

Ordinal Diffusion Models for Color Fundus Images

Gustav Schmidt, Philipp Berens, and Sarah Müller

Hertie Institute for AI in Brain Health, University of Tübingen, Germany
`sar.mueller@uni-tuebingen.de`

Abstract. Generative image models such as diffusion models can improve performance on clinically relevant tasks by offering deep learning models supplementary training data. However, most conditional diffusion models treat disease stages as independent classes, ignoring the continuous nature of disease progression. This mismatch is problematic in medical imaging because continuous pathological processes are typically only observed through coarse, discrete but ordered labels as in ophthalmology for diabetic retinopathy (DR). We propose an ordinal latent diffusion model for generating color fundus images that explicitly incorporates the ordered structure of DR severity into the generation process. Instead of categorical conditioning, we used a scalar disease representation, enabling a smooth transition between adjacent stages. We evaluated our approach using visual realism metrics and classification-based clinical consistency analysis on the EyePACS dataset. Compared to a standard conditional diffusion model, our model reduced the Fréchet inception distance for four of the five DR stages and increased the quadratic weighted κ from 0.79 to 0.87. Furthermore, interpolation experiments showed that the model captured a continuous spectrum of disease progression learned from ordered, coarse class labels¹.

Keywords: Diffusion Models · Ordinal Labels · Disease Progression.

1 Introduction

Diabetic retinopathy (DR) is a leading cause of preventable vision loss worldwide, and it is routinely monitored using retinal color fundus photography (CFP). These images allow a direct, noninvasive view of the microvasculature and neural tissue; this is why fundus imaging is important in large-scale screening programs [20]. DR progression is continuous, with small anatomical and functional changes accumulating over time. However, clinicians generally follow a severity scale consisting of five ordinal stages: No DR (stage 0), mild non-proliferative DR (stage 1), moderate NPDR (stage 2), severe NPDR (stage 3) and proliferative DR (stage 4) (Fig. 1a). Although deep learning models have achieved excellent performance in DR detection in general [23], their development and application is restricted by access to sufficiently large diverse datasets, in particular with

¹ <https://github.com/berenslab/OrdinalDiffusionModels>




Fig. 1. Overview figure of the proposed approach. Our diffusion model generates fundus images with different disease severity (b) by conditioning on ordinal embeddings (a). Symbols indicate jointly trained (flame) and pretrained frozen (snowflake) components.

sufficient examples of late disease stages or from persons with underrepresented ethnicity [22]. Generative models may offer a solution by synthesizing realistic fundus images, which have been suggested to improve the performance of deep learning models in medical imaging providing additional data for training on underrepresented classes [16,1,4]. While many attempts at generative modeling for CFPs have mostly resorted to generative adversarial networks (GANs) [12,21,18], diffusion models have demonstrated superior stability, diversity and scalability [6,25,13]. For example, [13] used classifier-guidance for training a diffusion model for CFP generation, requiring robust classifiers to be trained in the first place. However, the ordinal disease labels could already provide sufficient information for generating clinically meaningful, realistic fundus images along the DR disease spectrum. In this work, we present an ordinally conditioned latent diffusion model for generating CFP images that explicitly models the continuous progression of DR (Fig. 1b). First, we explicitly encode the ordered structure of DR severity into the diffusion process, enabling smooth transitions between DR stages and modeling disease progression as a continuum rather than discrete categories. Second, we propose a dual-conditioning strategy that separates anatomical retinal structure from disease-specific pathology via a learned struc-

tural embedding, improving visual realism, disease consistency, and controllable interpolation across severity levels. Finally, we evaluate the model using both standard image quality metrics and classifier-based disease consistency analysis to verify that generated images preserve clinically meaningful severity ordering.

2 Related Work

Ordinal learning [27,31] has received comparably little attention in medical imaging. It has been studied in a discriminative setting to model ordered labels for disease severity grading [19,33,4]. These methods reduce clinically implausible errors by explicitly introducing ordinal constraints, leading to increased performance and interpretability [30,26]. Recently, a few others have developed generative models for medical data exploiting ordinal constraints [17,29]. For example, [29] suggest to employ geometric constraints on the difference of the noisy images to enforce ordinal ranks in latent space. Based on their evaluation, this approach did not consistently yield better performance in terms of realism than a simple diffusion model. We were unable to reproduce their findings because no code was provided, despite our best efforts to reimplement their approach. Similarly, [17] learn a high-dimensional look-up table to encode disease stages for endoscopic images based on a Stable Diffusion model. However, since no constraints are enforced on the look-up table entries, it is not clear how this approach can be considered ordinal at all. Thus, to our knowledge, our model is the first reproducible work that explicitly encodes ordinal geometric constraints into the conditioning representation of a latent diffusion model for disease staging.

3 Methods

3.1 Ordinal-Aware Latent Diffusion Model

Latent Autoencoding and Diffusion. We followed the standard latent diffusion framework [25]. An image $x \in \mathbb{R}^{H \times W \times 3}$ is mapped to a lower-dimensional latent using a Variational Autoencoder (VAE) with f as the downsampling factor:

$$E_\phi : \mathbb{R}^{H \times W \times 3} \rightarrow \mathbb{R}^{h \times w \times c}, \quad h = H/f, \quad w = W/f.$$

The encoder parametrizes a diagonal Gaussian posterior in latent space $E_\phi(x) = (\mu_\phi(x), \sigma_\phi(x))$, $q_\phi(z | x) = \mathcal{N}(z; \mu_\phi(x), \sigma_\phi^2(x)I)$, with $z_0 \sim q_\phi(z | x) = \mu_\phi(x) + \sigma_\phi(x) \odot \epsilon$, $\epsilon \sim \mathcal{N}(0, I)$, such that the diffusion model operates on stochastic latents. The VAE, consisting of encoder E_ϕ and decoder D_ψ , is trained via

$$\mathcal{L}_{\text{AE}} = \mathbb{E}_x \left[\|x - D_\psi(z_0)\|_1 + \lambda_{\text{perc}} \mathcal{L}_{\text{perc}}(x, D_\psi(z_0)) + \lambda_{\text{KL}} D_{\text{KL}}(q_\phi(z | x) \| \mathcal{N}(0, I)) \right],$$

where $\mathcal{L}_{\text{perc}}$ is a perceptual loss [14] and the KL term regularizes the latent space toward a standard normal. Forward diffusion is applied to z_0 over $t = 1, \dots, T$:

$$z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad \epsilon \sim \mathcal{N}(0, I),$$

with $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{k=1}^t \alpha_k$ and $\{\beta_t\}_{t=1}^T \subset (0, 1)$ as the noise schedule.

Conditioning the Diffusion Model on Disease Stage and CFP Structure. The denoiser ϵ_θ was conditioned on the timestep t , a representation of the disease grade c , and, optionally, the image structure s , and trained via

$$\mathcal{L}(\theta) = \mathbb{E}_{x_0, \epsilon, t} [\|\epsilon - \epsilon_\theta(z_t, t, c, s)\|_2^2], \quad t \sim \mathcal{U}\{1, \dots, T\},$$

with random conditioning dropout for classifier-free guidance [11]. As a baseline, we encoded the disease grade c with one-hot encoding; ordinal-aware representations are introduced below. To represent the structure of the image s independently of pathology, we followed [4] trained a structural encoder, $E_s : \mathbb{R}^{H \times W \times 3} \rightarrow \mathbb{R}^d$, $s = E_s(x)$, via contrastive learning, implemented as a ResNet-50 [9] trained with NT-Xent loss [3]. During sampling, we combined unconditional, disease-conditioned, and structure-conditioned predictions as

$$\epsilon_\theta(z_t, t, c, s) = \epsilon_\theta(z_t, t) + w_c(\epsilon_\theta(z_t, t, c) - \epsilon_\theta(z_t, t)) + w_s(\epsilon_\theta(z_t, t, s) - \epsilon_\theta(z_t, t)),$$

where w_c and w_s control disease and structural guidance strength.

Ordinal Disease Stage Conditioning. In addition to the standard categorical one-hot encoding for disease grade conditioning, we implemented two strategies that explicitly capture disease stage ordering:

1. **Equidistant margins:** disease stages are embedded on a 1D ordinal axis with equal spacing, $c_i = i$, $i \in \{0, \dots, K-1\}$.
2. **Learned margins:** relative spacings are learned via $c_i = \sum_{j=0}^i v_j^2$, $v \in \mathbb{R}^K$, where squaring enforces positive increments and guarantees monotonicity.

In all cases, c was passed through a two-layer fully connected MLP to produce a 1,024-dimensional embedding before being injected into the denoising network.

3.2 Evaluation of Image Realism and Disease Consistency

Visual realism was measured by the Fréchet inception distance (FID) score [10,24]. Because FID is highly sensitive to small sample sizes, class-specific FID comparisons using per-class reference pools would be unreliable. We therefore computed FID using a fixed reference pool of 10,000 randomly sampled real images from the full dataset, compared against 10,000 generated images from each DR grade. Disease consistency was assessed using a pretrained DR classifier (ResNet-50 with the CORAL ordinal regression loss [2]) on the generated images. We computed the quadratic weighted κ (QWK) between the predicted grades and the target stages used for generation. QWK penalizes disagreements according to their distance in the ordered disease-stage space.

3.3 Dataset and Preprocessing

We used the EyePACS dataset [5] (accessed 2022), larger than the public Kaggle subset. After quality filtering [8], 127,144 fundus images from 35,947 patients remained, with the following class distribution: Grade 0: 100,584; Grade 1: 9,948;

Grade 2: 14,823; Grade 3: 960; Grade 4: 829. Images were center-cropped and resized to 256×256 pixels. We used task-specific strong augmentations (random crops, rotations, color jitter) for the DR classifier and the structural encoder E_s , while anatomically preserving augmentations (small rotations, horizontal flips, mild color jitter) were used for diffusion training. Data were split by patient identifier to avoid leakage. For latent diffusion and autoencoder training, we used a 90/10 train/validation split. To train the DR classifier, data were split 80/10/10 into train/validation/test sets. Weighted random sampling mitigated class imbalance. The final DR classifier, used to assess disease consistency, achieved a QWK of 0.63 on the test set.

3.4 Training Details

All models were implemented in PyTorch [7] (see code repository for full configuration files) and trained on NVIDIA A100 GPUs. The autoencoder was trained for 100 epochs using Adam ($\text{lr} = 10^{-5}$, batch size 16, $\lambda_{\text{perc}} = 0.25$), and the structural encoder for 100 epochs using Adam ($\text{lr} = 10^{-4}$, batch size 256). Diffusion was performed in a $64 \times 64 \times 4$ ($f = 4$) latent space with a linear noise schedule over $T = 1,000$ steps ($\beta_{\text{start}} = 10^{-4}$, $\beta_{\text{end}} = 0.02$). The denoiser was optimized using AdamW ($\text{lr} = 10^{-5}$, batch size 16). Sampling employed DDIM [28] with 100 inference steps and classifier-free guidance for both conditioning signals. Guidance weights were set empirically: $(w_c, w_s) = (4.0, 0.0)$ for unconditional structural generation and $(7.0, 1.5)$ for structure-preserving synthesis.


4 Results

4.1 Ordinal Embedding Improved Realism and Disease Consistency

We evaluated six model variants: a baseline model with categorical one-hot encoding for disease stages and two ordinally conditioned models with equidistant disease stages or learned distance between disease stages, each trained with and without structural conditioning (Table 1). We found that for models *without structural conditioning* both ordinal variants showed more realism as measured by the FID score than the baseline model for all DR grades except severe DR, where the model with equidistant margins was only slightly better than the baseline and the model with learned margins was slightly worse. In addition, the model with equidistant margins resulted in a QWK of 0.85 compared to 0.79 for the baseline model, indicating that learned disease features were more consistent with those of real images, while learned margins performed similar to the baseline. Including *structural embeddings* into the model, we found again that the model with equidistant margins created more realistic images overall which were also more in line with disease features, except for proliferative DR, where apparently structural features interfered with disease features. The learned margin model did show much improved disease consistency, but the realism of the pictures was worse than without structural constraints. In the following analysis, we only considered model versions with structural conditioning.

Table 1. Quantitative evaluation: class-wise FID and QWK for all model variants.

Setting	FID ↓					QWK ↑
	No DR	Mild	Moderate	Severe	Prolif.	
<i>Baseline</i>						
w/o structural cond.	22	23	19	16	63	0.79
w/ structural cond.	23	22	23	18	22	0.82
<i>Equidistant Margins</i>						
w/o structural cond.	11	13	13	15	29	0.85
w/ structural cond.	12	12	12	16	38	0.87
<i>Learned Margins</i>						
w/o structural cond.	13	14	16	24	34	0.79
w/ structural cond.	15	17	21	42	54	0.85

**Fig. 2.** CFPs generated from the ordinal diffusion model with equidistant margins are realistic and capture disease related lesions.

4.2 Ordinal Diffusion Model Generates Images with Realistic Morphological Features and Disease Lesions

Across all classes, the model reproduced core retinal structures (Fig. 2). The optic disc had correct brightness, shape, and placement, and blood vessels had realistic branching and tapering. Disease-specific features increased with severity: “no DR” samples lacked lesions; “mild DR” showed sparse microaneurysms; moderate and severe stages had more hemorrhages and exudates; and proliferative DR featured dense, irregular vascular formations resembling neovascularization. Some proliferative samples contained artifacts (fragmented vessels or blurriness), consistent with the higher FID values and highlighting the challenge of modeling advanced pathology. The model showed no mode collapse, with each image presenting a unique vessel and lesion distribution, demonstrating realistic fundus generation across DR stages.

4.3 Structure Encoder Enables Counterfactual Image-to-Image Generation

To further investigate potential benefits of adding the structural conditioning to the model, we used the ordinal diffusion model for counterfactual image gen-




Fig. 3. The structure encoder enabled counterfactual image-to-image generation. Original CFP of healthy eye on the left, generated images with ascending disease severity to the right. Below: difference image for each generated image from the healthy image.

eration based on an input image. Starting from a healthy reference image, we generated fundus images with ascending DR severities (Fig. 3). The resulting images retained largely the same anatomical structures — vasculature, optic disc, and macula — while progressively exhibiting disease-specific lesions. Mean difference maps over the three color channels highlighted the added pathological features (Fig. 3, second row), demonstrating that the model could modify DR characteristics while maintaining consistent structural anatomy.

4.4 Ordinal Diffusion Model Captures Continuous Disease Structure

To examine how the model captured the continuous disease spectrum underlying the discrete DR stages, we interpolated across the ordinal spectrum and generated class-conditioned images on intermediate numbers between the disease stages (Fig. 4). Importantly, this interpolation is not possible with the categorical baseline, as one-hot encoding does not define a continuous conditioning space. For the equidistant-margin model, these values lay between adjacent discrete numbers encoding the disease stages and for the learned-margin model, they lay between the corresponding learned values used for class conditioning. We generated 256 images for each number and classified them using our pretrained DR classifier. We found that the model trained with equal margins showed a smooth progression between classes: as we interpolated from “No DR” toward “Mild”, proportion of images classified as healthy decreased while the “Mild” ratio increases, and this gradual transition continued across all stages. Intermediate images contained mixed features rather than sharp jumps, showing that the model learned a continuous spectrum of the disease rather than discrete categories. For the learned-margin model, interpolation revealed smaller gaps between early stages (“No DR” to “Moderate”) and larger gaps for advanced stages (“Severe” to “Proliferative”), while still maintaining smooth transitions.




Fig. 4. The ordinal diffusion model learns the continuous disease spectrum underlying ordinal labels. Images generated at intermediate class values for (a) the equidistant margin model and (b) the learned margins model were generated with a multi-class DR classifier. Resulting label proportions are shown.

The spacing reflects the model’s perception of disease severity: early transitions appeared more subtle, whereas later stages corresponded to more pronounced pathological changes such as neovascularization. Thus, while equidistant margins improved quantitative performance, learned margins provide additional insights into the clinical complexity of DR progression.

5 Discussion

Our model generates realistic interpolations between adjacent disease stages, indicating that it learns a continuous representation of disease severity consistent with the progressive nature of DR. Ordinal conditioning thus improves both disease-consistent generation and visual realism. Despite these strengths, several limitations remain. Learned margins make no assumptions about inter-stage distances, allowing the model to learn inter-stage geometry from data, but may suffer from training instability due to the additional learnable parameters and insufficient gradient signal from minority classes. While performing on par with the baseline, learned margins provided improved interpretability through clinically meaningful margins. Beyond architectural choices, class imbalance poses a persistent challenge: higher FID values were observed for minority classes, which were likely primarily driven by class imbalance, as the model receives substantially less gradient signal from these stages. While the model reproduces major vessel structures, finer vessel branching remains imperfect, particularly in advanced stages where neovascularization introduces complex topology. These generation artifacts point to a broader methodological challenge: evaluating medical im-

age generators remains difficult. FID measures realism but does not guarantee clinical correctness [32]. To address this, we evaluated a DR classifier on the generated images to assess disease consistency. Interestingly, the DR classifier performed better on generated images than on real data, suggesting that the generator may emphasize dominant class features or that the EyePACS dataset contains substantial label noise, as reported previously [15]. Furthermore, our evaluation mainly relied on automated metrics, in contrast to other recent work on generative diffusion models for CFPs [13]. Addressing these limitations motivates several directions for future work. Since ordinal conditioning is resolution-agnostic, higher-resolution synthesis is a natural next step to better capture fine pathological details. Future work should further include expert clinical assessment, quantitative evaluation of counterfactual images, and latent space analysis of the structural encoder. Finally, extending the framework to longitudinal datasets may enable explicit modeling of temporal disease progression.

Acknowledgments. This project was supported by the Hertie Foundation and by the Deutsche Forschungsgemeinschaft under Germany’s Excellence Strategy with the Excellence Cluster 2064 “Machine Learning — New Perspectives for Science”, project number 390727645. PB is a member of the Else Kröner Medical Scientist Kolleg “ClinbrAIn: Artificial Intelligence for Clinical Brain Research”. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting SM.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Akrouit, M., Gyepesi, B., Holló, P., Poór, A., Kincsó, B., Solis, S., Cirone, K., Kawahara, J., Slade, D., Abid, L., et al.: Diffusion-based data augmentation for skin disease classification: Impact across original medical datasets to fully synthetic images. In: International conference on medical image computing and computer-assisted intervention. pp. 99–109. Springer (2023)
2. Cao, W., Mirjalili, V., Raschka, S.: Rank consistent ordinal regression for neural networks with application to age estimation. *Pattern Recognition Letters* **140**, 325–331 (2020)
3. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PmLR (2020)
4. Cheng, Y., Ying, H., Hu, R., Wang, J., Zheng, W., Zhang, X., Chen, D., Wu, J.: Robust image ordinal regression with controllable image generation. *arXiv preprint arXiv:2305.04213* (2023)
5. Cuadros, J., Bresnick, G.: EyePACS: An Adaptable Telemedicine System for Diabetic Retinopathy Screening. *Journal of Diabetes Science and Technology* **3**(3), 509–516 (May 2009). <https://doi.org/10.1177/193229680900300315>
6. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. *Advances in neural information processing systems* **34**, 8780–8794 (2021)
7. Falcon, W., The PyTorch Lightning team: PyTorch Lightning (Mar 2019). <https://doi.org/10.5281/zenodo.3828935>

8. Gervelmeyer, J., Müller, S., Huang, Z., Berens, P.: Fundus Image Toolbox: A Python package for fundus image processing. *Journal of Open Source Software* **10**(108), 7101 (Apr 2025). <https://doi.org/10.21105/joss.07101>
9. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
10. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems* **30** (2017)
11. Ho, J., Salimans, T.: Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598* (2022)
12. Hou, B.: High-fidelity diabetic retina fundus image synthesis from freestyle lesion maps. *Biomedical Optics Express* **14**(2), 533–549 (2023)
13. Ilanchezian, I., Boreiko, V., Kühlewein, L., Huang, Z., Seçkin Ayhan, M., Hein, M., Koch, L., Berens, P.: Development and validation of an AI algorithm to generate realistic and meaningful counterfactuals for retinal imaging based on diffusion models. *PLOS Digital Health* **4**(5), e0000853 (2025)
14. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *European conference on computer vision*. pp. 694–711. Springer (2016)
15. Ju, L., Wang, X., Wang, L., Mahapatra, D., Zhao, X., Zhou, Q., Liu, T., Ge, Z.: Improving medical images classification with label noise using dual-uncertainty estimation. *IEEE transactions on medical imaging* **41**(6), 1533–1546 (2022)
16. Khosravi, B., Li, F., Dapamede, T., Rouzrokh, P., Gamble, C.U., Trivedi, H.M., Wyles, C.C., Sellergren, A.B., Purkayastha, S., Erickson, B.J., et al.: Synthetically enhanced: unveiling synthetic data’s potential in medical imaging research. *EBioMedicine* **104** (2024)
17. Kurt, M.M., Çağlar, Ü.M., Temizel, A.: Progressive Disease Image Generation with Ordinal-Aware Diffusion Models. *Diagnostics* **15**(20), 2558 (2025)
18. Lang, O., Gandelman, Y., Yarom, M., Wald, Y., Elidan, G., Hassidim, A., Freeman, W.T., Isola, P., Globerson, A., Irani, M., et al.: Explaining in style: training a GAN to explain a classifier in stylespace. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 693–702 (2021)
19. Liu, X., Zou, Y., Song, Y., Yang, C., You, J., K Vijaya Kumar, B.: Ordinal regression with neuron stick-breaking for medical diagnosis. In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops*. pp. 0–0 (2018)
20. London, A., Benhar, I., Schwartz, M.: The retina as a window to the brain—from eye research to cns disorders. *Nature Reviews Neurology* **9**(1), 44–53 (Nov 2012). <https://doi.org/10.1038/nrneuro1.2012.227>
21. Müller, S., Koch, L.M., Lensch, H.P., Berens, P.: Disentangling representations of retinal images with generative models. *Medical Image Analysis* p. 103628 (2025)
22. Nakayama, L.F., Zago Ribeiro, L., Novaes, F., Miyawaki, I.A., Miyawaki, A.E., de Oliveira, J.A.E., Oliveira, T., Malerbi, F.K., Regatieri, C.V.S., Celi, L.A., et al.: Artificial intelligence for telemedicine diabetic retinopathy screening: a review. *Annals of Medicine* **55**(2), 2258149 (2023)
23. Nielsen, K.B., Lautrup, M.L., Andersen, J.K., Savarimuthu, T.R., Grauslund, J.: Deep learning-based algorithms in screening of diabetic retinopathy: a systematic review of diagnostic performance. *Ophthalmology Retina* **3**(4), 294–304 (2019)
24. Parmar, G., Zhang, R., Zhu, J.Y.: On aliased resizing and surprising subtleties in gan evaluation. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 11410–11420 (2022)

25. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10684–10695 (2022)
26. Schmidt, G., Heidrich, H., Berens, P., Müller, S.: Learning Disease State from Noisy Ordinal Disease Progression Labels. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 284–293. Springer (2025)
27. Shin, N.H., Lee, S.H., Kim, C.S.: Moving window regression: A novel approach to ordinal regression. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 18760–18769 (2022)
28. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. arXiv preprint arXiv:2010.02502 (2020)
29. Takezaki, S., Uchida, S.: An ordinal diffusion model for generating medical images with different severity levels. In: 2024 IEEE International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2024)
30. Wang, J., Chen, J., Liu, J., Tang, D., Chen, D.Z., Wu, J.: A Survey on Ordinal Regression: Applications, Advances and Prospects. arXiv preprint arXiv:2503.00952 (2025)
31. Wang, J., Cheng, Y., Chen, J., Chen, T., Chen, D., Wu, J.: Ord2seq: Regarding ordinal regression as label sequence prediction. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5865–5875 (2023)
32. Wu, Y., Liu, F., Yilmaz, R., Konermann, H., Walter, P., Stegmaier, J.: A Pragmatic Note on Evaluating Generative Models with Fréchet Inception Distance for Retinal Image Synthesis. arXiv preprint arXiv:2502.17160 (2025)
33. Yu, Q., Xie, J., Nguyen, A., Zhao, H., Zhang, J., Fu, H., Zhao, Y., Zheng, Y., Meng, Y.: CLIP-DR: Textual knowledge-guided diabetic retinopathy grading with ranking-aware prompting. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 667–677. Springer (2024)