhang, P.; Zhang, H.; Yang, J.; Li, C.; Zhong, Y.; Wang, L.; Yuan, L.; Zhang, L.; Hwang, J.-N.; et al. 2022. Grounded language-image pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10965–10975.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, 740–755. Springer.
- Liu, F.; Shareghi, E.; Meng, Z.; Basaldella, M.; and Collier, N. 2020. Self-alignment pretraining for biomedical entity representations. *arXiv preprint arXiv:2010.11784*.
- Liu, L.; Ren, Y.; Lin, Z.; and Zhao, Z. 2022. Pseudo numerical methods for diffusion models on manifolds. *arXiv preprint arXiv:2202.09778*.
- Pinaya, W. H.; Graham, M. S.; Gray, R.; Da Costa, P. F.; Tudosiu, P.-D.; Wright, P.; Mah, Y. H.; MacKinnon, A. D.;

- Teo, J. T.; Jager, R.; et al. 2022a. Fast unsupervised brain anomaly detection and segmentation with diffusion models. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VIII*, 705–714. Springer.
- Pinaya, W. H.; Tudosi, P.-D.; Dafflon, J.; Da Costa, P. F.; Fernandez, V.; Nachev, P.; Ourselin, S.; and Cardoso, M. J. 2022b. Brain imaging generation with latent diffusion models. In *Deep Generative Models: Second MICCAI Workshop, DGM4MICCAI 2022, Held in Conjunction with MICCAI 2022, Singapore, September 22, 2022, Proceedings*, 117–126. Springer.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Reynaud, H.; Qiao, M.; Dombrowski, M.; Day, T.; Razavi, R.; Gomez, A.; Leeson, P.; and Kainz, B. 2023. Feature-Conditioned Cascaded Video Diffusion Models for Precise Echocardiogram Synthesis. arXiv:2303.12644.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-Resolution Image Synthesis With Latent Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E.; Ghasemipour, S. K. S.; Ayan, B. K.; Mahdavi, S. S.; Lopes, R. G.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv preprint arXiv:2205.11487*.
- Saharia, C.; Ho, J.; Chan, W.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2021. Image super-resolution via iterative refinement. *arXiv:2104.07636*.
- Schuhmann, C.; Beaumont, R.; Vencu, R.; Gordon, C.; Wightman, R.; Cherti, M.; Coombes, T.; Katta, A.; Mullis, C.; Wortsman, M.; et al. 2022. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35: 25278–25294.
- Tang, R.; Pandey, A.; Jiang, Z.; Yang, G.; Kumar, K.; Lin, J.; and Ture, F. 2022. What the daam: Interpreting stable diffusion using cross attention. *arXiv preprint arXiv:2210.04885*.
- Wang, Z.; Simoncelli, E. P.; and Bovik, A. C. 2003. Multi-scale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, 1398–1402. Ieee.
- Wolleb, J.; Bieder, F.; Sandkühler, R.; and Cattin, P. C. 2022. Diffusion models for medical anomaly detection. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VIII*, 35–45. Springer.
- Yan, A.; McAuley, J.; Lu, X.; Du, J.; Chang, E. Y.; Gentili, A.; and Hsu, C.-N. 2022. RadBERT: Adapting transformer-based language models to radiology. *Radiology: Artificial Intelligence*, 4(4): e210258.
- Zhang, Q.-s.; and Zhu, S.-C. 2018. Visual interpretability for deep learning: a survey. *Frontiers of Information Technology & Electronic Engineering*, 19(1): 27–39.
- Zhang, Y.; Jiang, H.; Miura, Y.; Manning, C. D.; and Langlotz, C. P. 2022. Contrastive learning of medical visual representations from paired images and text. In *Machine Learning for Healthcare Conference*, 2–25. PMLR.