Diffusion Models in Dermatological Education: Flexible High Quality Image Generation for VR-based Clinical Simulations

Leon Pielage[1], Dr. med. Paul Schmidle[2], Prof. Dr. med. Bernhard Marschall[3], Prof. Dr. Benjamin Risse[1]

[1] Institute for Geoinformatics, University of Münster, Germany
[2] Department of Dermatology, University Hospital Münster, Germany
[3] Institute of Medical Education, University of Münster, Germany

medical training

link.cvmls.org/Xmyyp

Paper

Training medical students to accurately recognize malignant melanoma is a crucial competence and part of almost all medical curricular. We here present a pipeline to generate realistic high-resolution imagery of nevus and melanoma skin lesions by using diffusion models. To ensure the required quality and flexibility we introduce three novel guidance strategies and an adapted upsampling approach which enable the generation of user-specified shapes and to integrate the lesions onto pre-defined skin textures. We use our results in a virtual reality (VR) simulation for clinical education. The main advantages of synthetic over real images are the ability to facilitate adjustable learning scenarios and the preservation of patient privacy.

link.cvmls.org/RMBDa

Video

Training

Our architecture is based on the Denoising Diffusion Probabilistic Model (DDPM) with multiple adjustments, utilizing a U-Net-like neural network model ε_θ for the reverse diffusion process with a fixed size of 128×128 pixels and four channels (RGB + segmentation mask). We used separate models for melanoma and nevus classes due to better performance. The pipeline (*Figure 1*) involves training on the HAM10000 dataset from which we sampled random patches between 128×128 and 450×450 pixels. To preserve color information crucial for diagnosis, color normalization utilizing the segmentation masks is applied.

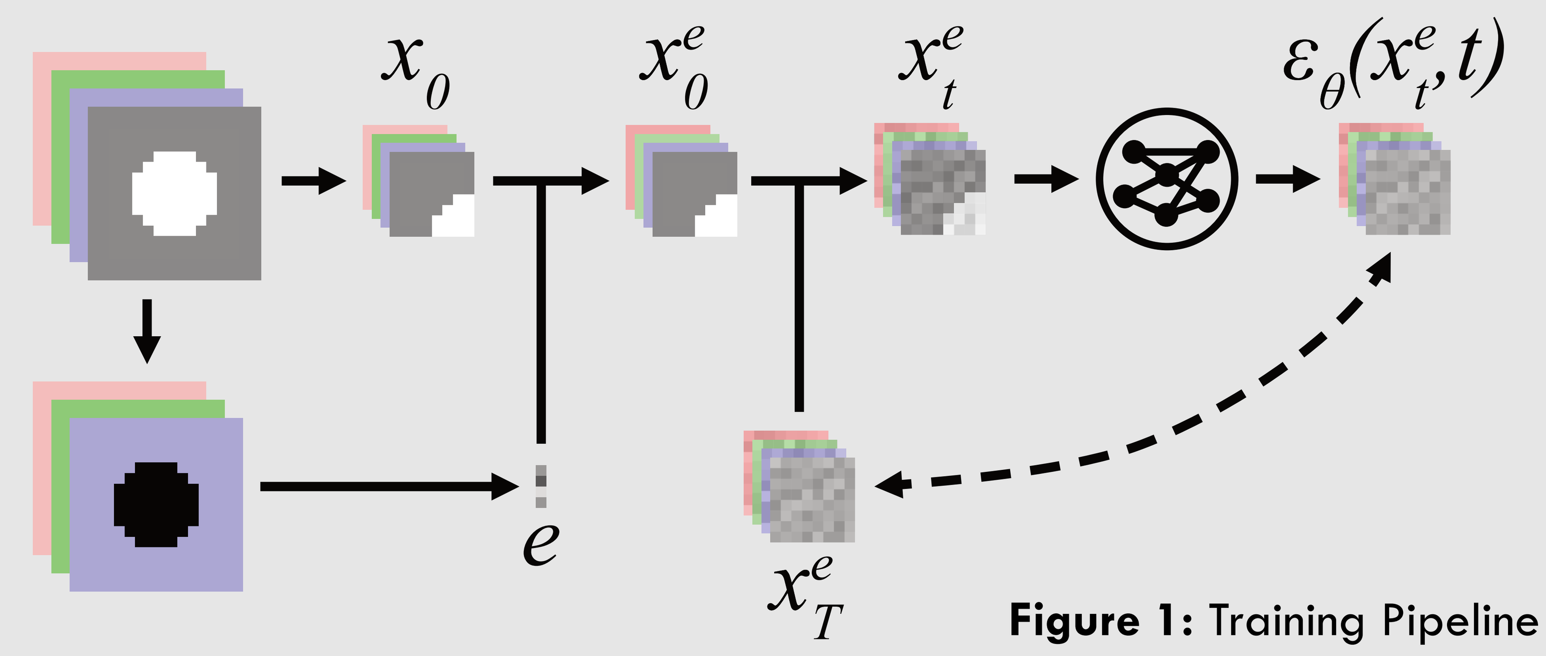

Figure 1: Training Pipeline

Guided Sampling

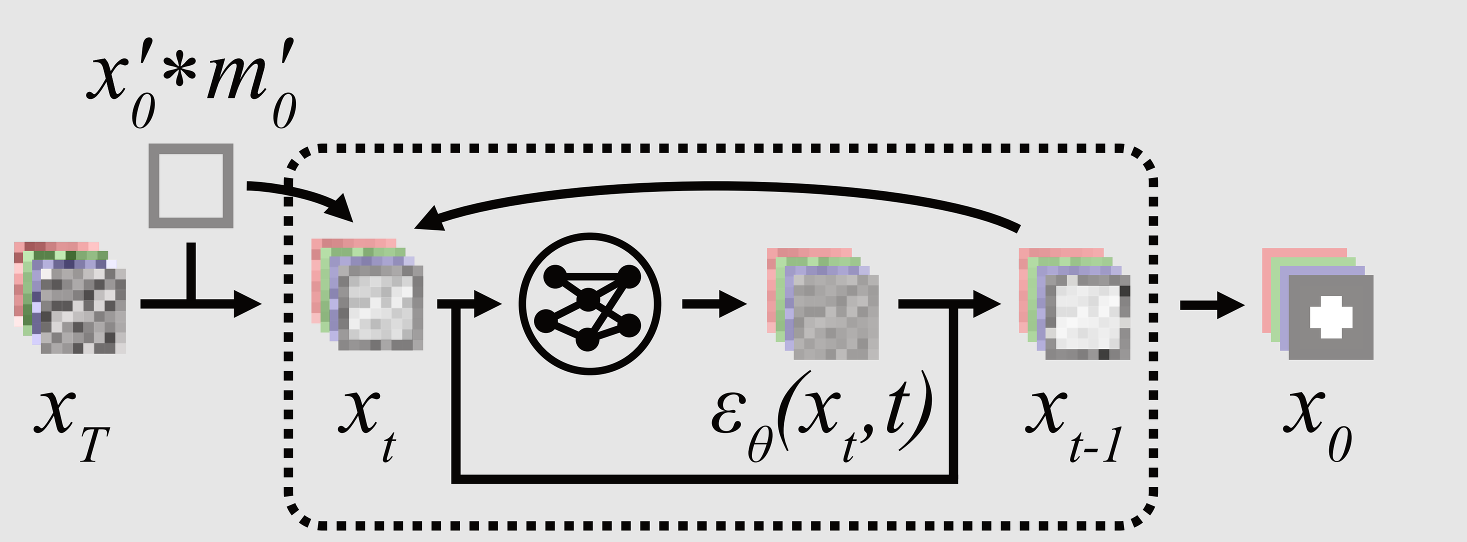

Figure 2: Enforced Class Border Guiding

During sampling we introduce three guiding mechanisms all based on the assumption that a certain region x_0' in the final image x_0 is already known. Given a mask m_0' for this region the unguided sampling process can be extended as shown in *Figure 2*. Thanks to them, it is also possible to place lesions on the virtual agent (*Figure 3*).

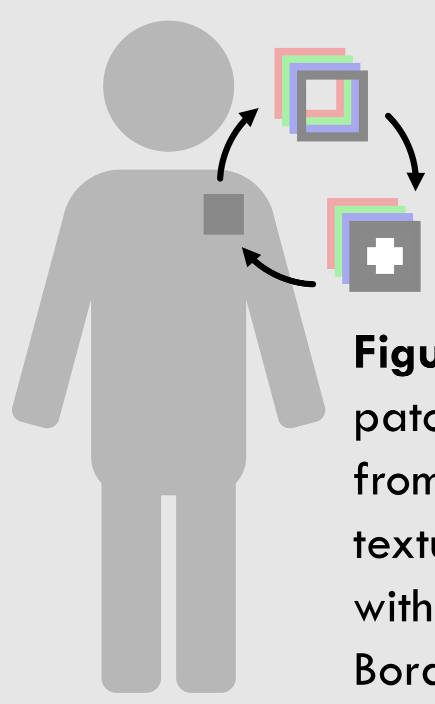

Figure 3: Skin patch cropped from the agent's texture and filled with lesion using Border Guiding

Border Guiding

- Place on existing skin texture
- Pre-defined border region constant
- Only pixels within this border are generatively filled during the sampling process

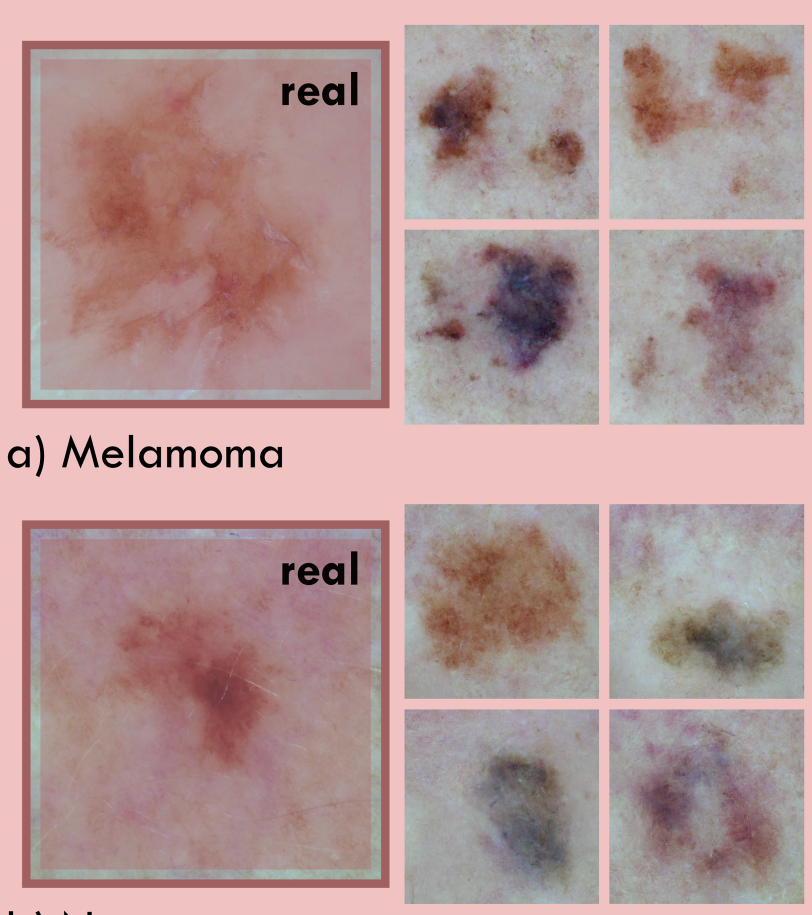

Figure 4: Border Guiding Results

Segmentation Guiding

- Generate user-specified lesion shapes
- Segmentation mask channel constant
- Only color channels are generatively filled during the sampling process

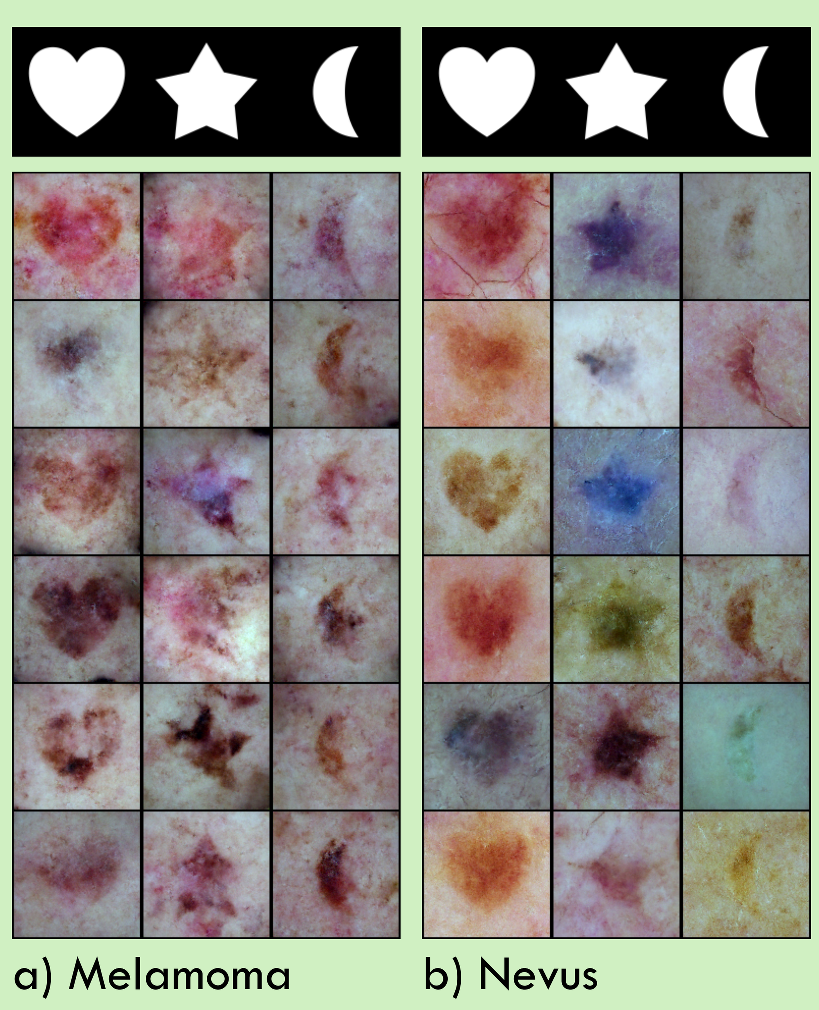

Figure 5: Segmentation Guiding Results

Enforced Class Border Guiding

- Generation of images with non-cropped lesions
- Only the values in the segmentation mask of the border pixels are kept constant

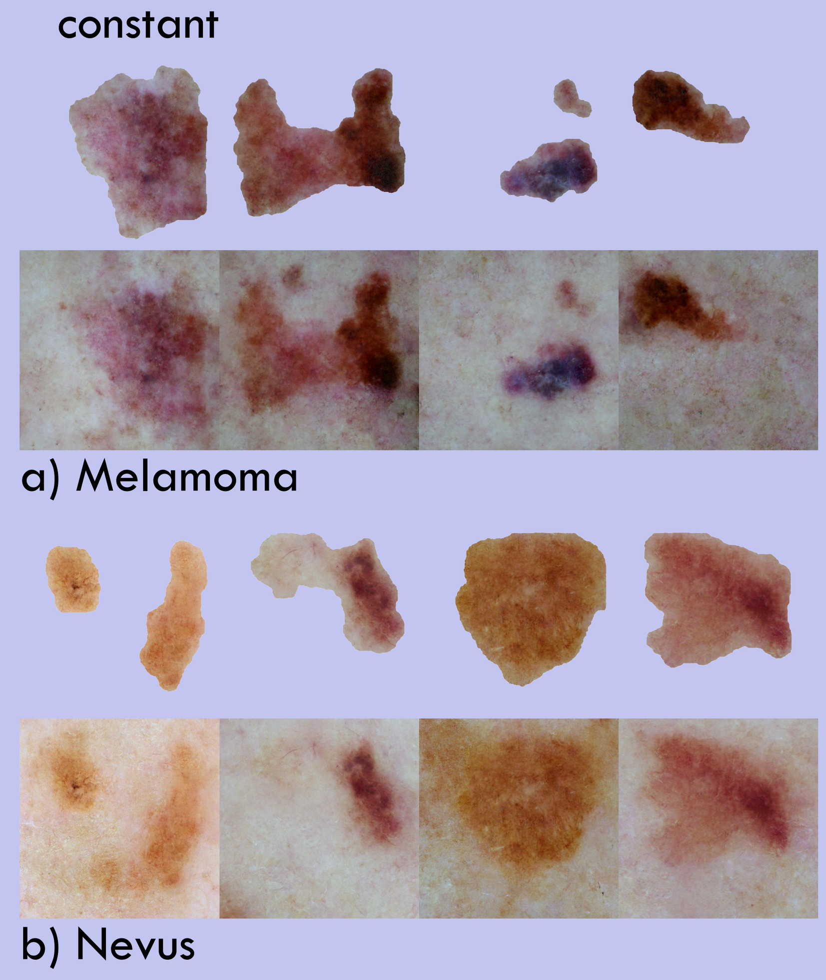

Figure 6: Class Enforced Border Guiding Results

Tile Upsampling

We introduce Tile Upsampling, a strategy that recursively increases the resolution up to 454×454 by splitting the image into tiles, bilinear upsampling, adding noise, and denoising through DDPM iterations with the same trained model used for the initial generation (*Figure 7*). Our Border Guiding strategy is used to prevent seams in overlapping regions, and a novel Pool Guiding (*Figure 8*) addresses color consistency by comparing a downsampled noisy intermediate image with the low-resolution input version for guided sampling.

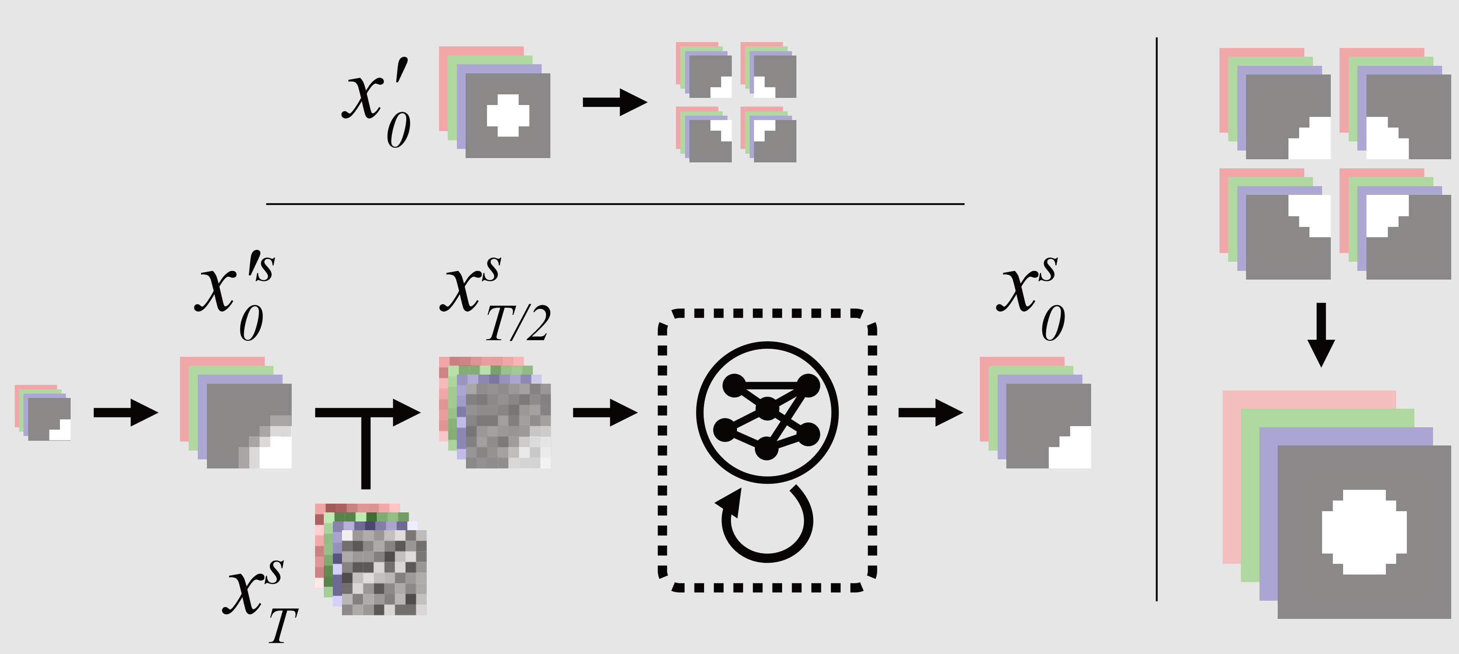

Figure 7: Tile Upsampling

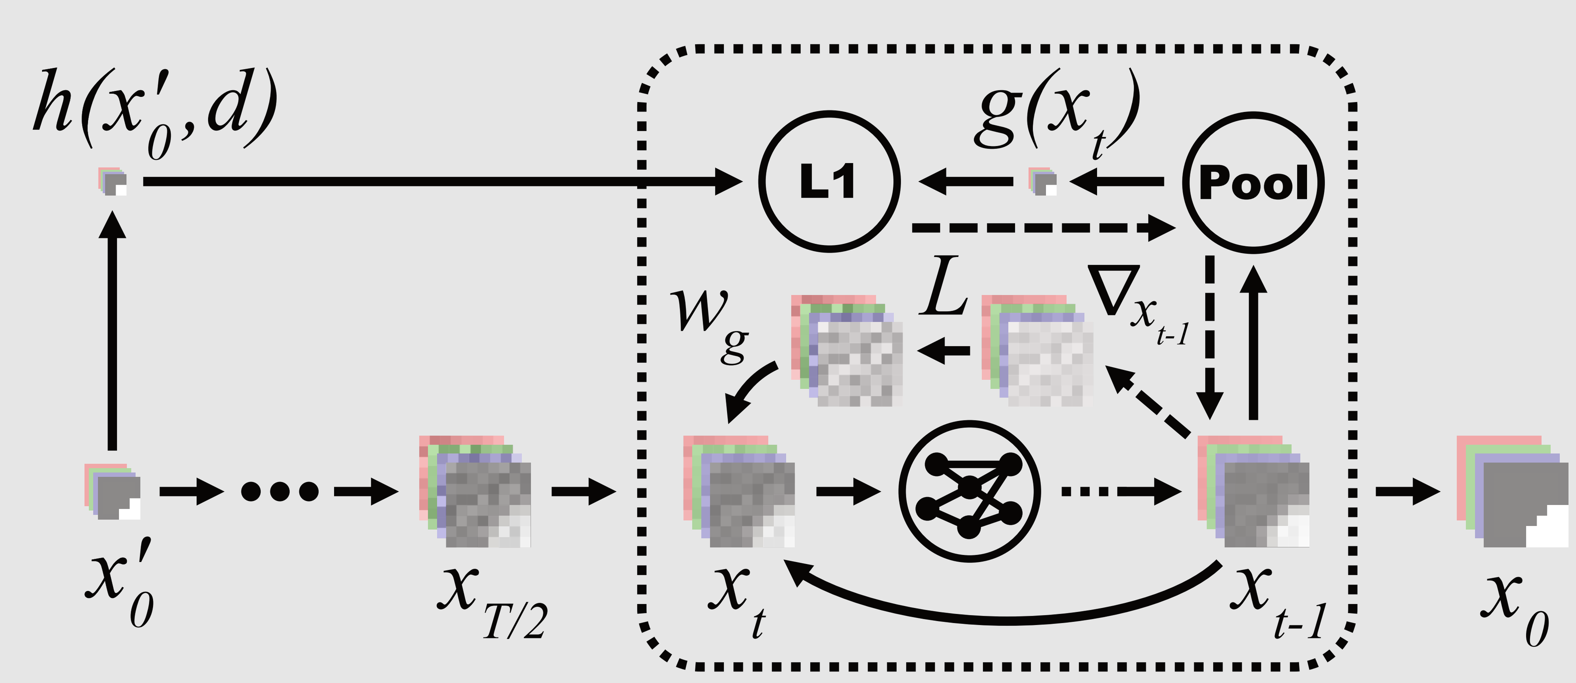

Figure 8: Pool Guiding during upsampling of one tile

Universität Münster | UKM Universitätsklinikum Münster | ifgi Institute for Geoinformatics University of Münster | CVMLS Computer Vision & Machine Learning Systems | link.cvmls.org/k7Xx4 | Gefördert vom Bundesministerium für Bildung und Forschung | Ministerium für Kultur und Wissenschaft des Landes Nordrhein-Westfalen | Ministerium der Finanzen und für Wissenschaft Saarland

leon.pielage@uni-muenster.de | b.risse@uni-muenster.de | www.medical-training-project.de