

Sketch-guided Image Inpainting with Partial Discrete Diffusion Process

Supplementary Material

1. Qualitative Case Studies

1.1. How does the proposed approach work for sketch queries of varying levels of styles and details?

In Figure 1, we conduct a qualitative study to evaluate the effectiveness of our proposed approach in handling out-of-domain non-synthetic sketches with diverse levels of detail and style. Specifically, we mask a portion of the bedroom image and perform inpainting using two sketches depicting a table and two sketches representing a bed. These sketches are selected such that they exhibit significantly different styles and levels of detail. The results demonstrate the robustness of our approach to varying levels of styles.


1.2. How does the proposed approach work for partial sketch queries?

Our research is centered on **object-level sketch-guided** image inpainting. Unlike methods that work on a sketch or partial stroke guidance in sketch manipulation tasks, our work is designed for whole-object sketches, thus impacting its performance on partial sketch query due to our model’s training paradigm. During training, our model is expected to reconstruct the masked region using the sketch of the whole object. So, the model interprets that every input sketch represents the sketch of a full object, which needs to be inpainted fully into the masked region. So, at the test time, when we mask regions and provide a partial sketch, the model interprets that the provided sketch represents the whole object and does not provide appropriate inpainting. This may be addressed by training the proposed model using the partial sketch. However, this substantially digresses from the original goal of this work and hence is beyond the scope of this paper.

1.3. Inpainting using Edgemaps

In some cases, the sketches generated by the proposed method exhibit abstract representations with limited object details. To address this, we explored an alternative approach by utilizing the edge map of the masked image obtained by applying a Laplacian kernel. The quantitative results of inpainting using the edge-map as a query are presented in Table 1. The generated edge maps often capture more intricate object details, contributing to improved inpainting results when compared to inpainting using sketch queries. We continued training the sketch-guided model from the 150th epoch by conditioning it on edge maps until convergence.

Additional Qualitative Results: In Figure 2, we show more qualitative results to contrast the results of our baselines with our method.



Masked Image




Figure 1. Evaluation of our model with non-synthetic sketches with varying styles and levels of detail. Sketches are sourced from the web.

Guidance	Synthetic Sketch	Edge Map
FID	7.72	7.59
LPIPS	0.107	0.106

Table 1. Comparison of Inpainting Results using Edge-Maps as a query/guidance.




Figure 2. Additional results for the qualitative comparison of the proposed method with competing baselines.