

Sinkhorn-based Quantization for Reversed Disease Progress: Further Investigations

Chih-Chieh Chen¹

JACKFRANK@GMAIL.COM

Chang-Fu Kuo^{1,2,3}

ZANDIS@GMAIL.COM

¹ *Center for Artificial Intelligence in Medicine, Chang Gung Memorial Hospital, Taoyuan, Taiwan*

² *Medical Education Department, Chang Gung Memorial Hospital, Taoyuan, Taiwan*

³ *Division of Rheumatology, Allergy and Immunology, Chang Gung Memorial Hospital, Taoyuan, Taiwan*

Editors: Under Review for MIDL 2026

Abstract

Interpretability plays a pivotal role in the collaboration between artificial intelligence (AI) systems and clinicians. It enables clinicians to critically reassess the rationale underlying AI-generated predictions. Moreover, translating these interpretations into clinically meaningful quantifications is feasible even for more granular algorithms, thereby potentially reducing the extensive annotation efforts typically required. Recently, a novel approach was introduced to generate reversed disease progression trajectories by applying conditional flow matching within the latent space of an autoencoder, jointly training a linear classifier. However, the architectural design, training procedures, and objective functions associated with the flow matching network warrant further investigation and refinement. In the present study, we implement this concept utilizing a recently proposed vector-quantized autoencoder framework incorporating Sinkhorn-based quantization. Our findings indicate that reversed disease progression can be consistently generated even in the absence of joint classifier training. Additionally, the method preserves strong spatial correspondences between the pixel domain and latent representations, enabling the synthesis of desired images through a CutMix-inspired algorithm. We demonstrate the efficacy of our approach by applying it to the weakly supervised quantization of midline shift distances.

Keywords: Interpretable AI, Medical Image, Clinical Diagnosis.

1. Introduction

The clinical context presents several distinctive challenges. Firstly, there exists an extensive volume of data, the majority of which is unlabeled and can only be derived from clinical reports. The generation of fine-grained annotations necessitates specialized domain knowledge, rendering such annotations largely unavailable in most cases. This issue is further exacerbated in the case of rare diseases, where each institution typically possesses only a few hundred annotated instances. Under these circumstances, training a robust classifier or generative model becomes particularly difficult (Karras et al., 2020; Kadkhodaie et al.). Secondly, the requirement for interpretability in clinical settings is generally more stringent than in other fields. Clinicians prefer model interpretations that are comprehensible to facilitate the reexamination of predictions. Moreover, it is advantageous to provide quantifiable information that can be integrated into various scoring systems, thereby alleviating the workload of clinical practitioners.

Recently, (Chih-Chieh and Chang-Fu, 2025) introduced a method for quantifying disease progression by generating a reversed trajectory, specifically transforming pathological images back to their normal counterparts, as depicted in Fig. 1. Their approach employs conditional flow matching within the latent space of an autoencoder, which is concurrently trained alongside a linear classifier. Although this concept demonstrates considerable potential, certain aspects require correction and warrant more comprehensive examination.

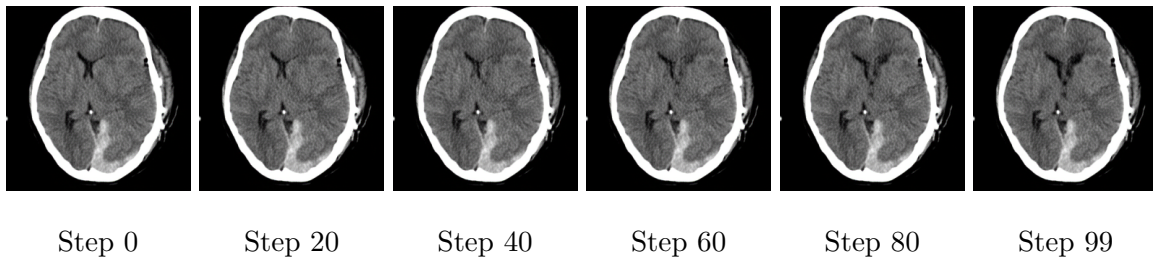


Figure 1: Reversed disease progression generated by our proposed architecture.

Empirical observations indicate that flow matching techniques can be readily applied to both the pixel space and the latent space of variational autoencoders (see Appendix E for more details). However, these methods fail to achieve successful training when applied to the standard VQGAN model (Esser et al., 2021). This limitation is likely attributable to the low codebook utilization rate inherent in VQGAN. Therefore, the authors of (Chih-Chieh and Chang-Fu, 2025) introduced an auxiliary classification head within the latent space of the autoencoder, which is trained jointly to distinguish between class-related and class-unrelated image tokens. For broader applicability, it would be advantageous to train the autoencoder and the classifier independently; specifically, one could train the autoencoder on a large dataset and subsequently train a classifier using images associated with specific target pathologies.

Moreover, the flow matching model faces challenges in regulating irrelevant features when such features are more prevalent in images from the source domain compared to those from the target domain. This issue is particularly common in clinical contexts. For instance, consider the example of midline shift (Broder, 2010). Typically, midline shift arises as a stress response to brain hematoma; however, not all patients with brain hematoma exhibit midline shifts. Moreover, it is generally easier to obtain images lacking both midline shift and brain hematoma than to acquire images without midline shift but with brain hematoma. Consequently, during the transfer of images to the target domain, these features may inadvertently be transferred alongside other characteristics.

In this study, we aim to examine these challenges and explore strategies to enhance performance. Our investigation builds upon the recently developed OptVQ framework (Zhang et al., 2024), which attains a high codebook utilization rate by relaxing the vector quantization (VQ) objective—specifically, by minimizing the distances between latent vectors and their nearest image tokens—through the application of a Sinkhorn-like algorithm. Notably, we observed that when employing OptVQ, pathological images can often be precisely transformed into normal images, even in cases exhibiting minimal midline shift distances. To address the issue of data imbalance, we discovered a pronounced spatial correspondence between pixel space and the latent space within vector-quantized autoencoders. In particular,

we found that applying a CutMix-inspired algorithm (Yun et al., 2019) within the latent space frequently yields realistic reconstructed images. Motivated by this observation, we propose a novel approach that involves substituting segments of the latent codes of target images with those from source images at locations characterized by high pixel intensity. We refer to this method as intensity shuffling.

We also aim to examine the underlying factors contributing to the enhanced performance observed when employing Sinkhorn-based quantization. Our analysis reveals that OptVQ exhibits a more pronounced spatial correlation between the pixel domain and the latent representation, as it effectively preserves local patches even when there is significant noise in adjacent regions. Motivated by this finding, we propose a method to transform images exhibiting midline shifts into normal images by substituting the cropped regions corresponding to the frontal horns with those from normal images. We refer to this approach as patch shuffling.

In summary, the contributions of this study are as follows:

- We elucidate the technical challenges associated with simulating the reverse progression of disease as introduced by (Chih-Chieh and Chang-Fu, 2025).
- To enhance the quality of generated image sequences, we apply the same conceptual framework to the recently proposed OptVQ model, achieving notable improvements.
- We conduct a detailed investigation into the role of Sinkhorn-based quantization methods and introduce two cost-effective data augmentation techniques—intensity shuffling and patch shuffling—designed to address data imbalance issues and facilitate the transformation of pathological images into normal ones, respectively.
- We evaluate the performance of our proposed methods using the midline shift dataset to demonstrate their efficacy.

2. Related Works

2.1. Basics for Midline Shift

In accordance with the methodology presented by Liao et al. (Liao et al., 2018), this subsection provides a concise overview of the concept of midline shift distance. The ideal midline, depicted as the blue line in Fig. 2 (a), is defined as the straight line connecting the anterior and posterior falx cerebri. Conversely, the actual midline, represented by the red curve, corresponds to the anatomical boundary that genuinely separates the left and right cerebral hemispheres. In a neurologically normal individual, these two lines are expected to be closely aligned. However, pathological conditions such as stroke, hematoma, or other brain injuries can induce increased intracranial pressure, resulting in deformation of the actual midline. The quantification of this intracranial pressure can be achieved by measuring the midline shift distance (MLS), which is defined as the displacement between the ideal and actual midlines at the level of the foramen of Monro (FM)—the anatomical channel connecting the frontal horns of the third ventricle.

(Liao et al., 2010) developed a landmark-based automatic midline shift measurement. Deep learning based methods are proposed by (Pisov et al., 2019; Wang et al., 2020; Wei

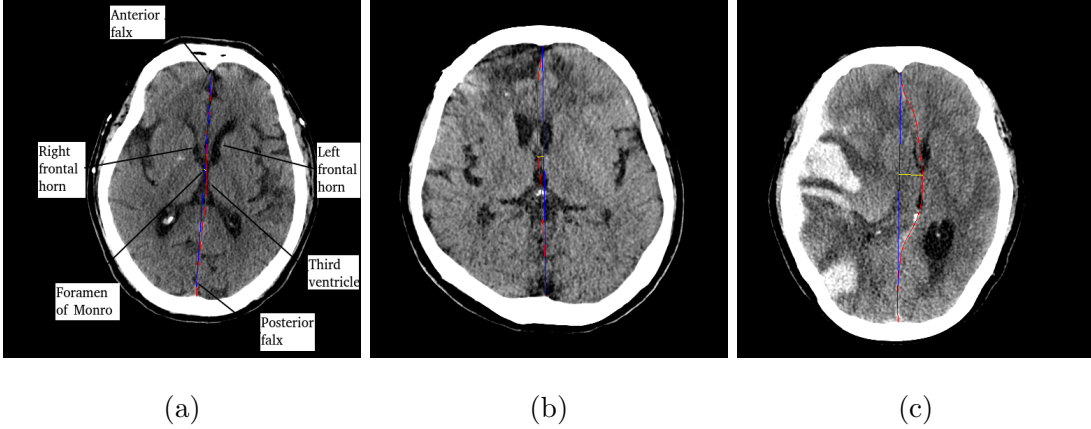


Figure 2: Illustrations of CT images (a) without midline shift, (b) with midline shift distance ≤ 5 mm (c) with midline shift distance > 5 mm.

et al., 2019; Nguyen et al., 2021) to handle the extreme cases which rule-based methods do not work well. (Pisov et al., 2019; Wang et al., 2020; Wei et al., 2019) transferred the midline shift measurement problem as the midline segmentation problem. (Nguyen et al., 2021) applied 2D-UNet and 3D-UNet to segment the anterior falx, posterior falx, and the foramen of Monro. Where the performances are promising, lots of annotations are required.

(Gong et al., 2023) proposed a semi-supervised learning approach: Instead of only outputting a single midline shift distance value, the authors attempted to predict a deformation field and predict the midline shift distance as the maximal deformation distance. The annotations for each image are the location where clinicians measure the midline shift distance, and the midline shift distance. The authors combined labeled and unlabeled images to train a class-conditional diffusion model. For labeled images, the authors first used classifier-free guidance (Ho and Salimans, 2021) to generate the image guided by the normal class label. The original and generated images are combined as the input of their designed deformation network to generate predicted the deformation field.

Our approach is similar to (Gong et al., 2023). Instead of implicitly using the generated images as input features to train a deformation field, we directly compared the original image and generated image, and generate the deformation field by image registration.

2.2. Vector Quantization

The study on the deep discrete representation is initiated in (Van Den Oord et al., 2017), the authors proposed VQ-VAE, an autoencoder with discrete latent representation. The authors proposed a trainable codebook $e \in \mathbf{R}^{K \times D}$, where K is the number of latent codes and D is the dimension of each code vector. For an input x , encoder output $z_e(x) \in \mathbf{R}^{L \times D}$, let e_x be the quantized vector of x , namely, for $i = 1, \dots, L$,

$$e_{x,i} = \arg \min_k \|z_e(x)_i - e_k\|_2, \quad (1)$$

In order to train $z_e(x)$, let $z_q(x)$ be the straight-through estimator (Bengio et al., 2013) of e_x with respect to $z_e(x)$, i.e., $z_q(x) = z_e(x) + \text{sg}[e_x - z_e(x)]$, where sg stands for the stop gradient operator. Then $z_q(x)$ is passed through the decoder for reconstructions and the training objective is

$$L = \log p(x|z_q(x)) + \| \text{sg}[z_e(x)] - e_x \|_2^2 + \| z_e(x) - \text{sg}[e_x] \|_2^2. \quad (2)$$

Vanilla Variational Autoencoders (VAEs) are known to suffer from the problem of posterior collapse (Wang et al., 2021), wherein the latent space becomes excessively noisy, causing the decoder or autoregressive model to neglect the fine-grained details encoded in the latent representations. On the other hand, the authors empirically demonstrate that the learned discrete latent codes retain meaningful information, and varying combinations of these tokens produce distinctly different outputs. To enhance the quality of image generation, (Esser et al., 2021) introduced VQGAN, which integrates the training objective defined in Equation 2 with additional discriminator and perceptual loss terms.

However, from Equation 2 we see that code vectors nearest to $z_e(x)$ are further adjusted to be even closer to it, which consequently increases the likelihood that these vectors will be selected in subsequent training iterations. This dynamic leads to a low utilization rate of the codebook e , thereby limiting the representational capacity of the model. To address this limitation, (Zhu et al., 2024) introduced VQGAN-LC. Rather than initializing the codebook randomly, (Zhu et al., 2024) constructed a fixed codebook by extracting patch-level features from datasets characterized by rich texture information (e.g., ImageNet or FFHQ) using a robust pretrained backbone network such as CLIP. Employing the same training objective as described in Equation 2, the authors then trained a projector to align the codebook with $z_e(x)$. (Zhu et al., 2025) introduced SimVQ, which reparameterizes the latent codebook through a learnable linear layer. Their findings demonstrate that this straightforward modification effectively mitigates the issue of codebook collapse.

(Zhang et al., 2024) proposed OptVQ, a method inspired by the optimal transport theory. Instead of replacing each embedding vector by the closest codebook vector, (Zhang et al., 2024) proposed to view each batch of latent embeddings $z_e(x)$ and the codebook e as two distributions of vectors, and transport $z_e(x)$ to e with minimal 2-Wasserstein distance, i.e., the distance in Equation 7. Since the marginal of π on e is exactly the probabilistic distribution of e , which means that every code vector of e will possibly be assigned by one of the embedding vector in $z_e(x)$, hence the utility rate of e is increased. Computationally, Sinkhorn algorithm (Peyré et al., 2019) is adopted for approximating the transport map: given a small number $\epsilon > 0$, initializing the assignment transport matrix C as the pairwise similarity matrix between pairs of vectors in $z_e(x)$ and e . Let $u = \mathbf{1} \in \mathbf{R}^K, v = \mathbf{1} \in \mathbf{R}^L$, where K, L are the sizes of $z_e(x), e$, respectively. Let $K = e^{-\frac{C}{\epsilon}}$, at each iteration, u, v are updated as the following way:

$$u \leftarrow \frac{z_e(x)}{Kv}, v \leftarrow \frac{e}{Ku}. \quad (3)$$

The assignment matrix will be given by $\text{diag}(u)K\text{diag}(v)$, where $\text{diag}(u), \text{diag}(v)$ are the diagonal matrices in $\mathbf{R}^{K \times K}, \mathbf{R}^{L \times L}$, respectively. From Remark 4.6 of (Peyré et al., 2019), $O(\frac{n^2 \log(n)}{\epsilon^3})$ iterations are enough to achieve ϵ -approximation, where $n = \max(K, L)$. In the




Figure 3: Illustrations of the effects of intensity shuffling on OptVQ with different numbers of patches.

work (Zhang et al., 2024), $\epsilon = 1$ and the authors only use 3 Sinkhorn iterations. However, empirically the authors demonstrated the codebook utility rates are already close to 100% in all the experiments. In our work, we adopt OptVQ as our autoencoder.

3. Method

3.1. Fair Reversed Disease Progress by Intensity Shuffling

As demonstrated in Section 2.1, images exhibiting midline shift distance are typically associated with stroke or hematoma. Nevertheless, in certain instances, the midline shift is not attributable to these conditions or is only partially caused by them. Conversely, empirical observations indicate that the majority of the collected normal CT scans do not present with stroke or hematoma. Given that our approach employs conditional flow matching to transform pathological images into normal images, it follows that most features related to hematoma are eliminated, including regions that do not contribute to the midline shift. For instance, in the source image of the first row depicted in Fig. 3, there are two hemorrhagic components; however, only the component adjacent to the frontal horns can be considered the causal factor for the midline shift. Despite this, we observed that the flow matching model tends to remove all hemorrhagic regions, likely because the training dataset contains very few normal images exhibiting hemorrhage.

To address this challenge, we observed that although the default configurations of VQ-GAN and OptVQ incorporate attention layers within the bottleneck layers of both the encoder and decoder, a pronounced spatial correspondence persists between the input image patches and the latent vectors at their respective locations. Consequently, as depicted in Fig. 3, we implemented a data augmentation technique whereby segments of patches corresponding to hemorrhage locations are directly shuffled during the training of the flow matching model. We refer to this augmentation approach as intensity shuffling.




Figure 4: Illustrations of the reconstructions of masked images. From above to below: original images, masks, reconstructed images by OptVQ, reconstructed images by VQGAN.

3.2. Masked Image Reconstruction as a Test of Spatial Correlation

Conceptually, Sinkhorn-based quantization differs from nearest neighbor-based quantization in that the final quantized value of each patch is influenced not only by the patch itself but also by its correlations with other patches within the same image. To empirically validate this hypothesis, and drawing inspiration from the works of (He et al., 2022; Li et al., 2022), we replaced 85% of the patches with the average pixel value of the original image and evaluated the reconstruction performance. As depicted in Fig. 4, the VQGAN model frequently produces fragmented patches under these conditions. In contrast, the OptVQ model reconstructs the patches at the unmasked locations with notable fidelity. Although similar reconstruction quality can be achieved using Variational Autoencoders (VAE) or other vector quantization techniques (refer to Fig. 11 and Fig. 14 for further details), we observed an intriguing phenomenon: for patches in masked regions, VQGAN and VAE tend to generate patches with darker coloration, potentially treating these patches as outliers. Conversely, OptVQ generates patches whose color closely approximates the average value of the original image. While the implications of this observation warrant further investigation, it suggests that OptVQ may possess superior capabilities in reconstructing unseen or outlier data. This insight motivated the introduction of patch shuffling, a CutMix-like data augmentation method, which is discussed in the subsequent subsection.

3.3. Patch Shuffling

As demonstrated in Fig. 4, our findings indicate that OptVQ is capable of reconstructing patches even when the neighboring patches are unrelated or untrained. In our study,

all images were rigidly registered, ensuring that the locations of the frontal horns remain proximate, even in images exhibiting midline shifts. This experiment, depicted in Fig. 4 prompted us to investigate the following question: if patches are cropped from the frontal horn regions of one latent code and subsequently pasted onto another latent code, will the resulting generated image be semantically and perceptually coherent? The results for OptVQ, presented in Fig. 5, reveal that the method reconstructs the frontal horns with near-perfect accuracy and produces minimal artifacts at the boundaries of the cropped regions. For comparative purposes, the corresponding results obtained using a VAE are shown in Fig. 10.

This form of data augmentation is referred to as patch shuffling. Notably, it is possible to directly produce plausible images without employing a flow matching model (refer to Appendix B for further details). In the present study, during the inference phase, patch shuffling is applied to the input and generated images to ensure that unrelated features remain constant throughout the reversed disease progression generation process.




Figure 5: Illustrations of the patch shuffling. From above to below: source images, target images, source images after patch shuffling with target images, target images after patch shuffling with source images.

4. Experiments

The training and testing datasets were from the Chang Gung Research Database (CGRD) (Tsai et al., 2017). For the training data, we used 1177 slices with midline shifts and 2221 normal slices from 907 Brain CT images. During the training of the flow matching network, the slices are further rigid registered to minimize rotational discrepancies involved in the reversed disease progression path. For testing data, we used 346 normal slices (referred to "normal" in Table 1), 178 slices with midline shift distances ranging from $2mm$ and $5mm$ (referred to "small MLS" in Table 1), and 119 slices with midline shift distances greater than $5mm$ (referred to "large MLS" in Table 1). To further quantify our results, we apply diffeomorphic registration to measure the displacement between the original and transformed images, which serves as an estimate of the midline shift distance of the original image. All the experiments were conducted with single Nvidia V100 GPU. For the flow matching model, we train for 20000 steps with a batch size of 80, using the Adam optimizer with a learning rate $2e - 4$.

During the training process, intensity shuffling is employed. Fig. 6 presents the generated images obtained with and without the application of patch shuffling during inference. Our findings indicate that the proposed models accurately restore the frontal horns to the optimal midline position, including cases exhibiting minimal midline displacement, as demonstrated in the third and fourth rows of Fig. 6. Furthermore, when patch shuffling is applied during inference, all unrelated regions remain unaltered throughout the reversal of disease progression.

The quantitative results are summarized in Table 1. Additionally, we employ the B-spline diffeomorphic registration technique implemented in AirLab (Sandkühler et al., 2018) to register the input images with the generated images. The resulting deformation fields are utilized to estimate the midline shift distance. It is noted that the mean absolute error of our method is relatively elevated for images exhibiting large midline shifts. Some of these errors may be attributed to unsuccessful image generation. However, we observe that numerous cases, particularly those involving hemorrhage, edema, or intensity variations, demonstrate inconsistencies between the deformation fields and the actual reversed progression of the disease. These instances are detailed in Appendix G. Further investigation into these cases is planned for future research.

Table 1: Mean absolute errors of unsupervised midline shift quantizations.

	Normal	Small MLS	Large MLS	Average (mm)
w/o patch shuffling	1.54	2.26	5.12	2.40
with patch shuffling	1.41	2.09	5.19	2.30

5. Conclusion

The generalization capabilities of flow matching models remain insufficiently understood and warrant further examination (Bertrand et al., 2025; Kamb and Ganguli, 2024; Biroli et al., 2024; Gao and Li, 2024; Scarvelis et al., 2023; Li et al., 2024). In the present study, we conduct a comprehensive analysis to determine the extent to which generalization arises from the autoencoder component itself. Specifically, we focus on the recently introduced OptVQ model, which employs Sinkhorn-based quantization to produce image tokens. Our findings




Figure 6: Illustrations of the generations of reverse disease progress. The first row and third row illustrate the flow matching model trained with the intensity shuffling, and the second and the fourth row represent the input images augmented by the generated images by using the patch shuffling during the inference time.

indicate that OptVQ is capable of reconstructing images with minimal artifacts when latent vectors are augmented using a CutMix-inspired augmentation technique. Motivated by this observation, we propose two novel augmentation strategies: intensity shuffling and patch shuffling. These approaches can be regarded purely as methods for generating novel images and have the potential to enhance the performance of conditional flow matching generation.

In this study, we employ intensity shuffling and patch shuffling techniques to enhance the fairness and stability of the generative process modeling the reversed disease progression. Given the limited number of images exhibiting hemorrhage in the collected dataset, intensity shuffling is utilized to augment normal images during the training of the flow matching model. Additionally, when reconstructing the reverse disease progression, it is preferable to maintain unrelated features unchanged; however, controlling the flow matching model to achieve this proves challenging, as minor alterations occur even in features that are entirely unrelated. To address this challenge, we augment the input images with generated images produced through patch shuffling applied to the regions corresponding to the frontal horns. Our findings indicate that the generated images exhibit minimal artifacts, and all unrelated features remain stable.

Although vector-quantized autoencoders are predominantly employed for autoregressive generation, the present study offers a novel perspective by demonstrating their suitability for image generation using latent codes enhanced through CutMix-like augmentation techniques. Future research will explore the potential applications of vector-quantized autoencoders within clinical contexts.

Acknowledgments

Both authors acknowledge funding from the Center for Artificial Intelligence in Medicine at Chang Gung Memorial Hospital, via grant agreement no. CLRPG3H0016.

References

- Michael Samuel Albergo and Eric Vanden-Eijnden. Building normalizing flows with stochastic interpolants. In *The Eleventh International Conference on Learning Representations*, 2023. URL <https://arxiv.org/abs/2209.15571>.
- Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.
- Quentin Bertrand, Anne Gagneux, Mathurin Massias, and Rémi Emonet. On the closed-form of flow matching: Generalization does not arise from target stochasticity. *arXiv preprint arXiv:2506.03719*, 2025.
- Giulio Biroli, Tony Bonnaire, Valentin De Bortoli, and Marc Mézard. Dynamical regimes of diffusion models. *Nature Communications*, 15(1):9957, 2024.
- Joshua Seth Broder. Head computed tomography interpretation in trauma: a primer. *Psychiatric Clinics*, 33(4):821–854, 2010.
- Chen Chih-Chieh and Kuo Chang-Fu. Weakly-supervised midline shift quantification through simulating the reversed disease progression. MIDL 2025 short paper, 2025.
- Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12873–12883, 2021.
- Ruiqi Gao, Emiel Hoogetboom, Jonathan Heek, Valentin De Bortoli, Kevin P. Murphy, and Tim Salimans. Diffusion meets flow matching: Two sides of the same coin. 2024. URL <https://diffusionflow.github.io/>.
- Weiguo Gao and Ming Li. How do flow matching models memorize and generalize in sample data subspaces? *arXiv preprint arXiv:2410.23594*, 2024.
- Shizhan Gong, Cheng Chen, Yuqi Gong, Nga Yan Chan, Wenao Ma, Calvin Hoi-Kwan Mak, Jill Abrigo, and Qi Dou. Diffusion model based semi-supervised learning on brain hemorrhage images for efficient midline shift quantification. In *International Conference on Information Processing in Medical Imaging*, pages 69–81. Springer, 2023.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022.

- Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on Deep Generative Models and Downstream Applications*, 2021.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Zahra Kadkhodaie, Florentin Guth, Eero P Simoncelli, and Stéphane Mallat. Generalization in diffusion models arises from geometry-adaptive harmonic representations. In *The Twelfth International Conference on Learning Representations*.
- Mason Kamb and Surya Ganguli. An analytic theory of creativity in convolutional diffusion models. *arXiv preprint arXiv:2412.20292*, 2024.
- Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training generative adversarial networks with limited data. *Advances in neural information processing systems*, 33:12104–12114, 2020.
- Sixu Li, Shi Chen, and Qin Li. A good score does not lead to a good generative model. *arXiv preprint arXiv:2401.04856*, 2024.
- Siyuan Li, Di Wu, Fang Wu, Zelin Zang, Stan Li, et al. Architecture-agnostic masked image modeling—from vit back to cnn. *arXiv preprint arXiv:2205.13943*, 2022.
- Chun-Chih Liao, Furen Xiao, Jau-Min Wong, and I-Jen Chiang. Automatic recognition of midline shift on brain CT images. *Computers in biology and medicine*, 40(3):331–339, 2010.
- Chun-Chih Liao, Ya-Fang Chen, and Furen Xiao. Brain midline shift measurement and its automation: a review of techniques and algorithms. *International journal of biomedical imaging*, 2018(1):4303161, 2018.
- Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.
- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations*, 2023.
- Nguyen P Nguyen, Youngjin Yoo, Andrei Chekkoury, Eva Eibenberger, Thomas J Re, Jyotipriya Das, Abishek Balachandran, Yvonne W Lui, Pina C Sanelli, Thomas J Schroepel, et al. Brain midline shift detection and quantification by a cascaded deep network pipeline on non-contrast computed tomography scans. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 487–495, 2021.
- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.

- Maxim Pisov, Mikhail Goncharov, Nadezhda Kurochkina, Sergey Morozov, Victor Gombolevskiy, Valeria Chernina, Anton Vladzmyrskyy, Ksenia Zamyatina, Anna Chesnokova, Igor Pronin, et al. Incorporating task-specific structural knowledge into cnns for brain midline shift detection. In *International Workshop on Multimodal Learning for Clinical Decision Support*, pages 30–38. Springer, 2019.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022.
- Robin Sandkühler, Christoph Jud, Simon Andermatt, and Philippe C Cattin. Airlab: autograd image registration laboratory. *arXiv preprint arXiv:1806.09907*, 2018.
- Christopher Scarvelis, Haitz Sáez de Ocáriz Borde, and Justin Solomon. Closed-form diffusion models. *arXiv preprint arXiv:2310.12395*, 2023.
- Alexander Tong, Kilian FATRAS, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrod Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2023.
- Ming-Shao Tsai, Meng-Hung Lin, Chuan-Pin Lee, Yao-Hsu Yang, Wen-Cheng Chen, Geng-He Chang, Yao-Te Tsai, Pau-Chung Chen, and Ying-Huang Tsai. Chang Gung Research Database: A multi-institutional database consisting of original medical records. *biomedical journal*, 40(5):263–269, 2017.
- Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. *Advances in neural information processing systems*, 30, 2017.
- Shen Wang, Kongming Liang, Yiming Li, Yizhou Yu, and Yizhou Wang. Context-aware refinement network incorporating structural connectivity prior for brain midline delineation. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 208–217. Springer, 2020.
- Yixin Wang, David Blei, and John P Cunningham. Posterior collapse and latent variable non-identifiability. *Advances in neural information processing systems*, 34:5443–5455, 2021.
- Hao Wei, Xiangyu Tang, Mingqing Zhang, Qingfeng Li, Xiaodan Xing, Xiang Sean Zhou, Zhong Xue, Wenzhen Zhu, Zailiang Chen, and Feng Shi. Regression-based line detection network for delineation of largely deformed brain midline. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 839–847. Springer, 2019.
- Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6023–6032, 2019.

Borui Zhang, Wenzhao Zheng, Jie Zhou, and Jiwen Lu. Preventing local pitfalls in vector quantization via optimal transport. *arXiv preprint arXiv:2412.15195*, 2024.

Lei Zhu, Fangyun Wei, Yanye Lu, and Dong Chen. Scaling the codebook size of vq-gan to 100,000 with a utilization rate of 99%. *Advances in Neural Information Processing Systems*, 37:12612–12635, 2024.

Yongxin Zhu, Bocheng Li, Yifei Xin, Zhihua Xia, and Linli Xu. Addressing representation collapse in vector quantized models with one linear layer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22968–22977, 2025.

Appendix A. Pseudocode for the Intensity Shuffling

The pseudocode for the intensity shuffling is presented in Algorithm 1.

Algorithm 1: Pseudocode for Intensity Shuffling

Input: Image pair x_0, x_1 . Encoder Enc . Pixel intensity threshold t_{min}, t_{max} . Number of pixels for shuffling n . Indicator function for the top- k values of tensors in the spatial dimensions `Select_Mask`.

Output: Augmented latent vector x_1^e

```

 $x_0^e \leftarrow Enc(x_0)$ ,  $x_1^e \leftarrow Enc(x_1)$ 
 $h \leftarrow x_0^e.size()[-2]$ ,  $w \leftarrow x_0^e.size()[-1]$ 
 $x_0 \leftarrow torch.mean(x_0, dim = 1).unsqueeze(1)$ 
 $x_0 \leftarrow torch.nn.functional.adaptive_avg_pool2d(x_0.float(), (h, w))$ 
 $x_0 \leftarrow torch.where(x_0 > t_{min}, 1, 0) \cdot torch.where(x_0 < t_{max}, 1, 0)$ 
 $x_0 \leftarrow Select\_Mask(x_0, n)$ 
 $x_1^e \leftarrow x_0 \cdot x_0^e + (1 - x_0) \cdot x_1^e$ 
return  $x_1^e$ 

```

Appendix B. Pseudocode for the Patch Shuffling

Algorithm 2 outlines the pseudocode for the patch shuffling procedure. Notably, the application of Algorithm 2 alone can occasionally produce images that are semantically coherent, as demonstrated in the first two columns of Fig. 7. However, for images exhibiting substantial midline shift distances, minor semantic inaccuracies tend to occur, as observed in the last four columns of Fig. 7.

Appendix C. Condition Flow Matching

Diffusion models (Ho et al., 2020; Dhariwal and Nichol, 2021) have attained state-of-the-art performance in image generation tasks. Nonetheless, these models typically assume that the source distributions follow a Gaussian form. Flow matching (Lipman et al., 2023; Liu et al., 2023; Albergo and Vanden-Eijnden, 2023; Tong et al., 2023) be interpreted as a generalization of diffusion models to accommodate arbitrary source distributions (Gao et al.,

Algorithm 2: Pseudocode for the Patch Shuffling

Input: latent code pair x_0, x_1 . Crop coordinates $h_{min}, h_{max}, w_{min}, w_{max}$.

Output: Augmented latent vector x_0^f

$M \leftarrow \text{torch.zeros_like}(x_0)$, $M[h_{min} : h_{max}, w_{min} : w_{max}] = 1$

$x_0^m \leftarrow M \cdot x_0$, $x_1^m \leftarrow M \cdot x_1$

$x_0^m \leftarrow \text{torch.flatten}(x_0^m, \text{start_dim} = 1)$, $x_1^m \leftarrow \text{torch.flatten}(x_1^m, \text{start_dim} = 1)$

$d \leftarrow \text{torch.cdist}(x_0^m, x_1^m)$

$d \leftarrow \text{torch.argmax}(d, \text{dim} = 1)$

$x_1 \leftarrow x_1[d]$

$x_0^f \leftarrow M \cdot x_0 + (1 - M) \cdot x_1$

return x_0^f




Figure 7: Illustrations of the generations of reverse disease progress by purely using the patch shuffling.

2024). (Tong et al., 2023) introduced conditional flow matching and reformulated (Lipman et al., 2023; Liu et al., 2023; Albergo and Vanden-Eijnden, 2023) within this framework. This subsection provides a concise overview of the conditional flow matching formulation.

Given a probabilistic density $p_0(x) \in \mathbb{R}^d$, we would like to construct a continuous family of probabilistic distributions $p_t(x), t \in [0, 1]$ from the solution of the following differential equation

$$dx = u(t, x)dt, \quad u(t, x) \in C^1([0, 1] \times \mathbb{R}^d, \mathbb{R}^d). \quad (4)$$

That is, let $\phi_t(x)$ be the solution of equation 4, and , we will construct p_t as the push-forward measure $p_t := [\phi_t]_{\#}(p_0)$. The probability path $p : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}$ such that $p(t, x) = p_t(x)$ is characterized as the famous continuity equation

$$\frac{dp}{dt} = -\nabla \cdot (p_t u_t), \quad (5)$$

where $u_t(x) := u(t, x)$, with initial condition p_0 .

Given a sampler on the joint distribution of source and target distributions $q(z)$, a conditional probabilistic distribution $p_0(x|z)$, and a computable $u_t(x|z)$, the conditional flow matching objective is defined as

$$\mathcal{L}_{CFM}(\theta) = \mathbb{E}_{t, q(z), p_t(x|z)} \|v_\theta(t, x) - u_t(x|z)\|^2, \quad (6)$$

where $p_t(x|z) := [\phi_t]_{\#}(p_0)$, ϕ_t is the solution of differential equation $dx = u(t, x|z)dt$, and v_θ is the neural network to be trained. The reformulations of (Lipman et al., 2023; Liu et al., 2023; Albergo and Vanden-Eijnden, 2023) are presented in (Tong et al., 2023) Table 1. For example, for (Liu et al., 2023), $q(z)$ is the independent coupling $q(x_0)q(x_1)$, $u(t, x|z) = x_1 - x_0$, and $p_t(x|z) = t \cdot x_1 + (1 - t) \cdot x_0$.

Let Π be the set of probabilistic distributions whose marginals are $q_0 := q(x_0)$ and $q_1 := q(x_1)$, and let 2-Wasserstein distance between distributions $q(x_0)$ and $q(x_1)$ with respect to Euclidean metric defined as the following:

$$W_2(q_0, q_1) := \inf_{\pi' \in \Pi} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|x - y\|^2 d\pi'(x, y), \quad (7)$$

and let $\pi(x_0, x_1) \in \Pi$ be the distribution achieving the optimal in Equation 7. (Tong et al., 2023) proposed to use $q(z) := \pi(x_0, x_1)$, $u(t, x|z) = x_1 - x_0$, and $p_t(x|z) = \mathcal{N}(x|t \cdot x_1 + (1 - t) \cdot x_0, \sigma^2)$. (Tong et al., 2023) call this method optimal transport CFM (OT-CFM) and faster training and inference compared with other baselines are demonstrated.

Appendix D. Feature Interpolation

In this study, we explore the application of conditional flow matching to transform pathological images into normal counterparts and to quantify the differences between these image pairs. Conditional flow matching assumes that the final trajectory comprises points obtained through linear interpolation between data points. To validate the appropriateness of employing this method within the latent space of a vector quantized autoencoder, we perform interpolations between two distinct images in the pixel space and in the latent




Figure 8: Illustrations of feature interpolations of different methods. From above to below: image interpolation, VAE, VQGAN, OptVQ.

spaces of VAE, VQGAN, and OptVQ, respectively. We select one pathological image x_0 and one normal image x_1 , and visualize their interpolations defined by $x_t = (1 - t)x_0 + tx_1$ on $t = 0.2, 0.4, 0.6, 0.8$. The resulting interpolated images are presented in Fig. 8.

In the direct interpolation of two images, as illustrated in the first row of Fig. 8, numerous artifacts are evident, particularly at $t = 0.4$, the ventricles from the two images appear conflated, with noticeable artifacts surrounding the skull region, and the patches near the left frontotemporal area do not resemble those originating from a single coherent image. These issues are mitigated when interpolation is performed within the latent space of VAE. At $t = 0.4$, two right ventricles are present in the same image, and the VAE attempts to interpret features from both the left and right ventricles of the source image as a single left ventricle in the interpolated output. Additionally, no artifacts are observed around the skull in this case. VQGAN demonstrates remarkable interpolation performance, even in the absence of conditional flow matching. At every interpolation step, the generated ventricles appear realistic, with minimal artifacts. However, possibly due to limitations related to codebook utilization, the reconstruction capability is constrained, resulting in the omission of some fine-grained textures. Lastly, the OptVQ method seeks to balance these factors; although non-realistic reconstructions of ventricles persist at $t = 0.4$, the merging of the two ventricles from the source image into a single (albeit unrealistic) right ventricle is somewhat improved compared to the VAE-generated images. Furthermore, the reconstruction quality achieved by OptVQ surpasses that of VQGAN and is comparable to that of the VAE.




Figure 9: Illustrations of the effects of intensity shuffling on VAE with different numbers of patches.

Appendix E. Performances on the Variational Autoencoder

VAEs are frequently employed as the autoencoder architecture in subsequent diffusion or flow matching generative models (Rombach et al., 2022). In the present study, we train a VAE with a latent representation of dimensions $16 \times 16 \times 16$, followed by training a flow matching model on this latent space. As demonstrated in Fig. 12, the reversed disease progress can be generated in a smooth manner. Nevertheless, we observe that when using VAEs, CutMix-like augmentations can introduce artifacts into the generated images.

Fig. 9 illustrates the impact of intensity shuffling. It is observed that the pixel values corresponding to the hemorrhage regions in the generated image are excessively elevated, while the surrounding edema is entirely omitted and fails to be reconstructed. When intensity shuffling is employed during the training of the flow matching model, as demonstrated in Fig. 13, a comparison with Fig. 12 reveals the emergence of numerous artifacts surrounding the brain skulls.

In the context of patch shuffling, as demonstrated in Fig. 10, while the reconstruction of the frontal horns is nearly flawless, the patches within the cropped areas exhibit inconsistencies when compared to those adjacent to the cropped regions. Notably, shadows and patches of high intensity are also produced.

Appendix F. Performances with different vector quantizations

We also conducted the experiment described in Section 3.2 using SimVQ (Zhu et al., 2025), a method that enhances the performance and stability of VQGAN by incorporating a linear learnable layer into the latent codebook. The results are illustrated in Fig. 14. When compared to Fig. 11, we find the reconstructed images produced by SimVQ and VAE exhibit a high degree of similarity. Notably, while both SimVQ and VAE tend to reconstruct darker patches within the masked regions, OptVQ generates patches with pixel values that approximate the average intensities of the original images.




Figure 10: Illustrations of the patch shuffling on VAE. From up to below: source images, target images, source images after the patch shuffling with target images, target images after the patch shuffling with source images.




Figure 11: Illustrations of the reconstructions of masked images. From up to below: original images, masks, reconstructed images by VAE.




Figure 12: Illustrations of the generations of reverse disease progress on the latent space of VAE.




Figure 13: Illustrations of the generations of reverse disease progress trained with intensity shuffling augmented data on the latent space of VAE.




Figure 14: Illustrations of the reconstructions of masked images. From up to below: original images, masks, reconstructed images with SimVQ.

Appendix G. Failure Case Study

Fig. 15 depicts the most illustrative examples of unsuccessful cases. In the first two columns, although the flow matching model successfully reconstructs plausible normal images, the performance of the diffeomorphic registration is significantly compromised by the presence of edema in the right hemisphere, whose pixel intensity closely resembles that of the frontal horns. Consequently, the edema is incorrectly deformed to form the right frontal horn. In the third and fourth columns, the frontal horns in the source images are notably small, leading to their deformation toward adjacent pixels rather than the intended final ventricular structures.

In the fifth column, the left frontal horn of the source image is positioned near the septum pellucidum in the generated image, resulting in its deformation toward the septum pellucidum rather than toward the left frontal horn of the generated image. In the final column, it is likely that the substantial spatial separation between the frontal horns of the source and generated images prevents the frontal horns of the source image from being accurately warped to the corresponding regions in the generated image.




Figure 15: Illustrations of failure cases. From above to below: original images, reconstructed images, warped images.