

Generating Counterfactual Trajectories with Latent Diffusion Models for Concept Discovery

Payal Varshney^{*,1,2[0009-0009-6571-6134]}, Adriano Lucieri^{*,1,2[0000-0003-1473-4745]}, Christoph Balada^{1,2[0000-0003-0307-7866]},
Andreas Dengel^{1,2[0000-0002-6100-8255]}, and Sheraz Ahmed^{2[0000-0002-4239-6520]}

¹ Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau,
Kaiserslautern, Germany

² German Research Center for Artificial Intelligence GmbH (DFKI),
Kaiserslautern, Germany
`firstname.lastname@dfki.de`

Abstract. Trustworthiness is a major prerequisite for the safe application of opaque deep learning models in high-stakes domains like medicine. Understanding the decision-making process not only contributes to fostering trust but might also reveal previously unknown decision criteria of complex models that could advance the state of medical research. The discovery of decision-relevant concepts from black box models is a particularly challenging task. This study proposes *Concept Discovery through Latent Diffusion-based Counterfactual Trajectories* (CDCT), a novel three-step framework for concept discovery leveraging the superior image synthesis capabilities of diffusion models. In the first step, CDCT uses a Latent Diffusion Model (LDM) to generate a counterfactual trajectory dataset. This dataset is used to derive a disentangled representation of classification-relevant concepts using a Variational Autoencoder (VAE). Finally, a search algorithm is applied to identify relevant concepts in the disentangled latent space. The application of CDCT to a classifier trained on the largest public skin lesion dataset revealed not only the presence of several biases but also meaningful biomarkers. Moreover, the counterfactuals generated within CDCT show better FID scores than those produced by a previously established state-of-the-art method, while being 12 times more resource-efficient. Unsupervised concept discovery holds great potential for the application of trustworthy AI and the further development of human knowledge in various domains. CDCT represents a further step in this direction.

Keywords: Explainability · Counterfactuals · Concept Based Explanations · Latent Diffusion Models · Dermoscopy · Concept Discovery

1 Introduction

Deep learning (DL) algorithms have gained immense popularity for their outstanding performance in the past decade [26,34,25]. The inherent non-linearity

* These authors contributed equally to this work.

and over-parametrization of these algorithms make it particularly challenging to comprehend their reasoning processes. This, among other factors, contributes to a general lack of trust in such black-box models, constituting a major impediment to their application in safety-critical domains such as medicine. With the General Data Protection Regulation [38] and the upcoming European Artificial Intelligence (AI) Act [36], DL-based AI systems must now also comply with regulatory requirements directly concerning the topic of trust. Research on eXplainable AI (XAI) aims at increasing trust by providing more profound insights into the decision-making of opaque black-box models.

The correct interpretation of an explanation by the explainee is fundamental for yielding an understanding of the opaque decision-making process, which is a crucial prerequisite for building trust [23]. Many popular XAI methods only provide information about the relevance of individual input features [27,43,31]. The interpretation of these methods by explainees often suffers from a lack of context about the informative value of feature relevance. Concept-based explanation methods promise to remedy this by allowing to localize [21] human-aligned concepts and quantify [16] their relevance for decision-making. However, the acquisition of fine-grained concept annotations is expensive and inflexible. First efforts have been made towards unsupervised concept discovery [11,19,3], which could reduce the burden of collecting labels for concepts but also comes with the prospect of discovering new, previously unknown concepts. Existing methods are either based on strong structural assumptions [11] or make use of Generative Adversarial Networks (GANs) [19,3] that are usually challenging to train.

This work proposes *Concept Discovery through Latent Diffusion-based Counterfactual Trajectories* (CDCT), the first concept discovery framework based on counterfactual trajectories generated by state-of-the-art latent diffusion models (LDMs) [28]. The first step encompasses the generation of a counterfactual trajectory dataset using a text-conditioned LDM guided by the target classifier. This step aims to extract relevant semantic changes that describe the classifier’s decision boundaries. The counterfactual trajectory dataset is used in the subsequent step to compute a disentangled representation of the classifier-relevant features using a Variational Autoencoder (VAE). In the final step, the VAE’s latent space is exploited for concept discovery through a search algorithm that is built upon [19]. The contribution of our work is as follows:

- we propose CDCT, the first concept discovery framework that leverages latent diffusion-based counterfactual trajectories.
- we demonstrate improved counterfactual explanation capabilities by combining latent diffusion models with classifier guidance.
- we apply CDCT to a skin lesion classifier, demonstrating its ability to reveal classifier-specific concepts including biases and new biomarkers.

2 Related Work

The field of eXplainable AI can be subdivided by the output modality of a given XAI method. Several popular examples in the image domain are based

on the quantification of feature relevance through surrogate models [27], activations [43], or gradients [31]. Recently, human-centric XAI methods like counterfactuals [39,7,35] or concept-based explanation methods [16,18] struck particular interest in the community, as they focus primarily on facilitating the interpretation of explanations by stakeholders and thus achieving their trust.

2.1 Generative Counterfactual Explanations

Counterfactuals constitute alternative classification scenarios with different outcomes given a marginal modification of a reference image. Wachter et al. [39] first introduced a method for the generation of counterfactuals through optimization.

The closest related works use diffusion models (DMs) to generate counterfactuals [13,30,4,29,14]. Some works use an unconditional diffusion model for the generation of counterfactuals guided by the target classifier [13,30]. Augustin et al. [4] utilize adaptive parameterization and cone regularization of gradients. Sanchez et al. [29] made use of a conditional diffusion model to generate healthy counterfactual examples of brain images to localize abnormal lesions. In [14], an adversarial pre-explanation is improved by diffusion-based inpainting to generate minimal counterfactuals. All previous works on DM-based counterfactual generation operate in the pixel space. Our proposed work instead leverages the more recent latent diffusion model to make use of improved synthesis capabilities with better computational efficiency, while focusing on semantic changes through guidance in the latent instead of the pixel space.

2.2 Concept-based Explainability

Concept-based XAI methods aim at the quantification and localization of higher-level concepts that are aligned with human cognitive processes. Concept-based explanations in pedestrian detection might for instance highlight a group of pixels as belonging to the concept *face*. Supervised concept-based explanation methods [16,21,18,8] rely on concept annotations, which can be expensive to obtain. Moreover, in complex domains like medicine, often experts are required for such annotations. However, some works approached the unsupervised discovery of novel concepts. One line of work deals with the identification of concepts from a classifier’s latent representation [11,40,41,9,37]. Ghorbani et al. [11], for instance, proposed a framework for concept discovery based on the segmentation of images, clustering of their latent representations, and quantification of the conceptual influence of clusters. Fel et al. [9] extend the work of [11] and [41] to recursively decompose concepts across layers and localize them. In another approach, Achibat et al. [1] introduced Concept Relevance Propagation (CRP), combining local and global perspectives to address both the ‘where’ and ‘what’ questions for individual predictions. Additionally, Poeta et al. [24] provided a comprehensive survey on concept-based explainable AI, detailing various methodologies and their applications.

Another relevant line of work regards the classifier as a black box and aims at the discovery of concepts through the generative manipulation of input samples [19,3,10,33]. Lang et al. [19], for instance, modified the StyleGAN architecture to systematically examine the GAN’s style space for class-influential concept dimensions. Ghandeharioun et al. [10] approach concept discovery by enforcing the disentanglement of a GAN’s latent concepts, encouraging higher distances between dissimilar and proximity between similar concepts. In contrast to existing concept discovery frameworks, CDCT leverages the power of state-of-the-art, conditional latent diffusion models for the generation of counterfactual trajectories to disentangle and identify classifier-relevant information.

3 Methodology

CDCT is a three-step framework for unsupervised concept discovery. In the first step, a latent diffusion model with classifier guidance is used to generate a counterfactual trajectory dataset, capturing decision-relevant concepts of the target classifier. This dataset is used to obtain a disentangled representation of concepts with the help of a Variational Autoencoder. Finally, decision-relevant dimensions are identified within the VAE by analyzing their impact on the classifier’s output. An overview of the framework is presented in figure 1.




Fig. 1. CDCT is a three-step concept discovery framework. An LDM with classifier guidance is used to generate a counterfactual trajectory dataset. A VAE is trained on this trajectory dataset to disentangle decision-relevant features. Finally, class-relevant dimensions are identified by manipulating the VAE’s latent space and observing the target classifier’s output.

3.1 Generation of Counterfactual Trajectories

The first step aims to extract decision-relevant variations of the input image from the target classifier into a counterfactual trajectory dataset. In contrast to previous works [13,4,29], our proposed framework utilizes a latent diffusion model for a counterfactual generation to benefit from the improved generation efficiency and quality. Text conditioning is used to structure the latent space in accordance with known concepts. Figure 2 depicts the counterfactual generation process with detailed information about the guided step t .

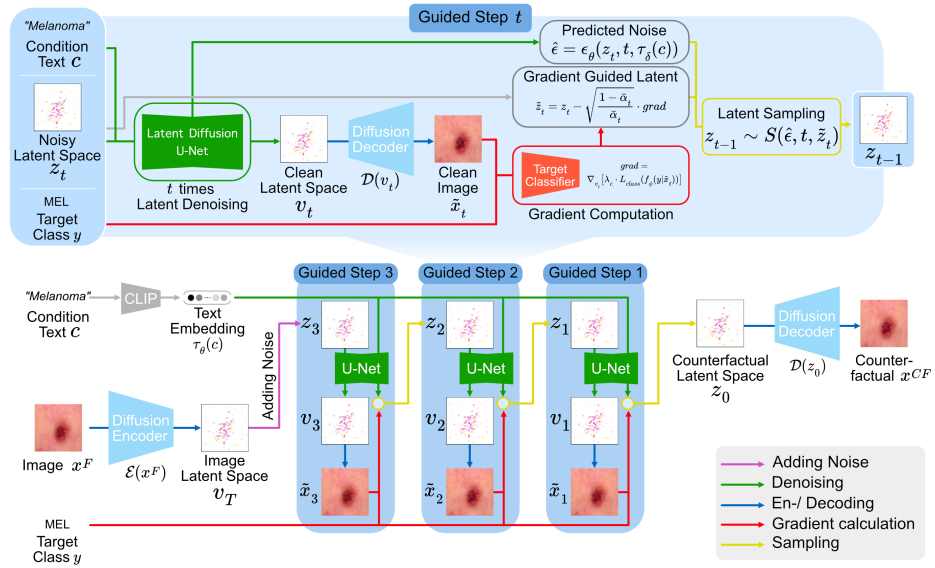


Fig. 2. The counterfactual generation process starts by encoding the image $\mathcal{E}(x^F)$ and perturbing it to obtain z_T (here $T = 3$). The text encoder transforms the condition (c) into an embedding $\tau_\theta(c)$. Both z_T and $\tau_\theta(c)$ are fed to the diffusion model for denoising. At guided step t , we denoise the noisy latent z_t , t times to produce the clean latent v_t , which is decoded into a clean image \tilde{x}_t to calculate the gradient of loss $\mathcal{L}_{\text{class}}$ for updating z_t . We sample the previous less noisy latent z_{t-1} from the estimated noise $\hat{\epsilon}$ and the updated noisy latent \tilde{z}_t . In the final guided step, z_0 is decoded to yield the counterfactual image (x^{CF}).

We start the process with an original image x^F and utilize the encoder \mathcal{E} to transform it into a lower-dimensional latent space. In the forward process, we take its latent representation to calculate a noisy version z_t with variance schedule $\beta_t \in (0, 1)$, where $\alpha_t = 1 - \beta_t$, $\bar{\alpha}_t = \prod_{k=1}^t \alpha_k$ and $1 \leq t \leq T$.

$$z_T = \sqrt{\bar{\alpha}_T} \mathcal{E}(x^F) + \epsilon_T \sqrt{1 - \bar{\alpha}_T}, \quad \text{where } \epsilon_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I}) \quad (1)$$

Subsequently, we traverse the reverse Markov chain to generate a counterfactual with adapted classifier guidance introduced by Jeanneret et al. [13] applied to the latent space. In the reverse process, we iteratively go through the following steps until $t = 0$:

1. The LDM produces the cleaned latent vector v_t by denoising z_t , t times.
2. Then, v_t is decoded to obtain its representation in the image space \tilde{x}_t .
3. The target classifier f_ϕ is then applied on \tilde{x}_t to calculate the gradients of the classification loss $\mathcal{L}_{\text{class}}$.

$$\nabla_{z_t} L(z_t, y) = \frac{1}{\sqrt{\alpha_t}} \nabla_{v_t} [\lambda_c \cdot L_{\text{class}}(f_\phi(y|\tilde{x}_t))] \quad (2)$$

4. The noisy latent z_t is guided using the gradient to produce \tilde{z}_t .
5. z_{t-1} is finally computed from \tilde{z}_t and the noise $\hat{\epsilon}$ predicted at timestamp t .

We utilize the PNDM [20] sampling method represented as $S(\hat{\epsilon}, t, \tilde{z}_t) \rightarrow z_{t-1}$ to get the less noisy sample at each timestamp t . After t repetitions of the aforementioned steps, the counterfactual image is generated by decoding z_0 using the decoder $x^{CF} = \mathcal{D}(z_0)$. A sequence of clean images produced throughout the guiding process, commencing with the factual image and concluding with the counterfactual image, is referred to as a counterfactual trajectory.

3.2 Semantic Space Disentanglement

Counterfactual explanations may simultaneously alter multiple features, complicating the discovery of single, class-relevant concepts. Therefore, the counterfactual trajectory dataset is used in the second step to derive a disentangled representation of decision-relevant biomarkers for the target classifier.

Variational Autoencoders [17] are known for their ability to learn disentangled representations [6]. To enhance the reconstruction fidelity of the VAE, three additional loss functions are integrated besides the default reconstruction loss (\mathcal{L}_{Rec}) and the regularization term (\mathcal{L}_{KLD}). The L1 loss (\mathcal{L}_{L1}) is added to minimize pixel-wise differences, while the Structural Similarity Index Measure (SSIM) [42] ($\mathcal{L}_{\text{SSIM}}$) ensures perceptual image similarity. Lastly, a perceptual loss ($\mathcal{L}_{\text{Perc}}$) based on all the layers of a pre-trained VGG19 is used to capture high-level semantic features. Weighting the KLD loss by w_{kld} balances the trade-off between reconstruction fidelity and disentanglement. The final objective function for training the VAE can be stated as:

$$\mathcal{L}_{\text{VAE}} = \mathcal{L}_{\text{Rec}} + w_{kld} \cdot \mathcal{L}_{\text{KLD}} + \mathcal{L}_{\text{L1}} + \mathcal{L}_{\text{SSIM}} + \mathcal{L}_{\text{Perc}}$$

The VAE is trained on the counterfactual trajectory dataset to disentangle the decision-relevant semantic features.

3.3 Discovery of Relevant Concepts

Finally, our framework modifies the AttFind [19] algorithm to identify decision-relevant concepts in the disentangled latent representation. Algorithm 1 identifies the top latent space dimensions L_y for a class with a given set of images. These dimensions yield the highest average increase in class probability when changing its value during reconstruction through the VAE. In contrast to [19], our approach modifies directions in the range $D_y \in [-3, +3]$ to account for the standard normal distribution, as high magnitudes do not yield artifacts or unrealistic changes but lead to properly emphasized concepts. Moreover, concepts that yield inconsistent change directions were not filtered out, as we argue that conceptual manifestation is not necessarily monotonous on a linear trajectory.

Algorithm 1 Identify Relevant Latent Space Dimensions for a Specific Class

Input: Target classifier f_ϕ , encoder \mathcal{E} , decoder \mathcal{D} ,
Set X of images with $\forall x \in X, f_\phi(x) \neq y$,
Output: Set L_y of top latent space dimensions & set D_y of their directions.

```

for  $x$  in  $X$  do
   $l \leftarrow \mathcal{E}(x), l \in \mathbb{R}^m$  // Encoding
  for  $i = 1, \dots, l$  do
    for  $d$  in  $[-3, -2, -1, 0, 1, 2, 3]$  do
       $l[i] = d$ 
       $\tilde{x} \leftarrow \mathcal{D}(l)$  // Decoding
      Set  $\Delta[x, l, d] = f_{\phi_y}(\tilde{x}) - f_{\phi_y}(\mathcal{D}(\mathcal{E}(x)))$  //Diff. softmax probability class  $y$ 
    end for
  end for
end for
Set  $\bar{\Delta}[l, d] = \text{Mean}(\Delta[x, l, d])$  over all  $x \in X$ 
 $L_y, D_y \leftarrow \text{argsort}(\bar{\Delta}[l, d], \text{descending order, top-K})$ 

```

4 Experiments & Results

4.1 Datasets & Classification Models

Data from the International Skin Imaging Collaboration (ISIC) challenges (2016-2020)³ is used for experimentation. The duplicate removal strategy proposed in Cassidy et al. [5] is followed, resulting in a consolidated dataset of 29,468 samples, further split into train, validation, and test sets. The dataset consists of eight distinct classes, namely *Melanocytic Nevus* (NV), *Squamous Cell Carcinoma* (SCC), *Benign Keratosis* (BKL), *Actinic Keratosis* (AK), *Basal Cell Carcinoma* (BCC), *Melanoma* (MEL), *Dermatofibroma* (DF), and *Vascular Lesions* (VASC). The ISIC 2016, ISIC 2017, PH2 [22], and Derm7pt [15] datasets

³ The data is available at <https://challenge.isic-archive.com/challenges/>.

were used for the training of concept classifiers to provide missing conditioning labels. The selected dermoscopic concepts are *Streaks*, *Pigment Networks*, *Dots & Globules*, *Blue-Whitish Veils*, *Regression Structures*, and *Vascular Structures*.

In this study, a ResNet50 [12] is trained on the consolidated ISIC training dataset, serving as the target classifier. The individual concept classifiers are based on the target classifier, fine-tuned on the respective concept dataset.

4.2 Generation of Counterfactual Trajectories

Stable Diffusion (SD) 2.1⁴ is used as a baseline for the latent diffusion model. To enhance conditioning information for fine-tuning, captions are generated for training images including diagnostic details and skin lesion concepts. Concept classifiers were used to predict the presence or absence of concepts wherever no annotations were available. Only present concepts were used as conditioning information. The SD model is fine-tuned on the consolidated ISIC training dataset with text conditioning for 15K steps and a learning rate of $1e^{-5}$, generating realistic skin lesion samples.

The methodology described in section 3.1 is applied to generate counterfactuals. Experimentation revealed that a gradient loss scale of $\lambda_c > 10$ introduces unrealistic artifacts on the counterfactual images. A value of $\lambda_c = 4$ has therefore been chosen for all results. High values for the initial noise level t compromised the image quality, while $t = 10$ yielded a good trade-off between generative capacity and image quality. Figure 3 shows exemplary counterfactuals from one source image to all other target classes, showcasing the ability of the LDM to generate class-selective biomarkers. More examples can be found in the supplementary material A.3.




Fig. 3. An image of the *Nevus* class alongside its counterfactual images in all other target classes. The second row shows difference maps, providing an easy way to identify the areas of alteration between the original image and each counterfactual.

Counterfactuals generated by different versions of CDCT are quantitatively compared with the results from DiME [14] in table 1. The flip ratio measures the

⁴ Available at <https://huggingface.co/stabilityai/stable-diffusion-2-1>.

frequency at which generated counterfactual images are classified as the target class by the classification model. A qualitative comparison can be found in the supplementary material A.4.

Table 1. L1, L2, FID, and Flip Ratio (FR) scores for counterfactuals generated by different versions of the proposed approach, compared with results from Jeanneret et al. [13]. Subscripts *ft* describe LDM instances fine-tuned on ISIC, while *wo-ft* refers to pre-trained LDMs. All reported results are computed on the consolidated ISIC test dataset for *MEL* and *NV* classes. The best values are highlighted in bold.

Method	L1	L2	FID	FR
DiME [13]	0.02969	0.00176	60.79377	0.86710
Unconditional CDCT _{wo-ft}	0.01930	0.00078	17.36892	1.0
Unconditional CDCT _{ft}	0.01858	0.00073	13.45541	1.0
Conditional CDCT _{wo-ft}	0.01930	0.00078	17.31825	1.0
Conditional CDCT _{ft}	0.01857	0.00073	13.41800	1.0

In every version, CDCT shows a significant improvement in overall metrics, indicating the superiority of counterfactuals generated by LDMs, over traditional approaches. The fine-tuning of the diffusion model yielded a notable improvement in FID, L1, and L2 measures. This aligns with the discovery, that the pre-trained LDM is completely incapable of synthesizing skin lesion images (see supplementary material A.2). The conditional generation of counterfactuals using the target class yielded further minor improvements. It is notable that in comparison to DiME, CDCT achieves a perfect flip ratio, indicating a stronger generative capacity of the LDM. Due to its superior performance, the conditional, fine-tuned CDCT is chosen for all further experiments.

Counterfactual trajectories capture the progression from a factual image to a counterfactual, as shown in figure 4. The counterfactual trajectory dataset is generated by computing trajectories for each ISIC training image into each alternate target class. Each trajectory consists of intermediate images from the 10 guided steps as well as the final counterfactual for 7 target classes, resulting in 78 images per sample. With 23,868 samples in the ISIC training dataset, the resulting counterfactual trajectory dataset contains 1,861,704 samples.

4.3 Semantic Space Disentanglement

A VAE was trained on the counterfactual trajectory dataset to capture all variations reflected by the data. The encoder comprises six ResNet-based down-sampling blocks and final convolution layers, and the decoder architecture mirrors the encoder executing the reverse process of the encoding phase. An extensive hyperparameter search has been conducted, using different loss combinations, weightings, batch sizes, learning rates, and corresponding schedules. The experiments indicate that there is an inherent trade-off between the disentanglement provided by the KLD loss and the reconstruction quality of the VAE.




Fig. 4. Counterfactual trajectory for a *Nevus* with target class *Melanoma*. Along the process, the manifestation of darker, atypical pigment structures can be observed.


All reported results are based on a VAE trained for 23 epochs with the Adam optimizer using a batch size of 128, a learning rate of $2.5e^{-5}$, and an exponential decay learning rate scheduler with a decay factor γ of 0.95. ISIC test images along with their reconstructions are shown in supplementary material A.5.

4.4 Discovery of Relevant Concepts


VAE latent dimensions relevant for the target classifier are determined by algorithm 1 using the ISIC validation dataset. Figure 5 shows a small selection of the most relevant dimensions identified by CDCT. The success rate of a dimension is provided in the sub-caption for each concept and describes the fraction of test cases where the target class probability increased when altering said dimension.

For *Melanoma*, it has been found that many dimensions are related to darkening the skin lesion area while brightening the surrounding skin (e.g., figure 5b). Moreover, some dimensions (e.g., figure 5a) added darker spots and texture within the skin lesion to promote the prediction of *Melanoma*. Another dimension adds dark corner artifacts to the images, evident in figure 5c. The inverse of that dimension was found to promote the prediction of *Nevus*.


Similarly, for other classes, mostly concepts related to overall color are identified. Several relevant dimensions for *Nevus* altered the overall hue of the image to a red color (see figure 5d). For *Basal Cell Carcinoma*, the most relevant dimensions were related to the brightening of the whole image (see figure 5e). For *Dermatofibroma*, the most relevant dimension added white structures in the center of the lesions (see figure 5f). The relevance of the concepts is also emphasized by the high success rates of concepts, as shown in the sub-captions.




(a) Success Rate = 83.85 % - Darkening lesion and surrounding skin.




(b) Success Rate = 79.99 % - Darkening lesion and brightening surrounding skin.




(c) Success Rate = 72.96 % - Adding dark corner artifacts.



(d) Success Rate = 83.74 % - Reddening of lesion and surrounding skin area.



(e) Success Rate = 80.27 % - Whitening of lesion and surrounding skin area.



(f) Success Rate = 65.18 % - Adding central white patch on the lesion.

Fig. 5. Examples of discovered concepts using CDCT with the ISIC dataset. Each row shows two examples of a concept, where the original image, the reconstruction, and the manipulated reconstruction are aligned from left to right.

5 Discussion

In the previous section, the proposed three-step framework has been successfully applied for the global explanation of a skin lesion classifier. The results strongly suggest that CDCT successfully outperforms DiME [13] regarding the quality of generated counterfactuals. Using an LDM allowed for a reduction of the required denoising steps leading to a significant performance improvement in the generation of counterfactuals as compared to DiME (at least 12x speed-up). Moreover, the reduced number of denoising steps together with the generation in the latent space allowed the omission of the perceptual and the L1 loss employed in [13] while yielding better counterfactuals. The conditioning of the LDM for the generation of counterfactual trajectories turned out to yield insignificant improvements in the quality of counterfactuals, but might aid the generation of more diverse and realistic counterfactuals.

Through CDCT several disentangled dimensions were identified. The success rates of the presented dimensions confirm their relevance for the target classifier, and therefore their utility for explanation. By manual inspection, the plausibility of some high-performing dimensions could be established. For instance, in the case of *Melanoma*, the discovered concept of dark corner artifacts is a bias inherent to ISIC, previously highlighted in the literature as well [32]. Other dimensions potentially associated with dataset biases include redness, which tends to increase the prediction confidence of *Nevus*, image brightening, which enhances the likelihood of predicting *Basal Cell Carcinoma*, and image darkening, which promotes the probability of *Melanoma* prediction. Upon nearer investigation of the dataset, it can be observed that these phenomena also reflect in the dataset’s statistics (see supplementary material A.1).

Apart from biases, CDCT has also yielded concepts that can be clinically proven. One dimension that led to a particularly high increase in *Dermatofibroma* prediction probability resulted in the introduction of white structures to the center of the lesion. This concept has been already documented in medical literature as a white, scar-like structure [2] that is highly indicative of the diagnosis, but not yet proven in DL-based classifiers.

6 Limitations & Future Work

Although CDCT was successfully used to discover previously unknown concepts in this work, some limitations remain. Despite the automated selection of classifier-relevant dimensions, a persisting challenge in concept discovery is the manual interpretation of concepts that often requires the help of domain experts. Future works might address this limitation by providing more sophisticated analysis and visualization methods that segment, aggregate, and cluster changes induced by the manipulation of concept dimensions. Moreover, the exploration of new concepts is drastically hampered by the existence of significant biases in the dataset. The linear process of concept discovery might have to be replaced in practice by an iterative approach that involves a debiasing step,

which would facilitate the discovery of meaningful concepts and improve the classifier’s validity. Finally, the utilization of the VAE architecture for deriving a disentangled representation of counterfactual trajectories might be further improved in the future. The VAE training presented a crucial trade-off between reconstruction fidelity and disentanglement through the weight of the KLD loss. As skin lesion classification relies on fine-grained features like subtle textures, reconstruction was favored over disentanglement. This, however, led to one entangled dimension that stood out for its comparatively high relevance throughout all classes, introducing inconclusive changes to the images. Future work should address this issue by applying more sophisticated disentanglement techniques with better fine-grained reconstruction capability, such as StyleGAN.

7 Conclusion

In high-stakes scenarios such as healthcare, DL models play a pivotal role in early disease detection, potentially saving lives. It is essential to provide human-aligned explanations which offer insights into the complex decision-making of models to gain the trust of their users. These insights can potentially reveal new biomarkers relevant in clinical practice. This study provides a new automated framework for the discovery of concepts based on the synthesis of counterfactual trajectories using LDMs.

The proposed CDCT framework is the first to use LDMs with classifier guidance to generate counterfactual explanations. Its proposed counterfactual generation step yields better FID scores compared to the previous DiME [13] method while being up to 12 times more resource-efficient. A counterfactual trajectory dataset is constructed, reflecting relevant semantic changes along the decision boundaries of the target classifier. A disentangled representation of these classifier-relevant cues is derived using a VAE. The automatic and unsupervised exploration of this latent representation yielded valuable insights into the decision-making behavior of a skin lesion classifier, revealing not only biases but also previously unknown biomarkers, supported by first evidence in the medical literature. CDCT can be applied to arbitrary application domains such as radiology and histology, providing a combination of local and global explanations, and therefore paving the way to trustworthy AI finding its way into clinical practice.

8 Acknowledgements

The project was funded by the Federal Ministry for Education and Research (BMBF) with grant number 03ZU1202JA.

References

1. Achtabat, R., Dreyer, M., Eisenbraun, I., Bosse, S., Wiegand, T., Samek, W., Lapuschkin, S.: From attribution maps to human-understandable explanations

- through concept relevance propagation. *Nature Machine Intelligence* **5**(9), 1006–1019 (2023)
2. Agero, A.L.C., Taliercio, S., Dusza, S.W., Salaro, C., Chu, P., Marghoob, A.A.: Conventional and polarized dermoscopy features of dermatofibroma. *Archives of dermatology* **142**(11), 1431–1437 (2006)
 3. Atad, M., Dmytrenko, V., Li, Y., Zhang, X., Keicher, M., Kirschke, J., Wiestler, B., Khakzar, A., Navab, N.: Chexplaining in style: Counterfactual explanations for chest x-rays using stylegan. *arXiv preprint arXiv:2207.07553* (2022)
 4. Augustin, M., Boreiko, V., Croce, F., Hein, M.: Diffusion visual counterfactual explanations. *Advances in Neural Information Processing Systems* **35**, 364–377 (2022)
 5. Cassidy, B., Kendrick, C., Brodzicki, A., Jaworek-Korjakowska, J., Yap, M.H.: Analysis of the isic image datasets: Usage, benchmarks and recommendations. *Medical image analysis* **75**, 102305 (2022)
 6. Chen, R.T., Li, X., Grosse, R.B., Duvenaud, D.K.: Isolating sources of disentanglement in variational autoencoders. *Advances in neural information processing systems* **31** (2018)
 7. Dandl, S., Molnar, C., Binder, M., Bischl, B.: Multi-objective counterfactual explanations. In: *International Conference on Parallel Problem Solving from Nature*. pp. 448–469. Springer (2020)
 8. Espinosa Zarlenga, M., Barbiero, P., Ciravegna, G., Marra, G., Giannini, F., Diligenti, M., Shams, Z., Precioso, F., Melacci, S., Weller, A., et al.: Concept embedding models: Beyond the accuracy-explainability trade-off. *Advances in Neural Information Processing Systems* **35**, 21400–21413 (2022)
 9. Fel, T., Picard, A., Bethune, L., Boissin, T., Vigouroux, D., Colin, J., Cadène, R., Serre, T.: Craft: Concept recursive activation factorization for explainability. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2711–2721 (2023)
 10. Ghandeharioun, A., Kim, B., Li, C.L., Jou, B., Eoff, B., Picard, R.W.: Dissect: Disentangled simultaneous explanations via concept traversals. *arXiv preprint arXiv:2105.15164* (2021)
 11. Ghorbani, A., Wexler, J., Zou, J.Y., Kim, B.: Towards automatic concept-based explanations. *Advances in neural information processing systems* **32** (2019)
 12. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
 13. Jeanneret, G., Simon, L., Jurie, F.: Diffusion models for counterfactual explanations. In: *Proceedings of the Asian Conference on Computer Vision*. pp. 858–876 (2022)
 14. Jeanneret, G., Simon, L., Jurie, F.: Adversarial counterfactual visual explanations. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16425–16435 (2023)
 15. Kawahara, J., Daneshvar, S., Argenziano, G., Hamarneh, G.: Seven-point checklist and skin lesion classification using multitask multimodal neural nets. *IEEE journal of biomedical and health informatics* **23**(2), 538–546 (2018)
 16. Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F., et al.: Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In: *International conference on machine learning*. pp. 2668–2677. PMLR (2018)
 17. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)

18. Koh, P.W., Nguyen, T., Tang, Y.S., Mussmann, S., Pierson, E., Kim, B., Liang, P.: Concept bottleneck models. In: International conference on machine learning. pp. 5338–5348. PMLR (2020)
19. Lang, O., Gandelsman, Y., Yarom, M., Wald, Y., Elidan, G., Hassidim, A., Freeman, W.T., Isola, P., Globerson, A., Irani, M., et al.: Explaining in style: Training a gan to explain a classifier in stylespace. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 693–702 (2021)
20. Liu, L., Ren, Y., Lin, Z., Zhao, Z.: Pseudo numerical methods for diffusion models on manifolds. arXiv preprint arXiv:2202.09778 (2022)
21. Lucieri, A., Bajwa, M.N., Dengel, A., Ahmed, S.: Explaining ai-based decision support systems using concept localization maps. In: International Conference on Neural Information Processing. pp. 185–193. Springer (2020)
22. Mendonça, T., Ferreira, P.M., Marques, J.S., Marcal, A.R., Rozeira, J.: Ph 2-a dermoscopic image database for research and benchmarking. In: 2013 35th annual international conference of the IEEE engineering in medicine and biology society (EMBC). pp. 5437–5440. IEEE (2013)
23. Palacio, S., Lucieri, A., Munir, M., Ahmed, S., Hees, J., Dengel, A.: Xai handbook: towards a unified framework for explainable ai. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3766–3775 (2021)
24. Poeta, E., Ciravegna, G., Pastor, E., Cerquitelli, T., Baralis, E.: Concept-based explainable artificial intelligence: A survey. arXiv preprint arXiv:2312.12936 (2023)
25. Ranjan, R., Sankaranarayanan, S., Bansal, A., Bodla, N., Chen, J.C., Patel, V.M., Castillo, C.D., Chellappa, R.: Deep learning for understanding faces: Machines may be just as good, or better, than humans. *IEEE Signal Processing Magazine* **35**(1), 66–83 (2018)
26. Rank, N., Pfahringer, B., Kempfert, J., Stamm, C., Kühne, T., Schoenrath, F., Falk, V., Eickhoff, C., Meyer, A.: Deep-learning-based real-time prediction of acute kidney injury outperforms human predictive performance. *NPJ digital medicine* **3**(1), 139 (2020)
27. Ribeiro, M.T., Singh, S., Guestrin, C.: " why should i trust you?" explaining the predictions of any classifier. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining. pp. 1135–1144 (2016)
28. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
29. Sanchez, P., Kascenas, A., Liu, X., O’Neil, A.Q., Tsaftaris, S.A.: What is healthy? generative counterfactual diffusion for lesion localization. In: MICCAI Workshop on Deep Generative Models. pp. 34–44. Springer (2022)
30. Sanchez, P., Tsaftaris, S.A.: Diffusion causal models for counterfactual estimation. In: Conference on Causal Learning and Reasoning. pp. 647–668. PMLR (2022)
31. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE international conference on computer vision. pp. 618–626 (2017)
32. Sies, K., Winkler, J.K., Fink, C., Bardehle, F., Toberer, F., Kommoss, F.K., Buhl, T., Enk, A., Rosenberger, A., Haenssle, H.A.: Dark corner artefact and diagnostic performance of a market-approved neural network for skin cancer classification. *JDDG: Journal der Deutschen Dermatologischen Gesellschaft* **19**(6), 842–850 (2021)

33. Song, Y., Shyn, S.K., Kim, K.s.: *Img2tab: Automatic class relevant concept discovery from stylegan features for explainable image classification*. arXiv preprint arXiv:2301.06324 (2023)
34. Sturman, O., von Ziegler, L., Schläppi, C., Akyol, F., Privitera, M., Slominski, D., Grimm, C., Thieren, L., Zerbi, V., Grewe, B., et al.: *Deep learning-based behavioral analysis reaches human accuracy and is capable of outperforming commercial solutions*. *Neuropsychopharmacology* **45**(11), 1942–1952 (2020)
35. Van Looveren, A., Klaise, J.: *Interpretable counterfactual explanations guided by prototypes*. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. pp. 650–665. Springer (2021)
36. Veale, M., Zuiderveen Borgesius, F.: *Demystifying the draft eu artificial intelligence act—analysing the good, the bad, and the unclear elements of the proposed approach*. *Computer Law Review International* **22**(4), 97–112 (2021)
37. Vielhaben, J., Bluecher, S., Strodthoff, N.: *Multi-dimensional concept discovery (mcd): A unifying framework with completeness guarantees*. arXiv preprint arXiv:2301.11911 (2023)
38. Voigt, P., Von dem Bussche, A.: *The eu general data protection regulation (gdpr). A Practical Guide, 1st Ed.*, Cham: Springer International Publishing **10**(3152676), 10–5555 (2017)
39. Wachter, S., Mittelstadt, B., Russell, C.: *Counterfactual explanations without opening the black box: Automated decisions and the gdpr*. *Harv. JL & Tech.* **31**, 841 (2017)
40. Wang, B., Li, L., Nakashima, Y., Nagahara, H.: *Learning bottleneck concepts in image classification*. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10962–10971 (2023)
41. Zhang, R., Madumal, P., Miller, T., Ehinger, K.A., Rubinstein, B.I.: *Invertible concept-based explanations for cnn models with non-negative concept activation vectors*. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 35, pp. 11682–11690 (2021)
42. Zhao, H., Gallo, O., Frosio, I., Kautz, J.: *Loss functions for neural networks for image processing*. arXiv preprint arXiv:1511.08861 (2015)
43. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: *Learning deep features for discriminative localization*. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2921–2929 (2016)

A Supplementary Material

A.1 Color Biases in the ISIC dataset

The concept discovery using the proposed CDCT framework highlighted that color plays a pivotal role in the classification of ISIC. This suggests the presence of biases in the dataset. Table 2 shows the average values of red (in RGB) and saturation (in HSV) channels of images from the different ISIC classes. The values have been computed separately for the lesion region as well as the surrounding skin using a segmentation algorithm. It can be observed, that the *Basal Cell Carcinoma* has the lowest average saturation value on both lesion and skin and that the skin in ISIC’s *Nevus* images exhibit the highest average value for the channel red. This provides first strong evidence for the existence of color-related biases in ISIC. Exploratory analysis of the dataset further reveals that there are large subsets of the *Nevus* class, that exhibit a similar, reddish color hue (see figure 6).

Table 2. Average color values for red (in RGB space) and saturation (in HSV space) for the ISIC dataset.

Area	Value	NV	DF	VASC	MEL	BCC	BKL	AKIEC
Lesion	Red	169	193	171	152	189	167	194
	Saturation	122	081	101	094	067	073	075
Skin	Red	208	201	205	188	196	181	199
	Saturation	070	052	057	046	046	048	055




Fig. 6. Exemplary nevus images from ISIC with typical red hue.

A.2 Generation of Synthetic Samples

Synthetic samples corresponding to skin lesion classes were generated utilizing a pre-trained LDM. The generation process involved using the class names as prompts to produce the samples. However, it was observed that the LDM, without being fine-tuned on a skin lesion dataset, cannot generate realistic images of skin lesions as shown in figure 7.

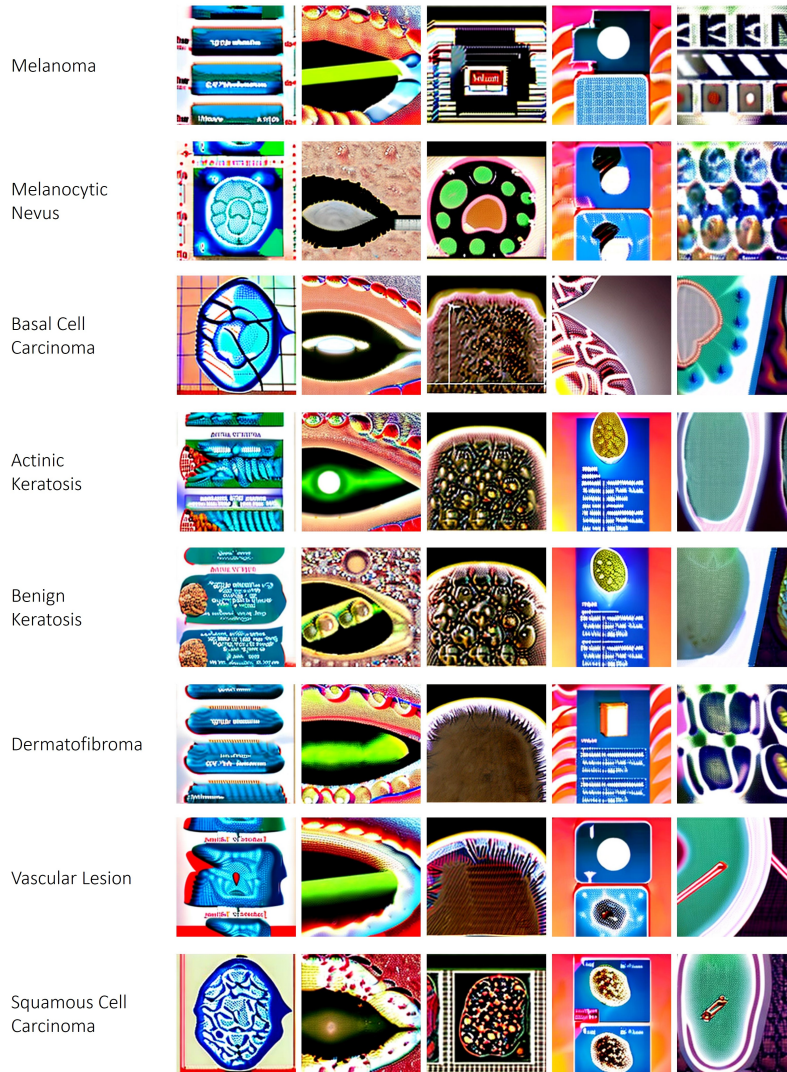


Fig. 7. Synthetic images produced by a pre-trained LDM. The text prompt utilized for generating samples for each skin class label is shown on the left side.

To address this limitation, the LDM was subsequently fine-tuned using the consolidated training dataset from the ISIC challenge. After fine-tuning, samples were generated using text prompts corresponding to the specific classes of skin lesions, with the resultant images being shown in figure 8.




Fig. 8. Synthetic images produced by the LDM fine-tuned on the consolidated ISIC challenge dataset. The text prompt utilized for generating samples for each skin class label is shown on the left side.

A.3 Counterfactual Generation Results

Counterfactuals can be created from a single source class to each target class, showcasing the distinct variations in the generated images for each class. To comprehensively illustrate these differences, we generate counterfactuals from one class to all other classes. For instance, if a sample image is initially classified as *Melanoma*, counterfactuals for this image can be generated for each of the remaining target classes, which include *Melanocytic Nevus*, *Basal Cell Carcinoma*, *Actinic Keratosis*, *Benign Keratosis*, *Dermatofibroma*, *Vascular Lesion*, and *Squamous Cell Carcinoma*, effectively illustrating the discriminative features of each class. Figures 9 to 16 display the counterfactual images generated from each source class showing the distinctive variations when transitioning across all target classes.




Fig. 9. Melanoma to all other target classes.

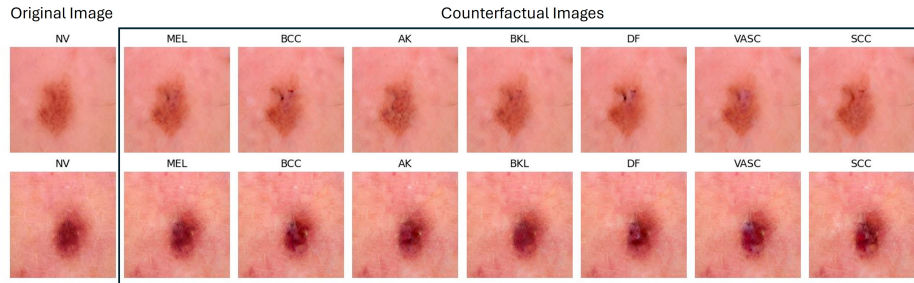


Fig. 10. Nevus to all other target classes.




Fig. 11. Basal Cell Carcinoma to all other target classes.




Fig. 12. Actinic Keratosis to all other target classes.




Fig. 13. Benign Keratosis to all other target classes.




Fig. 14. Dermatofibroma to all other target classes.




Fig. 15. Vascular Lesion to all other target classes.




Fig. 16. Squamous Cell Carcinoma to all other target classes.

A.4 Comparison of Qualitative Result

Figure 17 and 18 show a qualitative comparison of counterfactual results generated with CDCT and the previous state-of-the-art method DiME [13].




Fig. 17. Comparison of counterfactuals generated for eight exemplary images from the *Nevus* class to *Melanoma* using CDCT and DiME [13].




Fig. 18. Comparison of counterfactuals generated for eight exemplary images from the *Melanoma* class to *Nevus* using CDCT and DiME [13].

A.5 VAE reconstruction Results

Figure 19 depicts the reconstructed images generated by the trained VAE. While the reconstruction captures a substantial amount of information, it is noteworthy that the process does not entirely preserve detailed information. While conveying the essence of the original data, the reconstructed images exhibit a level of abstraction where finer details may not be reproduced. Despite losing some fine-grained details, the reconstructed images demonstrate the VAE’s ability to reproduce critical characteristics from the training data.




Fig. 19. Pairs of random ISIC test images (top) and their corresponding reconstructed image (bottom) generated by the trained VAE.