

Fuzzy-Conditioned Diffusion and Diffusion Projection Attention Applied to Facial Image Correction

Conference Paper

Author(s): El Helou, Majed

Publication date: 2023

Permanent link: https://doi.org/10.3929/ethz-b-000624777

Rights / license: In Copyright - Non-Commercial Use Permitted

Originally published in: https://doi.org/10.1109/ICIP49359.2023.10223103

This page was generated automatically upon download from the <u>ETH Zurich Research Collection</u>. For more information, please consult the <u>Terms of use</u>.

FUZZY-CONDITIONED DIFFUSION AND DIFFUSION PROJECTION ATTENTION APPLIED TO FACIAL IMAGE CORRECTION

Majed El Helou

Media Technology Center, ETH Zürich, Switzerland

ABSTRACT

Image diffusion has recently shown remarkable performance in image synthesis and implicitly as an image prior. Such a prior has been used with conditioning to solve the inpainting problem, but only supporting binary user-based conditioning.

We derive a fuzzy-conditioned diffusion, where implicit diffusion priors can be exploited with controllable strength. Our fuzzy conditioning can be applied pixel-wise, enabling the modification of different image components to varying degrees. Additionally, we propose an application to facial image correction, where we combine our fuzzy-conditioned diffusion with diffusion-derived attention maps. Our map estimates the degree of anomaly, and we obtain it by projecting on the diffusion space. We show how our approach also leads to interpretable and autonomous facial image correction.

Index Terms— Image-conditioned diffusion, fuzzy conditioning, diffusion projection, autonomous image correction

1. INTRODUCTION

Denoising diffusion methods [1] are effective probabilistic image generation models [2] that learn to synthesize images by gradually recovering them from noisier distributions. They have good sample quality due to their pixel-wise operation, for instance compared with variational auto-encoders [3, 4], or optimal transport distribution mapping [5]. Denoising diffusion currently remains an active research topic, with recent work on balancing the advantages of regular and residual latents [6] and for balancing the trade-off between diversity and high photorealism [7]. For a more extensive overview, we refer the reader to the survey by Croitoru *et al.* [8].

As image generators, diffusion models implicitly learn image distribution representations that can be exploited in a multitude of applications [9], similar to adversarially trained priors [10]. For instance, their implicit prior is useful for image-to-image translation to map across different domains [11, 12], for image super-resolution [13], as well as for synthesizing images with text-based conditioning [14]. Their probabilistic nature also enables diverse image completion through randomized sampling [15]. Stochastic differential equation models were recently used for image editing where the user can control the degree of projection on the perturbed space [16]. A direct inpainting application is presented in [17], providing both state-of-the-art results as well as different diverse outputs due to the stochasticity of the diffusion prior. In contrast, we propose a generalization of inpainting to the non-binary setting, in an effort towards prior contribution control [18, 10] and control over network hallucination [19].

We present a fuzzy-conditioned diffusion where inpainting is a special case. Our model can synthesize novel images from different stochastic samples, while conditioning on a given input image in an adjustable way. The conditioning can be spatially varying and can also be fuzzy (non-binary) in nature. In other words, the diffusion model can be conditioned on an input image to a varying degree, and this degree can change across the image, thus exploiting the underlying pixel-based formulation of denoising diffusion.

Diffusion models can also be used for out-of-distribution detection [20, 21, 22]. The first challenge in out-of-distribution detection is to learn a good image prior for the in-distribution data, which diffusion models can effectively do. Another challenge is to then classify out-of-distribution test cases. The approach in [22], inspired by previous work using deep auto-encoders [23], classifies test cases based on the full stack of reconstruction errors across a sequence of projections on all latent spaces of a diffusion model. A second approach proposes a combination of a diffusion reconstruction with a learned discriminator network to obtain the final classification [20]. Another more medically-oriented approach aims to detect and quantify deviations on the pixel-wise level with a non-Gaussian diffusion [21]. Building on these methods, we construct a diffusion projection attention map on a standard denoising diffusion model. We then exploit this map in our fuzzy-conditioned diffusion to guide an autonomous image correction on a facial image application.

We summarize our contributions as follows. We derive a novel approach to condition diffusion models on images, such that the conditioning can be spatially varying as well as fuzzy in strength. This enables both a generalization of inpainting where the degree of hallucination is controlled, and a generalization of image conditioning to be spatially adaptive. We additionally provide a projection-based anomaly map, enabling our model to perform autonomous face correction.

https://github.com/majedelhelou/FC-Diffusion

2. PROPOSED METHOD

We exploit diffusion models [1] to propose a diffusion with a fuzzy image conditioning (Sec. 2.2) and to obtain a guide-free attention map through projection (Sec. 2.3), the combination of which is applied for autonomous facial image correction.

2.1. Mathematical formulation

We begin with a brief overview of DDPM [2], which learns to map a Gaussian distribution to a data distribution, similar to other generative methods. Assuming the image data distribution to be given by $x_0 \sim q(x_0)$, we follow a Markovian noising process to gradually add Gaussian noise to the data

$$q(x_t|x_{t-1}) \coloneqq \mathcal{N}(x_t; \sqrt{1 - \beta_t} x_{t-1}, \beta_t \mathbb{1}), \qquad (1)$$

where β_t dictates a variance schedule across iterations. We can sample in the diffusion latent space $x_t \sim q(x_t|x_0)$ directly with a single procedure from the distribution

$$q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\overline{\alpha}_t} x_0, (1 - \overline{\alpha}_t)\mathbb{1}), \qquad (2)$$

where $\alpha_t \coloneqq 1 - \beta_t$ and $\overline{\alpha}_t \coloneqq \prod_{d=0}^t \alpha_d$. This leads to $x_t = \sqrt{\overline{\alpha}_t} x_0 + \epsilon \sqrt{1 - \overline{\alpha}_t}, \epsilon \sim \mathcal{N}(0, 1)$, where ϵ is a sampled noise tensor of same dimensions as x_0 . The conditional posterior distribution is derived with Bayes' theorem as

$$q(x_{t-1}|x_t, x_0) = \mathcal{N}(x_{t-1}; \tilde{\mu}(x_t, x_0), \tilde{\beta}_t \mathbb{1}), \qquad (3)$$

where the Gaussian parameters are respectively given by

$$\tilde{\mu}(x_t, x_0) \coloneqq \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1-\bar{\alpha}_t} x_0 + \frac{\sqrt{\bar{\alpha}_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t} x_t,$$

$$\tilde{\beta}_t \coloneqq \frac{1-\bar{\alpha}_{t-1}}{1-\bar{\alpha}_t} \beta_t.$$
(4)

The true $q(x_{t-1}|x_t)$ is estimated with a learned distribution

$$p_{\theta}(x_{t-1}|x_t) \coloneqq \mathcal{N}(x_{t-1}; \mu_{\theta}(x_t, t), \Sigma_{\theta}(x_t, t)), \qquad (5)$$

where the mean μ_{θ} and the covariance matrix Σ_{θ} are learned and predicted by a θ -parameterized network, and the covariance matrix can be fixed to a constant in practice [2]. Learning to predict $\tilde{\mu}$ amounts to predicting x_0 , or equivalently ϵ due to the conditioning on x_t in the posterior, leading to a learning objective of minimizing $E_{t,x_0 \sim q(x_0), \epsilon \sim \mathcal{N}(0,1)}[||\epsilon - \epsilon_{\theta}(x_t, t)||^2]$, where ϵ_{θ} is predicted by a neural network and the norm being a pixel-wise ℓ_2 norm.

2.2. Fuzzy-conditioned image diffusion

To condition the diffusion, the synthetic image x_{t-1} and the projected input image x_{t-1}^r (re-noised conditioning input) can be fused with a weight map m throughout the diffusion process. The projection x_{t-1}^r is sampled with Eq. (2) at t-1, and x_{t-1} is obtained from our modified diffusion process. Unlike

Algorithm I Our fuzzy-conditioned image diffusion
1: $x_T \sim \mathcal{N}(0, \mathbb{1})$
2: for $t = T,, 1$ do
3: for $j = 1,, J$ do
4: $\epsilon \sim \mathcal{N}(0, \mathbb{1}) \text{ if } t > 1 \text{ else } \epsilon = 0$
5: $x_{t-1}^r = \sqrt{\overline{\alpha}_{t-1}}x + \epsilon\sqrt{1-\overline{\alpha}_{t-1}}$
6: $\epsilon_2 \sim \mathcal{N}(0, \mathbb{1}) \text{ if } t > 1 \text{ else } \epsilon_2 = 0$
7: $x_{t-1} = \frac{1}{\sqrt{\overline{\alpha}_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \overline{\alpha}_t}} \right) \epsilon_\theta(x_t, t) + \tilde{\beta}_t * \epsilon_2$
8: $x_m = m \cdot x_{t-1}^r + (1-m) \cdot x_{t-1} - \sqrt{\overline{\alpha}_{t-1}} x$
9: $x_m = \sqrt{\overline{\alpha}_{t-1}}x + x_m/\sqrt{1 - 2 \cdot m + 2 \cdot m^2}$
10: if $j < J$ and $t > 1$ then
$x_t \sim \mathcal{N}(\sqrt{1 - \beta_t} x_{t-1}, \beta_t \mathbb{1})$
11: end if
12: end for
13: end for
14: return x_0

inpainting, our map is fuzzy rather than binary and the fusion step requires distribution matching such that each resulting pixel corresponds to the correct Markovian $q(x_t|x_{t-1})$ distribution. The next forward step to obtain x_t is preceded by the following transformation over x_{t-1}

$$x_m \coloneqq \sqrt{\overline{\alpha}_{t-1}}x + \frac{m \cdot x_{t-1}^r + (1-m) \cdot x_{t-1} - \sqrt{\overline{\alpha}_{t-1}}x}{\sqrt{1 - 2 \cdot m + 2 \cdot m^2}},$$
(6)

using a fuzzy weight map m that can be obtained from Sec. 2.3 or can be any map of same dimensions as x such as its values are in [0, 1]. x is the conditioning input image that becomes the ground-truth target when $m \rightarrow 1$ and gets ignored when $m \rightarrow 0$. The standardization coefficient $\sqrt{1-2\cdot m+2\cdot m^2}$ comes from the fact that the resulting variance at t-1 of x_m is given by $m^2 V\{x_{t-1}^r\} + (1-1)^{-1}$ $(m)^2 V\{x_{t-1}\}$ where both variances $V\{x_{t-1}^r\}$ and $V\{x_{t-1}\}$ are made equivalent by design, matching $(1 - \overline{\alpha}_t)$. We then obtain x_t by a forward step over x_m instead of over x_{t-1} . Lastly, we use the harmonization resampling technique presented in [17], to promote high-level content consistency. The harmonization resampling consists of incorporating backward steps J in the forward diffusion process, to bring information from future more refined steps back to earlier steps. This is achieved by resampling $x_t \sim \mathcal{N}(\sqrt{1-\beta_t}x_{t-1},\beta_t\mathbb{1})$ for a fixed number of steps, and restarting the diffusion from that iteration. We summarize our method in Algorithm 1.

2.3. Diffusion projection attention

We exploit a diffusion network to obtain an inference-time attention map that we describe in this section. We project an



Fig. 1. Top: Input with known mask, and generated image samples with uniform maps $m_i > m_j$ ($\forall i > j$) each multiplied by the inpainting mask. $m_5 = 1$ is standard inpainting, $m_0 = 0$ is an image synthesis, and in-between columns inpaint with fuzzy conditioning. We see that our approach enables varying the conditioning strength outside the mask, when a regular diffusion inpainting [17] collapses. Bottom: our diffusion synthesis with spatially uniform and gradually decreasing input conditioning.

image x onto the latent space of the diffusion model following Eq. (2) $\forall t \in PS$, where PS is a set of projection sampling iterations. We reconstruct the \hat{x}_t images with the diffusion model, and study the discrepancies between the input image and its reconstructions across all iterations. In particular, we consider the deviations in the discrepancies relative to standard discrepancies that can be expected from data within $q(x_0)$. The projection attention map we compute is given by

$$A(x) \coloneqq \frac{1}{N} \sum_{t \in PS} \mathcal{H}\left(\frac{\phi(x - \hat{x}_t) - \mu_t\{\mathcal{V}\}}{\sigma_t\{\mathcal{V}\}}\right), \qquad (7)$$

where x is the input image (with potential anomaly or degradation), \hat{x}_t is the full forward-diffusion reconstruction of x after t sampling iterations, $\phi(\cdot)$ denotes a discrepancy measure which we set to the ℓ_2 -norm in our experiments, \mathcal{H} is a truncation function that limits the normalized maps to the range $[\sigma, 6\sigma]$ such that $\mathcal{H}(x) = min(max(x, 1), 6)$, and $(\mu_t\{\mathcal{V}\}, \sigma_t\{\mathcal{V}\})$ are respectively defined as

$$\mu_t\{\mathcal{V}\} \coloneqq \frac{1}{|\mathcal{V}|} \sum_{z \in \mathcal{V}} \phi(z - \hat{z}_t),$$

$$\sigma_t\{\mathcal{V}\} \coloneqq \sqrt{\frac{1}{|\mathcal{V}|} \sum_{z \in \mathcal{V}} (\phi(z - \hat{z}_t) - \mu_t\{\mathcal{V}\})^2},$$
(8)

where \mathcal{V} is a validation set corresponding to the training data of the diffusion model, and \hat{z}_t is defined in an analogous way to \hat{x}_t above, with the same diffusion projection approach. All operations are executed following the diffusion model's formulation in a pixel-wise manner. For all experiments, *PS* is the set given by {300, 400, 500, 600} and the final map *m* is given by $m \coloneqq (1 - A(x))^2$ to promote anomaly removal and preserve regions that are within one standard deviation.

3. EXPERIMENTAL EVALUATION

3.1. Experimental setup

We experimentally evaluate each of our contributions separately, and our proposed application framework. In our application, we exploit fuzzy-conditioned image diffusion, combined with diffusion projection attention as a guiding map, for autonomous face correction. Given an input facial image, we autonomously detect anomalies and correct them using our weighted diffusion prior. All experiments use DDPM models trained on CelebA [24] data with 256×256 image resolution.

3.2. Fuzzy-conditioned image diffusion

We illustrate how fuzzy-conditioned diffusion enables the generation of novel images with varying degrees of condi-



Fig. 2. Top row: Expected value of $\phi(z - \hat{z}_t)$ estimated over \mathcal{V} , for increasing diffusion projection depths t. Middle and bottom rows: Respectively a degraded input x, normalized projection deviations $\phi(x - \hat{x}_t)$, and the anomaly map m(x) (*shown without scale adjustment for better visualization).

tioning. We show in the top of Fig. 1 a gradually decreasing degree of conditioning from left to right, coupled with an inpainting area (0 conditioning). Each column can be sampled practically *infinitely many times*. RePaint [17], which cannot support fuzzy conditioning, fails on non-binary maps due to diffusion probability distribution mismatch. The bottom of Fig. 1 shows synthesis examples with a uniform and gradually decreasing conditioning weight on the input.

3.3. Diffusion projection attention

The top row of Fig. 2 shows the mean maps at different diffusion projection depths. The maps are computed with $|\mathcal{V}| =$ 1000, which is empirically sufficient for stable convergence. The middle and bottom rows show two example images that contain anomalous pixel-dependent perturbations (we add a different offset on each pixel to reach a random threshold, within a randomly selected rectangle). For each example, we show the normalized projection deviations at different projection depths, as well as the overall map. We note how the map is able to autonomously detect the perturbations, without any prior knowledge of the test-time degradation model.

3.4. Diffusion fuzzy-guided diffusion

We apply a similar degradation algorithm as in the previous section, compute our conditioning map m(x), and use it as a guide for our fuzzy-conditioned diffusion. Sample results are



Fig. 3. From left to right: A degraded input image x, the restored output with 400 iterations of DDPM [2], our restored output and our human-interpretable fuzzy correction map m(x) (*shown with the same visualization as Fig. 2).

shown in Fig. 3, compared against DDPM projection [2] that is also a fully degradation-agnostic image correction. We can see that our results correct the degradations that overlap with the face (second row) while remaining more visually faithful to the input images. Furthermore, our final outputs are more human interpretable, as the guidance map illustrates the degree of hallucination per pixel. We lastly note that our method is conditioned on a facial image prior and, as such, is prone to treating non-facial occlusions as degradations to be removed (see the finger removal in the bottom row, blue frame).

4. CONCLUSION

We derive a novel fuzzy-conditioned image diffusion capable of generating random images with pixel-wise varying degrees of image conditioning. We also present a diffusion projection attention that creates an anomaly detection map. We use this map as fuzzy guidance to our conditioned diffusion and show how it can autonomously correct facial images. We note that our application is trained on faces, and a flexible learning for scale adjustment would be needed to extend to other domains. Future work can also explore a frequency/latent-domain version of our method, for non-spatially-defined conditioning. **Acknowledgement:** Align Technology, Ringier, TX Group, NZZ, SRG, VSM, Viscom, and the ETH Zurich Foundation.

5. REFERENCES

- J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *International Conference* on Machine Learning (ICML), 2015.
- [2] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," in Advances in Neural Information Processing Systems (NeurIPS), 2020.
- [3] K. Pandey, A. Mukherjee, P. Rai, and A. Kumar, "VAEs meet diffusion models: Efficient and high-fidelity generation," in Advances in Neural Information Processing Systems (NeurIPS) Workshop, 2021.
- [4] K. Pandey, A. Mukherjee, P. Rai, and A. Kumar, "DiffuseVAE: Efficient, controllable and high-fidelity generation from low-dimensional latents," *arXiv preprint arXiv:2201.00308*, 2022.
- [5] M. Daniels, T. Maunu, and P. Hand, "Score-based generative neural networks for large-scale optimal transport," in Advances in Neural Information Processing Systems (NeurIPS), 2021.
- [6] Y. Benny and L. Wolf, "Dynamic dual-output diffusion models," in *IEEE/CVF Conference on Computer Vision* and Pattern Recognition (CVPR), 2022.
- [7] V. Sehwag, C. Hazirbas, A. Gordo, F. Ozgenel, and C. Canton, "Generating high fidelity data from lowdensity regions using diffusion models," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [8] F.-A. Croitoru, V. Hondru, R. T. Ionescu, and M. Shah, "Diffusion models in vision: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence* (*TPAMI*), 2022.
- [9] H. Chung, J. Kim, M. T. Mccann, M. L. Klasky, and J. C. Ye, "Diffusion posterior sampling for general noisy inverse problems," *arXiv preprint arXiv:2209.14687*, 2022.
- [10] M. El Helou and S. Süsstrunk, "BIGPrior: Toward decoupling learned prior hallucination and data fidelity in image restoration," *IEEE Transactions on Image Processing (TIP)*, 2022.
- [11] H. Sasaki, C. G. Willcocks, and T. P. Breckon, "UNIT-DDPM: Unpaired image translation with denoising diffusion probabilistic models," *arXiv preprint arXiv:2104.05358*, 2021.
- [12] C. Saharia, W. Chan, H. Chang, C. Lee, J. Ho, T. Salimans, D. Fleet, and M. Norouzi, "Palette: Image-toimage diffusion models," in ACM SIGGRAPH, 2022.

- [13] C. Saharia, J. Ho, W. Chan, T. Salimans, D. Fleet, and M. Norouzi, "Image super-resolution via iterative refinement," *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2022.
- [14] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, B. K. Ayan, S. S. Mahdavi, R. G. Lopes, et al., "Photorealistic text-toimage diffusion models with deep language understanding," *arXiv preprint arXiv:2205.11487*, 2022.
- [15] D. Horita, J. Yang, D. Chen, Y. Koyama, and K. Aizawa, "A structure-guided diffusion model for large-hole diverse image completion," *arXiv preprint arXiv*:2211.10437, 2022.
- [16] C. Meng, Y. He, Y. Song, J. Song, J. Wu, J.-Y. Zhu, and S. Ermon, "SDEdit: Guided image synthesis and editing with stochastic differential equations," in *International Conference on Learning Representations (ICLR)*, 2021.
- [17] A. Lugmayr, M. Danelljan, A. Romero, F. Yu, R. Timofte, and L. Van Gool, "RePaint: Inpainting using denoising diffusion probabilistic models," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [18] M. El Helou and S. Süsstrunk, "Blind universal Bayesian image denoising with Gaussian noise level learning," *IEEE Transactions on Image Processing* (*TIP*), 2020.
- [19] M. El Helou, *Deep Image Restoration: Between Data Fidelity and Learned Priors*, Ph.D. thesis, EPFL, 2021.
- [20] L. Liu, Y. Ren, X. Cheng, and Z. Zhao, "Diffusion denoising process for perceptron bias in out-of-distribution detection," arXiv preprint arXiv:2211.11255, 2022.
- [21] J. Wyatt, A. Leach, S. M. Schmon, and C. G. Willcocks, "AnoDDPM: Anomaly detection with denoising diffusion probabilistic models using simplex noise," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, 2022.
- [22] M. S. Graham, W. H. Pinaya, P.-D. Tudosiu, P. Nachev, S. Ourselin, and M. J. Cardoso, "Denoising diffusion models for out-of-distribution detection," *arXiv preprint arXiv*:2211.07740, 2022.
- [23] D. Gong, L. Liu, V. Le, B. Saha, M. R. Mansour, S. Venkatesh, and A. Hengel, "Memorizing normality to detect anomaly: Memory-augmented deep autoencoder for unsupervised anomaly detection," in *International Conference on Computer Vision (ICCV)*, 2019.
- [24] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *IEEE International Conference on Computer Vision (ICCV)*, 2015.