

Directed Ordinal Diffusion Regularization for Progression-Aware Diabetic Retinopathy Grading

Huangwei Chen^{1,2,3,†}, Junhao Jia^{1,3,†}, Ruocheng Li¹, Cunyuan Yang¹, Wu Li¹, Xiaotao Pang², Yifei Chen⁴, Haishuai Wang¹, Jiajun Bu^{1(✉)}, and Lei Wu^{1,2(✉)}

¹ Zhejiang Key Laboratory of Accessible Perception and Intelligent Systems, College of Computer Science and Technology, Zhejiang University

² Hangzhou Pujian Medical Technology Co., Ltd

³ Hangzhou Dianzi University ⁴ Tsinghua University
{bjj, shenhai1895}@zju.edu.cn

Abstract. Diabetic retinopathy (DR) progresses as a continuous and irreversible deterioration of the retina, following a well-defined clinical trajectory from mild to severe stages. However, most existing ordinal regression approaches model DR severity as a set of static, symmetric ranks, capturing relative order while ignoring the inherent unidirectional nature of disease progression. As a result, the learned feature representations may violate biological plausibility, allowing implausible proximity between non-consecutive stages or even reverse transitions. To bridge this gap, we propose Directed Ordinal Diffusion Regularization (D-ODR), which explicitly models the feature space as a directed flow by constructing a progression-constrained directed graph that encourages forward disease evolution. By performing multi-scale diffusion on this directed structure, D-ODR imposes penalties on score inversions along valid progression paths, thereby effectively preventing the model from learning biologically inconsistent reverse transitions. This mechanism aligns the feature representation with the natural trajectory of DR worsening. Extensive experiments demonstrate that D-ODR yields superior grading performance compared to state-of-the-art ordinal regression and DR-specific grading methods, offering a more clinically reliable assessment of disease severity. Our code is available at <https://github.com/HovChen/D-ODR>.

Keywords: Diabetic Retinopathy · Disease Progression Modeling · Directed Graph Learning

1 Introduction

Diabetic retinopathy (DR) is a leading cause of preventable blindness among working-age adults, and large-scale screening commonly relies on grading retinal fundus photographs into discrete severity categories [5,21,2,25]. This grading task is frequently formulated as ordinal prediction due to the inherent ordering of the labels [26]. In practice, severity grades represent disease states along a

[†] Equal contribution; ✉ corresponding author.

temporal trajectory. Mild instances of DR may remain stable or progress, and the transition dynamics are typically asymmetric [19]. Consequently, the estimation of severity should extend beyond preserving ordinal rank and incorporate the direction of progression inherent to the pathology [7].

Most existing learning formulations overlook this directional constraint [22]. Methods based on regression, ordinal classification, and ranking typically enforce symmetric consistency by penalizing disagreement between two samples without verifying whether the relationship aligns with the disease course [20,28]. Under such symmetric supervision, constraints propagate equally from severe to mild cases and vice versa [17]. Consequently, the training objective may favor numerically smooth yet clinically inconsistent solutions [11,12,29].

We assert that ordinal supervision for DR should be intrinsically directional [23]. A valuable inductive bias is to encode forward-biased progression dynamics, where pairwise relationships utilized for supervision reflect feasible worsening trajectories under typical screening settings. From this perspective, the representation space should admit a manifold [10,9] where transitions are permitted only forward in severity, and the learning signal should enforce monotonicity along these forward transitions. Importantly, reverse relations should not impose constraints, as they correspond to biologically implausible trajectories and can introduce conflicting gradients that obscure the semantics of severity.

Motivated by this perspective, we propose Directed Ordinal Diffusion Regularization (D-ODR), integrating progression awareness into ordinal learning frameworks. During training, D-ODR constructs a directed graph within each mini-batch, establishing connections between feature-neighboring instances exclusively when the destination instance represents an equal or more advanced disease stage. The optimization objective enforces the predicted continuous severity scores to monotonically increase along these valid forward paths, which ensures that the learning signal strictly follows clinically valid trajectories.

Conceptually, the proposed approach models predicted severity as a scalar field over a directed manifold mirroring disease evolution. By propagating ordinal structure through multi-step diffusion, the regularizer captures extended progression chains while remaining anchored in local feature similarity [9,13,4]. This module operates only during training, introducing zero computational overhead at inference. Evaluations across multiple public datasets and backbones show D-ODR yields consistent improvements. The results validate that respecting disease progression irreversibility serves as a practical principle for DR grading.

2 Methods

2.1 Model Overview

Given a retinal fundus dataset $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$ with ordinal severity labels $y_i \in \{0, 1, 2, 3, 4\}$, our objective is to learn a model f_θ that predicts a continuous severity score. As illustrated in Fig. 1, f_θ comprises a feature encoder g_θ and a severity prediction head h_ϕ . For each input image x_i , the encoder extracts




Fig. 1. Overview of the proposed D-ODR framework. The model constructs a directed graph in the feature space using a directional mask and applies multi-scale diffusion to capture extended progression chains, thereby penalizing score inversion along valid forward trajectories. The regularization module is discarded during inference, so it incurs no additional computational cost.

a d -dimensional representation $z_i = g_\theta(x_i) \in \mathbb{R}^d$, which the head subsequently maps to a continuous score $s_i = h_\phi(z_i) \in \mathbb{R}$.

This architecture is governed by a fundamental clinical prior: DR progression is effectively irreversible, typically evolving from mild to severe stages. Consequently, disease samples in the representation space are expected to lie on a manifold with a directional constraint, where local transitions are valid only along paths of non-decreasing severity. To explicitly align model predictions with this forward progression, we introduce D-ODR. Within each training mini-batch, D-ODR constructs a label-driven directed graph in the feature space, enforcing ordinal consistency along the disease manifold by penalizing score inversions.

2.2 Directed Diffusion Graph and Regularization

To capture the local geometry of the disease manifold, we construct a directed weighted graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ over the samples in a mini-batch. We first define local neighborhoods by feature similarity. For each sample i , we compute its k -nearest neighbor set $\mathcal{N}_k(i)$, and adopt mutual kNN to keep only pairs that are neighbors of each other, which reduces noisy edges and improves robustness.

We then assign similarity weights using a locally scaled kernel:

$$\widetilde{W}_{ij} = \exp\left(-\frac{\|z_i - z_j\|^2}{\sigma_i \sigma_j}\right), \quad (1)$$

where σ_i denotes the local scale of sample i , which adapts the effective bandwidth to mitigate non-uniform feature density. To explicitly encode the irreversibility prior, we apply a directional mask $\mathbb{I}(y_i \leq y_j)$ that keeps edges within the same level or pointing to a more severe level, yielding the directed adjacency matrix

$$W_{ij} = \widetilde{W}_{ij} \cdot \mathbb{I}(y_i \leq y_j). \quad (2)$$

This construction removes reverse edges when $y_i > y_j$, which prevents the regularizer from enforcing constraints that contradict the progression direction. Since the directional mask can yield rows with $\sum_j W_{ij} = 0$, we add an identity self-loop $W \leftarrow W + I$ to guarantee nonzero row sums and thus a valid random-walk normalization.

Finally, we row-normalize W to obtain a directed transition matrix:

$$P_{ij} = \frac{W_{ij}}{\sum_k W_{ik}}, \quad (3)$$

where P_{ij} is the probability of diffusing from sample i to sample j along the progression direction.

A one-hop neighborhood may not capture longer progression chains, such as $i \rightarrow m \rightarrow j$. We therefore model multi-scale diffusion via t -step random walks [6]:

$$P^{(t)} = P^t, \quad t \in \{1, 2, 3\}. \quad (4)$$

Because every edge respects the non-decreasing severity constraint, $P_{ij}^{(t)} > 0$ implies that j is reachable from i through a valid path and thus implicitly indicates $y_i \leq y_j$. Multi-scale diffusion propagates local ordinal structure to a broader range, which strengthens global consistency of the regularization.

The key principle of D-ODR is that if sample i can reach sample j through diffusion, the graph construction enforces $y_i \leq y_j$, and the predictions should satisfy $s_i \leq s_j$. A violation $s_i > s_j$ indicates a reverse flow along the disease manifold and should be penalized. We define D-ODR as

$$\mathcal{L}_{\text{ODR}} = \sum_{t \in \mathcal{T}} \frac{1}{t} \sum_{i,j} P_{ij}^{(t)} \text{ReLU}\{s_i - s_j\}, \quad (5)$$

where $\text{ReLU}\{x\} = \max\{x, 0\}$ penalizes only ordering violations. The step-decayed weighting emphasizes local consistency while preventing overly strong constraints from long-range diffusion.

2.3 Supervision, Objective, and Inference

To learn continuous severity scores, we use ordinal supervision in a regression form with mean squared error (MSE):

$$\mathcal{L}_{\text{MSE}} = \frac{1}{N} \sum_{i=1}^N (s_i - y_i)^2. \quad (6)$$

The overall training objective is a weighted sum of MSE loss and D-ODR:

$$\mathcal{L} = \mathcal{L}_{\text{MSE}} + \lambda \mathcal{L}_{\text{ODR}}, \quad (7)$$

where λ controls the regularization strength.

At inference time, we compute $s = f_{\theta}(x)$ only. The D-ODR module is used solely for regularization during training and is not involved in testing, so it introduces no additional inference overhead and keeps $O(1)$ inference complexity.

3 Datasets and Implementation Details

Datasets. We evaluate our method on four public DR datasets: APTOS [14], Messidor [1,8], DeepDR [18], and DDR [16]. All datasets are preprocessed and split according to the official GDRBench [3] protocols.

Implementation Details. All experiments are conducted on four NVIDIA RTX A6000 GPUs with a batch size of 128. The initial learning rate is set to 1×10^{-4} and linearly decayed to zero. All images are resized and normalized using standard ImageNet statistics, and random horizontal flipping is applied as the only stochastic data augmentation during training. To ensure statistical reliability, each experiment is repeated three times with different random seeds, and we report the mean performance alongside the standard deviation.

Evaluation. Performance is evaluated using the Quadratic Weighted Kappa (QWK) and macro F1 score. We adopt a unified checkpoint selection protocol for all models. For each training run, we compute a composite score by summing the min-max normalized QWK and F1 on the validation set. The checkpoint corresponding to the highest composite score is selected for reporting.

4 Results

4.1 Comparison with State-of-the-Art Methods

Table 1 shows D-ODR achieves state-of-the-art performance across four benchmarks, proving that injecting explicit progression directionality into ordinal learning provides effective inductive bias for DR grading. With the ViT-B/16 backbone, D-ODR performs best nearly uniformly, showing the proposed regularizer particularly benefits generic vision backbones under standard ordinal supervision. ViT-based baselines mainly encode global rank structure, leaving local neighborhoods directionally ambiguous near adjacent grades. D-ODR resolves this by coupling local feature geometry with directed multi-step diffusion, enhancing transitive consistency. Using the RETFound backbone, D-ODR improves the strong baseline, showing it is backbone-agnostic and complementary to advanced ophthalmic pretraining.

However, gains vary across datasets, reflecting differences in label noise, acquisition protocols, and progression cue strength. Highly heterogeneous datasets

Table 1. Performance comparison on four datasets. † and ‡ indicate ViT-B/16 and RETFound backbones, respectively.

Model	APTOS		Messidor		DeepDR		DDR	
	<i>QWK</i>	<i>F1</i>	<i>QWK</i>	<i>F1</i>	<i>QWK</i>	<i>F1</i>	<i>QWK</i>	<i>F1</i>
ViT-B/16[27]	87.6 $\pm$ 1.1	65.8 $\pm$ 1.7	57.4 $\pm$ 2.3	51.2 $\pm$ 3.8	74.4 $\pm$ 0.5	56.6 $\pm$ 0.4	75.5 $\pm$ 1.0	54.2 $\pm$ 0.3
RETFound[32]	92.1 $\pm$ 0.2	72.3 $\pm$ 0.5	81.2 $\pm$ 0.8	70.5 $\pm$ 0.4	81.6 $\pm$ 1.2	62.8 $\pm$ 0.5	85.5 $\pm$ 0.5	60.2 $\pm$ 1.7
CORN [†] [24]	88.9 $\pm$ 0.2	68.8 $\pm$ 0.7	60.9 $\pm$ 2.3	49.2 $\pm$ 5.0	74.6 $\pm$ 2.4	56.8 $\pm$ 0.6	73.5 $\pm$ 1.8	53.8 $\pm$ 2.4
POEs [†] [17]	89.6 $\pm$ 0.7	67.4 $\pm$ 0.5	62.7 $\pm$ 1.1	48.1 $\pm$ 4.3	78.8 $\pm$ 1.4	56.4 $\pm$ 1.3	74.8 $\pm$ 0.8	52.4 $\pm$ 1.0
GOL [†] [15]	88.3 $\pm$ 1.1	64.6 $\pm$ 1.9	52.6 $\pm$ 2.9	45.6 $\pm$ 1.8	74.9 $\pm$ 1.2	51.3 $\pm$ 0.5	75.3 $\pm$ 1.5	50.7 $\pm$ 0.3
Ord2Seq [†] [28]	89.0 $\pm$ 0.5	67.6 $\pm$ 0.9	60.3 $\pm$ 2.1	50.0 $\pm$ 6.0	75.7 $\pm$ 1.2	56.9 $\pm$ 1.0	75.0 $\pm$ 1.1	54.1 $\pm$ 0.7
CLIP-DR [†] [30]	86.9 $\pm$ 0.4	62.8 $\pm$ 1.1	50.7 $\pm$ 1.6	40.9 $\pm$ 6.5	72.9 $\pm$ 0.6	51.3 $\pm$ 0.6	73.7 $\pm$ 2.0	50.4 $\pm$ 1.0
AOR-DR [†] [31]	89.4 $\pm$ 0.5	67.2 $\pm$ 0.7	58.2 $\pm$ 0.9	51.5 $\pm$ 4.0	75.0 $\pm$ 1.1	56.3 $\pm$ 1.6	75.3 $\pm$ 0.7	52.1 $\pm$ 0.7
AOR-DR [‡] [31]	85.9 $\pm$ 0.7	63.6 $\pm$ 1.3	71.0 $\pm$ 2.2	57.9 $\pm$ 1.0	72.0 $\pm$ 1.0	52.7 $\pm$ 0.6	77.5 $\pm$ 0.9	53.3 $\pm$ 1.0
D-ODR[†]	89.8 $\pm$ 0.8	67.0 $\pm$ 2.0	68.3 $\pm$ 1.4	53.0 $\pm$ 2.4	80.1 $\pm$ 0.7	57.4 $\pm$ 1.5	76.8 $\pm$ 0.3	55.7 $\pm$ 0.5
D-ODR[‡]	93.2$\pm$0.3	73.3$\pm$0.8	83.9$\pm$0.5	70.9$\pm$1.8	85.7$\pm$0.1	64.7$\pm$0.1	88.2$\pm$0.1	64.6$\pm$1.5

or those lacking inter-grade separability benefit most, as directional diffusion suppresses contradictory gradients and stabilizes decision boundaries by enforcing monotonicity along plausible transitions. Conversely, if pretrained representations already capture progression-related structure, D-ODR still demonstrates backbone-agnosticity and complementary improvements. These results highlight modeling irreversibility and directed reachability as a practical principle improving clinical reliability across diverse DR datasets.

4.2 Ablation Study

Table 2 evaluates each algorithmic component’s contribution to performance improvements. Removing the directional graph construction (w/o Dir.) reverts the relational manifold to standard laplacian smoothing. This inappropriately enforces bidirectional proximity between samples of varying severities, fails to encode the irreversible progression prior. Consequently, the network learns a clinically ambiguous representation, pulling mild and advanced cases together and washing out ordinal semantics. Distorting the asymmetric forward constraint (w/o FC) by improperly reversing the expected progression penalty induces extreme training instability, yielding drastically inflated F1 score standard deviations across multiple datasets.

Restricting regularization to one-hop neighborhoods (w/o MSD) limits representational capacity. Single-step transitions only enforce local smoothness, failing to capture the extended disease evolution sequence. Multi-scale diffusion aligns predictions with clinically plausible progression, establishing transitive ordinal relationships. Results across multiple public datasets show consistent gains on complementary metrics, indicating progression-constrained diffusion regularization effectively captures the progressive nature of retinal degradation.

Table 2. Ablation study of different regularization designs with two backbones. Dir.: direction constraint; FC: forward constraint; MSD: multi-scale diffusion.

Model	APTOS		Messidor		DeepDR		DDR	
	<i>QWK</i>	<i>F1</i>	<i>QWK</i>	<i>F1</i>	<i>QWK</i>	<i>F1</i>	<i>QWK</i>	<i>F1</i>
Backbone: ViT-B/16								
Baseline	87.7 $\pm$ 1.1	65.7 $\pm$ 1.7	57.4 $\pm$ 2.3	51.2 $\pm$ 3.8	74.4 $\pm$ 0.5	56.6 $\pm$ 0.4	75.5 $\pm$ 1.0	54.2 $\pm$ 0.3
w/o Dir.	88.7 $\pm$ 0.7	62.5 $\pm$ 1.2	64.7 $\pm$ 4.4	49.8 $\pm$ 0.4	79.0 $\pm$ 1.2	54.7 $\pm$ 0.3	76.2 $\pm$ 1.3	49.0 $\pm$ 0.3
w/o FC	87.7 $\pm$ 2.2	58.5 $\pm$ 13.2	62.4 $\pm$ 5.2	45.9 $\pm$ 11.9	77.4 $\pm$ 3.7	51.3 $\pm$ 13.2	75.0 $\pm$ 1.0	49.2 $\pm$ 4.0
w/o MSD	88.8 $\pm$ 1.7	65.9 $\pm$ 2.3	63.0 $\pm$ 2.4	46.0 $\pm$ 8.0	80.1 $\pm$ 1.7	56.9 $\pm$ 1.7	76.3 $\pm$ 0.7	52.5 $\pm$ 1.2
D-ODR	89.8$\pm$0.8	67.0$\pm$2.0	68.3$\pm$1.4	53.0$\pm$2.4	80.1$\pm$0.7	57.4$\pm$1.5	76.8$\pm$0.3	55.7$\pm$0.5
Backbone: RETFound								
Baseline	92.1 $\pm$ 0.2	72.3 $\pm$ 0.5	81.2 $\pm$ 0.8	70.5 $\pm$ 0.4	81.6 $\pm$ 1.2	62.8 $\pm$ 0.5	85.5 $\pm$ 0.5	60.2 $\pm$ 1.7
w/o Dir.	92.3 $\pm$ 0.5	70.9 $\pm$ 1.8	82.6 $\pm$ 2.0	69.5 $\pm$ 1.1	84.5 $\pm$ 1.2	64.0 $\pm$ 0.3	87.9 $\pm$ 0.2	63.5 $\pm$ 1.2
w/o FC	92.5 $\pm$ 1.1	68.1 $\pm$ 10.1	81.6 $\pm$ 1.4	64.1 $\pm$ 9.9	84.6 $\pm$ 2.5	55.5 $\pm$ 16.7	87.4 $\pm$ 0.9	59.5 $\pm$ 2.8
w/o MSD	93.0 $\pm$ 0.1	72.1 $\pm$ 0.2	82.9 $\pm$ 0.8	67.7 $\pm$ 1.3	85.0 $\pm$ 0.3	63.8 $\pm$ 1.1	87.8 $\pm$ 0.3	62.4 $\pm$ 0.4
D-ODR	93.2$\pm$0.3	73.3$\pm$0.8	83.9$\pm$0.5	70.9$\pm$1.8	85.7$\pm$0.1	64.7$\pm$0.1	88.2$\pm$0.1	64.6$\pm$1.5

The observed ablation trends remain consistent across both the general-purpose vision transformer and the ophthalmic foundational model. Consistent degradation upon removing any single module confirms conventional representation learning paradigms struggle to preserve strict temporal irreversibility. The proposed components synergistically correct this deficiency, transforming a static ordinal ranking objective into a continuous, biologically plausible geometric flow.

4.3 Sensitivity Analysis

Fig. 2 analyzes the sensitivity of the regularization weight λ , batch size (bs), and neighborhood size k . The model demonstrates strong robustness to λ across the range $[0.5, 2.0]$ and remains largely insensitive to variations in k between 15 and 35. However, while APTOS and DDR are stable across all batch sizes, performance on Messidor and DeepDR degrades noticeably at $bs = 256$, suggesting that excessively large batches may dilute the local manifold structure essential for diffusion. Based on these observations, we adopt $\lambda = 1.0$, $bs = 128$, and $k = 25$ as the default configuration.

4.4 Qualitative Analysis

To intuitively illustrate the representational improvements introduced by D-ODR, Fig. 3(a) visualizes the latent feature space via UMAP projections. The baseline ViT-B/16 yields a fragmented, heavily entangled topology, treating severity grades as isolated clusters and failing to capture the continuous etiology of diabetic retinopathy. In contrast, D-ODR sculpts the latent space into a contiguous, unidirectional manifold, where stages unfold along a clear progressive trajectory from healthy to advanced states. This provides compelling

visual evidence that the directed diffusion mechanism embeds the biological irreversibility of DR progression directly into the representation space geometry.




Fig. 2. Sensitivity analysis of hyperparameters λ , bs , and k on four datasets. The translucent shaded regions indicate $\pm$ one standard deviation.




Fig. 3. Qualitative evaluation of feature representations and score predictions. (a) 2D UMAP projections; (b) Predicted severity score distributions.

Furthermore, Fig. 3(b) shows the distributions of predicted continuous severity scores. For the baseline model, intermediate stages exhibit notable intra-class variance and inter-class overlap, highlighting ambiguity near clinical decision boundaries. Conversely, D-ODR refines these distributions, yielding severity intervals that better follow ordinal monotonicity, with median predictions progressing more consistently across stages. While some overlap remains between adjacent grades, improved inter-class alignment and slightly reduced intra-class variance suggest that eliminating contradictory backward gradients during training enables more consistent, biologically faithful severity grading.

5 Conclusion

This work proposes D-ODR, a progression-aware ordinal learning framework that incorporates the irreversibility of diabetic retinopathy progression into training. D-ODR constructs a directed diffusion structure and enforces ordinal consistency

only along forward trajectories, which aligns predictions with clinically plausible progression and avoids constraints induced by reverse or label-inconsistent relations. Experiments on multiple public datasets show consistent gains over existing methods across complementary metrics, indicating that progression-constrained diffusion regularization provides an effective and backbone-agnostic way to improve continuous severity estimation.

Acknowledgments. This work was supported by the National Natural Science Foundation of China (Grant No. 62372408) and Hangzhou Pujian Medical Technology Co., Ltd, China and ZJU-Pujian Research & Development Center of Medical Artificial Intelligence for Hepatobiliary and Pancreatic Disease.

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

References

1. Abràmoff, M.D., Folk, J.C., Han, D.P., Walker, J.D., Williams, D.F., Russell, S.R., Massin, P., Cochener, B., Gain, P., Tang, L., Lamard, M., Moga, D.C., Queleec, G., Niemeijer, M.: Automated analysis of retinal images for detection of referable diabetic retinopathy. *JAMA Ophthalmology* **131**(3), 351–357 (2013)
2. Arcadu, F., Benmansour, F., Maunz, A., Willis, J., Haskova, Z., Prunotto, M.: Deep learning algorithm predicts diabetic retinopathy progression in individual patients. *npj Digital Medicine* **2**(1), 92 (2019)
3. Che, H., Cheng, Y., Jin, H., Chen, H.: Towards generalizable diabetic retinopathy grading in unseen domains. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 430–440. Springer (2023)
4. Chen, H., Chen, Y., Yan, Z., Ding, M., Li, C., Qin, F., et al.: MMLNB: Multi-modal learning for neuroblastoma subtyping classification assisted with textual description generation. *Biomedical Signal Processing and Control* **126**, 110798 (2026)
5. Cheung, N., Mitchell, P., Wong, T.Y.: Diabetic retinopathy. *The Lancet* **376**(9735), 124–136 (2010)
6. Coifman, R.R., Lafon, S.: Diffusion maps. *Applied and Computational Harmonic Analysis* **21**(1), 5–30 (2006)
7. Dai, L., Sheng, B., Chen, T., Wu, Q., Liu, R., Cai, C., Wu, L., Yang, D., Hamzah, H., Liu, Y., Wang, X., Guan, Z., Yu, S., Li, T., Tang, Z., Ran, A., et al.: A deep learning system for predicting time to progression of diabetic retinopathy. *Nature Medicine* **30**(2), 584–594 (2024)
8. Decencière, E., Zhang, X., Cazuguel, G., Lay, B., Cochener, B., Trone, C., et al.: Feedback on a publicly distributed image database: the messidor database. *Image Analysis & Stereology* **33**(3), 231–234 (2014)
9. Jia, J., Liu, Y., Sun, Y., Chen, H., Qin, F., Wang, C., Peng, Y.: Geodesic prototype matching via diffusion maps for interpretable fine-grained recognition. In: *ICASSP 2026 - 2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 1786–1790 (2026)
10. Jia, J., Wang, J., Liu, Y., Jing, H., et al.: This looks distinctly like that: Grounding interpretable recognition in stiefel geometry against neural collapse (2026)
11. Jia, J., Wu, Y., Chen, H., Jing, H., Wang, H., Bu, J., Wu, L.: Unsupervised causal prototypical networks for de-biased interpretable dermoscopy diagnosis (2026)

12. Jia, J., Zheng, H., Wu, Y., Chen, H., Wang, H., Bu, J., Wu, L.: SCOUT: Selective coupling via optimal unbalanced transport for interpretable text classification. In: Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 6413–6431 (2026)
13. Jiang, S., Chen, Z., Xu, L., Zhu, Y., Wang, C., Zhang, J., Qin, F., Chen, Y., Zhu, Z.: Multimodal feature prototype learning for interpretable and discriminative cancer survival prediction (2025)
14. Karthik, Maggie, Dane, S.: Aptos 2019 blindness detection. Kaggle competition (2019)
15. Lee, S.H., Shin, N.H., Kim, C.S.: Geometric order learning for rank estimation. In: Proceedings of the 36th International Conference on Neural Information Processing Systems. pp. 1–13 (2022)
16. Li, T., Gao, Y., Wang, K., et al.: Diagnostic assessment of deep learning algorithms for diabetic retinopathy screening. *Information Sciences* **501**, 511–522 (2019)
17. Li, W., Huang, X., Lu, J., et al.: Learning probabilistic ordinal embeddings for uncertainty-aware regression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 13896–13905 (2021)
18. Liu, R., Wang, X., Wu, Q., Dai, L., Fang, X., Yan, T., Son, J., Tang, S., Li, J., Gao, Z., Galdran, A., Poorneshwaran, J.M., et al.: DeepDRiD: Diabetic retinopathy—grading and image quality estimation challenge. *Patterns* p. 100512 (2022)
19. Maturi, R.K., Glassman, A.R., Josic, K., et al.: Four-year visual outcomes in the protocol w randomized trial of intravitreal aflibercept for prevention of vision-threatening complications of diabetic retinopathy. *JAMA* **329**(5), 376–385 (2023)
20. Pitawela, D., Carneiro, G., Chen, H.T.: CLOC: Contrastive learning for ordinal classification with multi-margin n-pair loss. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 15538–15548 (2025)
21. Ruamviboonsuk, P., Tiwari, R., Sayres, R., et al.: Real-time diabetic retinopathy screening by deep learning in a multisite national screening programme: a prospective interventional cohort study. *The Lancet Digital Health* **4**(4), e235–e244 (2022)
22. Runje, D., Shankaranarayana, S.M.: Constrained monotonic neural networks. In: Proceedings of the 40th International Conference on Machine Learning. pp. 1–16. JMLR (2023)
23. Shaik, N.S., Cherukuri, T., et al.: Ordinal label-distribution learning with constrained asymmetric priors for imbalanced retinal grading. In: The Second Workshop on GenAI for Health: Potential, Trust, and Policy Compliance (2025)
24. Shi, X., Cao, W., Raschka, S.: Deep neural networks for rank-consistent ordinal regression based on conditional probabilities. *Pattern Analysis and Applications* **26**(3), 941–955 (2023)
25. Sun, Y., He, Y., Jia, J., Wang, J., Ge, R., Wang, C., Xu, H.: WDT-MD: Wavelet diffusion transformers for microaneurysm detection in fundus images (2025)
26. Toledo-Cortés, S., Useche, D.H., Müller, H., González, F.A.: Grading diabetic retinopathy and prostate cancer diagnostic images with deep quantum ordinal regression. *Computers in Biology and Medicine* **145**, 105472 (2022)
27. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: *Advances in Neural Information Processing Systems*. vol. 30, pp. 5998–6008 (2017)
28. 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 (ICCV). pp. 5865–5875 (2023)
29. Yang, C., Song, D., Pang, X., Shen, Q., et al.: LLM-bootstrapped targeted finding guidance for factual MLLM-based medical report generation (2026)

30. Yu, Q., Xie, J., Nguyen, A., Zhao, H., et al.: CLIP-DR: Textual knowledge-guided diabetic retinopathy grading with ranking-aware prompting. In: Medical Image Computing and Computer Assisted Intervention. pp. 667–677. Springer (2024)
31. Yu, Q., Zhou, W., Liu, H., Xu, Y., Wang, M., Zhao, Y., Fu, H., Ye, X., Zheng, Y., Meng, Y.: Parameterized diffusion optimization enabled autoregressive ordinal regression for diabetic retinopathy grading. In: Medical Image Computing and Computer Assisted Intervention. pp. 450–460. Springer (2025)
32. Zhou, Y., Chia, M.A., Wagner, S.K., et al.: A foundation model for generalizable disease detection from retinal images. *Nature* **622**(7981), 156–163 (2023)