Supplementary Material: Diffusion Models with Implicit Guidance for Medical Anomaly Detection

Cosmin I. Bercea[1,2], Benedikt Wiestler[1,3], Daniel Rueckert[1,3,5], and Julia A. Schnabel[1,2,4]

[1] Technical University of Munich, Munich, Germany
[2] Helmholtz AI and Helmholtz Center Munch, Munich, Germany
[3] Klinikum Rechts der Isar, Munich, Germany
[4] Kings College London, London, UK
[5] Imperial College London, London, UK
cosmin.bercea@tum.de

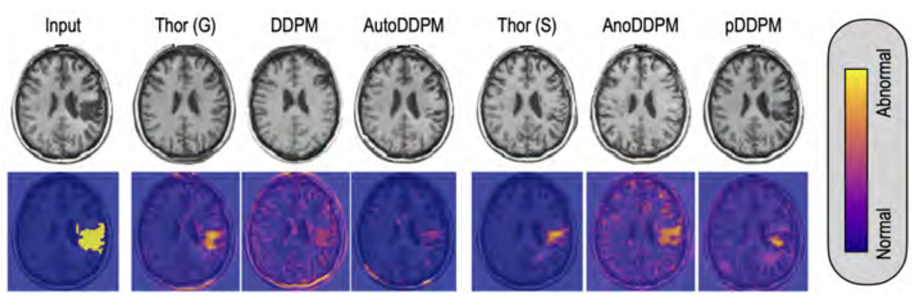

Fig. 1: Qualitative assessment of different diffusion-based models in Brain MRI. *THOR* refines the performance of both DDPM (Gaussian) and AnoDDPM (Simplex), resulting in more accurate reconstructions and enhanced segmentations.

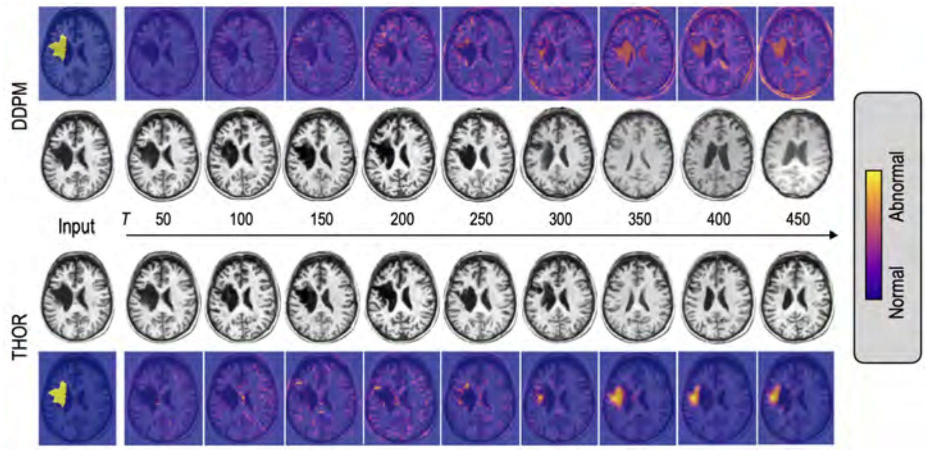

Fig. 2: Comparison of noise scales T for DDPM and *THOR* using Gaussian noise reveals that low noise levels (≤ 300) retain the anomaly's structure, leading to missed detections. At noise levels > 300, DDPM diverges significantly from the original image, affecting even pathology-free areas. Conversely, *THOR* applies temporal harmonization to generate outputs more closely aligned with the input, yet within the healthy spectrum, across all noise levels > 300.

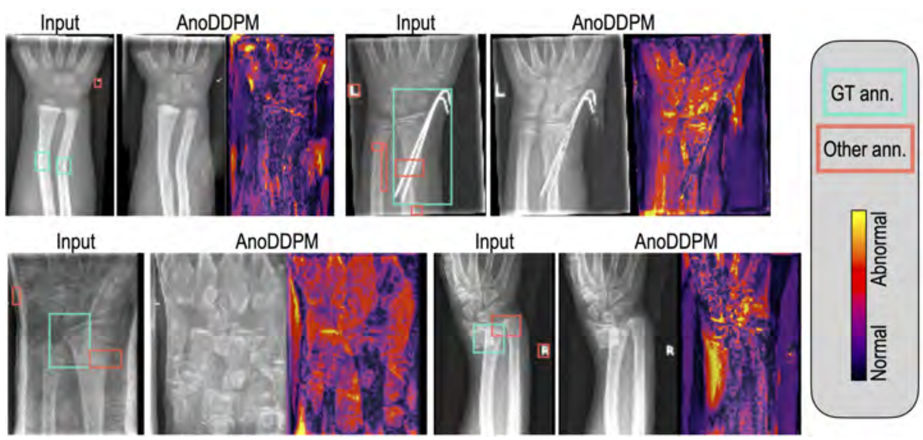

Fig. 3: AnoDDPM's performance with Simplex noise in detecting anomalies in wrist X-ray images demonstrates that the diffusion process reconstructs fractures, unnatural bone positions, and metal implants. Similar to observations in brain MRI experiments, Simplex noise struggles to accurately learn the healthy anatomy, instead, it aims at eliminating structures akin to its coarse noise patterns. Although literature cites its successful application in tumor segmentation, and our work confirms enhanced performance in stroke lesion segmentation, Simplex noise might fail in general anomaly detection tasks.