

Reliable and Interpretable Visual Field Progression Prediction with Diffusion Models and Conformal Risk Control

Wenwen Si^{*1}, Vivian Lin^{*1}, Bo Sun¹, Kuk Jin Jang¹, Rubo Xing¹, Almiqdad Saeed², Rina Nagatani², Oleg Sokolsky¹, Insup Lee¹, and Lama Al-Aswad^{1,2}

¹ Department of Computer and Information Science, University of Pennsylvania, Philadelphia, PA, USA
{wenwens, vilin}@seas.upenn.edu

² Department of Ophthalmology, University of Pennsylvania, Philadelphia, PA, USA

Abstract. Accurately predicting visual field progression is critical for early intervention and personalized treatment of glaucoma. However, existing methods struggle with both predictive accuracy and reliable uncertainty quantification. This paper introduces a framework that leverages diffusion models and conformal risk control to generate robust and interpretable forecasts of visual field deterioration. We first train a diffusion model to predict future visual fields based on a patient’s past examinations. To ensure trustworthy predictions, we design a novel archetypal-based conformal risk control method, which provides finite-sample coverage guarantees on intervals of archetypal contributions. This framework captures the underlying structures within uncertainty, enabling clinicians to interpret a range of potential progression patterns rather than a single deterministic outcome. Experimental results illustrate that our method achieves the target archetypal contribution coverage while providing tighter prediction intervals than baselines. Visualizations show how archetypal visual field patterns contribute to prediction uncertainty, offering interpretable insights into disease progression. By combining diffusion models with conformal methods, our framework enhances the reliability of AI-assisted visual field forecasting, ultimately supporting improved clinical decision-making. Our code is available at: <https://github.com/averysi224/abci.git>.

Keywords: Visual field · Diffusion models · Conformal Risk Control.

1 Introduction

Glaucoma is projected to affect over 100 million people worldwide by 2040 [21]. Caused by optic nerve damage, it can lead to irreversible functional vision loss and significantly impacts quality of life. Visual function in glaucoma patients is assessed through visual field (VF) tests, which provide differential light sensitivity at different points on the retina. For instance, a 24-2 VF test records

^{*}Equal contributions.

sensitivity in decibels (dB) across 54 points within the central 24 degrees of vision. These values form a grayscale image, where darker areas indicate worse sensitivity. Fig. 1 in Section 3 illustrates a case where the patient’s future VF shows defects in the lower hemisphere.

Predicting visual field progression is crucial for glaucoma interventions but is challenging due to limited clinical consensus on the disease’s pathology [14]. To assist clinicians, various methods have been proposed. Linear regression has been used to predict individual points [18] and the mean deviation (MD) metric [10], but this assumes linear progression over time, and MD does not capture spatial patterns. Deep learning approaches, including a variational autoencoder [5], a diffusion model [22], recurrent network [19], and convolutional model [24] have instead been used for predicting full VFs. However, existing deep learning methods do not incorporate clinical knowledge of archetypal vision loss patterns, which provide insight into patient quality of life (e.g., highly impactful on vision and quality of life), reveal structural changes in the eye [8], and help estimate rates of VF progression [16]. Moreover, deep learning models can fail to identify patients with worsening visual fields despite low overall prediction error [9].

Conformal inference [1,2,15,23] provides a statistical framework for quantifying prediction uncertainty with finite-sample guarantees. While extensively studied, its application to the image generation domain remains limited. The pixel-wise method im2im-uq [3] fails to capture spatial patterns, producing uninformative uncertainty bounds. Conffusion [11] uses quantile regression [13] to obtain quantile diffusion models as heuristic uncertainty intervals. However, its plausibility relies on a high-quality score network, which is hardly achievable with noisy medical datasets. PUQ [4] improves upon these methods by applying principal component analysis to the posterior uncertainty, yielding tighter intervals. All these approaches follow a risk-controlling paradigm [2] to ensure joint coverage under high-dimensional pixel dependencies.

Aiming to overcome these limitations, we propose Archetype Based Conformal Intervals (ABCI), a framework for reliable and clinically interpretable VF prediction that leverages diffusion models and uncertainty quantification of predicted archetypal loss patterns. We first use a diffusion model to predict VF progression. Diffusion models excel at modeling images, surpassing variational autoencoders and generative adversarial networks [7,20]. As probabilistic methods, they also capture the uncertainties inherent in the visual field data. We then apply a novel archetype-based conformal method to derive confidence intervals for significant archetype contributions to prediction uncertainty. These intervals offer marginal coverage guarantees, ensuring two key properties with high probability: first, the most significant archetypes contributing to prediction uncertainty are determined with low error, and second, their contributions to the future visual field are reliably estimated. Ours is the only approach that provides *joint rigorous guarantees* on the selection of significant archetypes and the estimation of their contributions to future VF. In practice, this offers clinicians a reliable range of patterns within which the future VF is expected to fall.

We validate our method on two real-world VF datasets: the University of Washington Humphrey Visual Field (UWHVF) dataset [17] and the Glaucoma Research Network (GRN) dataset. We compare it against two conformal baselines: pixel-wise conformal prediction (im2im-uq) [3] and principal uncertainty quantification (PUQ) [4]. Unlike the baseline methods, ABCI is the only approach that provides clinically interpretable and reliable visual field progression predictions. By enabling accurate risk assessment and actionable clinical insights, ABCI is well-suited for computer-aided glaucoma treatment.

2 Background and Problem Definition

In this section, we provide a brief but necessary overview of archetypal analysis and conformal risk control. The formal problem formulation follows.

2.1 Archetypal Patterns of Vision Loss

Visual fields are recorded as d measured points (e.g., $d = 52$ for 24-2 VFs, excluding the blindspot). Given VF records $D \subset \mathbb{R}^d$, archetypal analysis [6] identifies the extremal patterns of the data, located on the boundaries of the convex hull of D . Although derived mathematically, these archetypes align with clinically recognized loss patterns [8].

Given the fixed archetype matrix constructed from p archetypes of D ,

$$A = [v_1, v_2, \dots, v_p] \in \mathbb{R}^{d \times p}, \quad (1)$$

any $x \in D$ can be reconstructed by obtaining the coefficient vector $s \in \mathbb{R}_{>0}^p$ with $\sum_{i=1}^p s_i = 1$ that minimizes the residual sum of squares, $\min_s \|x - As\|^2$.

2.2 Conformal Risk Control (CRC)

CRC is a form of conformal inference that controls general risk metrics beyond standard coverage guarantees. For any covariate x , CRC additionally enforces constraints on a pre-defined risk function $\hat{L}(x, y)$, such as mis-coverage or prediction error, while maintaining statistical validity: $\Pr(\mathbb{E}[\hat{L}(x, y)] \leq \alpha) \geq 1 - \delta$, where α is the target risk level, and δ is the failure probability.

2.3 Problem Formulation

Formally, let $\mathcal{X}, \mathcal{Y} \subset \mathbb{R}^d$ represent the spaces of existing and future VF records, respectively, and let $\mathcal{T}$ denote the space of time spans between two records. Define an observation as $o = (x, t)$, where $x \in \mathcal{X}$ is the current visual field and $t \in \mathcal{T}$ is the prediction time span of interest. Denote $P_{o,y}$ as the joint probability distribution over $(\mathcal{X} \times \mathcal{T}) \times \mathcal{Y}$. Assume a diffusion model mapping $(\mathcal{X} \times \mathcal{T}) \rightarrow \mathcal{Y}$ and a fixed matrix $A \in \mathbb{R}^{d \times p}$ of p archetypal patterns are given.

Given an observation $o = (x, t)$, we aim to quantify the uncertainty of its VF progression prediction, as characterized by the posterior distribution estimated




Fig. 1. Overview of ABCI. a) At inference time, VF predictions are sampled from the diffusion model. Given α , β , and δ , with probability $1 - \delta$, conformal risk control estimates the most significant archetypes contributing to the future VF and intervals on these contributions. b) The intervals reliably include the ground-truth archetype contributions. The clinician can trust that the future VF will lie within this interval. c) The clinician can trust that the predicted most significant archetypes accurately represent the uncertainties in future VF loss patterns.

by the diffusion model, $\hat{P}_{y|o}$. Rather than defining intervals over individual pixels, we construct an interval-valued function $\mathcal{C}(o; \hat{A}(o))$ on archetypal contributions to the future VF, where $\hat{A}(o) = A \odot \hat{m}(o)$ and $\hat{m}(o)$ is a binary selection mask predicting the most significant archetypes contributing to uncertainty. For simplicity, we denote $\hat{A}(o)$ as $\hat{A}$ in the following demonstration.

Given a user-specified $\alpha \in (0, 1)$ and the ground-truth output $y \in \mathcal{Y}$, the interval $\mathcal{C}(o; \hat{A})$ should contain the projected ground-truth archetypal contributions with probability at least $1 - \alpha$, while ensuring the correctness of $\hat{m}$,

$$\mathbb{E} \left[\frac{1}{|\hat{m}(x)|} \sum_{\hat{m}(o)_i=1} \mathbb{1} \left\{ v_i^T y \in \mathcal{C}(o; \hat{A})_i \right\} \right] > 1 - \alpha. \quad (2)$$

3 Methodology

In this section, we present our Archetype Based Conformal Intervals (ABCI) method for constructing intervals with joint rigorous guarantees on both archetypal contributions to the ground truth and the selection of high-contributing archetypal patterns of progression uncertainty. Fig. 1 summarizes our approach.

3.1 Visual Field Prediction

First, we train a diffusion model $\hat{P}_{y|o}$ on a training set with (o, y) pairs. We build upon the previous state-of-the-art [22] in visual field prediction, employing a diffusion model with a transformer backbone to predict a patient's future visual field from a single past observation. Crucially, we introduce the input of the prediction horizon, encoded via sinusoidal positional embeddings.

3.2 Conformal Risk Control for Archetype Contributions

We use the trained diffusion model $\hat{P}_{y|o}$ as the score function. Building on this fixed diffusion model, we construct $\mathcal{C}(o; \hat{A})$ on the hold-out calibration set as

$$\mathcal{C}_\lambda(o; \hat{A}_i) := \left[v_i^T \hat{\mu}(o) - \lambda \hat{l}(o)_i, \quad v_i^T \hat{\mu}(o) + \lambda \hat{u}(o)_i \right]. \quad (3)$$

Here, $\hat{\mu}(o)$ is the average prediction, and $\lambda \in \mathbb{R}^+$ is a tunable parameter that scales the model-predicted uncertainty, $\hat{u}(o)$ and $\hat{l}(o)$.

The interval must satisfy two statistical constraints. First, we ensure that our selection mask $\hat{m}(o; \lambda_1)$ identifies key archetypal patterns that capture the essence of uncertainty. Denote $y_c = y - \hat{\mu}(o)$ as the *centered ground-truth* VF, and n as the sampling size for each o , we control the reconstruction error:

$$L_1(o, y; \lambda_1) := \hat{Q}_q \left(\{ |\bar{y}_c(o) - y_{c,i}| \}_{i=1}^n \right), \quad (4)$$

where $\bar{y}_c = \hat{A} \left(\hat{A}^T \hat{A} \right)^{-1} \hat{A}^T y_c$ is the archetypal reconstruction of y_c , with the corrections for non-orthogonality. This ensures that the conditionally selected archetypes recover more than q percent of the ground-truth value in y_c .

Specifically, we define the *model uncertainty* as $\hat{y}_c = \hat{y} - \hat{\mu}(o)$. For each o , we compute the normalized archetypal weights $\hat{w}$ as:

$$\hat{w}(o) = \frac{[\hat{\sigma}_1^2(o), \dots, \hat{\sigma}_p^2(o)]}{c}, \quad \text{where} \quad \hat{\sigma}_i(o) = v_i^T \hat{y}_c, \quad c = \sum_{j=1}^p \hat{\sigma}_j^2(o),$$

serving as scores indicating the predicted significance. To highlight the high-contributing archetypes, we reorder $\{v_i\}$ in non-increasing order of their weights to $\{v_{(i)}\}$. Then, the mask of $\hat{A}$ is constructed by adjusting λ_1 (Eq. (4)) to include the top K archetypal patterns with high predicted significance:

$$\hat{m}(o; \lambda_1)_i = 1, \quad \text{for } 1 \leq (i) \leq K; \quad K = \min \left\{ k \mid \sum_{(i)=1}^k \hat{w}_{(i)}(o) \geq \lambda_1 \right\}.$$

We then control the coverage risk L_2 for these archetypal patterns, where $\lambda_2 \in \mathbb{R}^+$ represents the calibration factor parameter:

$$L_2(o, y; \lambda_1, \lambda_2) := \frac{1}{|\hat{m}(o; \lambda_1)|} \sum_{i \text{ s.t. } \hat{m}_i(o)=1} \mathbb{1} \left\{ \hat{v}_i(o)^T y \notin \mathcal{C}_{\lambda_2}(o; \hat{A})_i \right\}. \quad (5)$$

Given a user-specified acceptable miscoverage rate α , a reconstruction error range β , and desired confidence level $1 - \delta$, λ_1 and λ_2 are selected as the extreme values that jointly satisfy the following coverage and reconstruction guarantee:

$$\Pr \left(\begin{array}{c} \mathbb{E}[L_1(o, y; \lambda_1)] \leq \beta \\ \mathbb{E}[L_2(o, y; \lambda_1, \lambda_2)] \leq \alpha \end{array} \right) \geq 1 - \delta. \quad (6)$$

In practice, since selecting high-contributing archetypes for prediction uncertainty controls the reconstruction risk, these archetypes sufficiently represent the patterns in the ground-truth visual progression uncertainty.

4 Experiments

Datasets. We verify the validity of our framework on two 24-2 visual field datasets, the UWHVF [17] and GRN datasets. The open-source UWHVF dataset contains 28,943 24-2 HVF tests from 3,871 patients and 7,428 eyes. The GRN dataset has 602,000 24-2 visual fields from 129,000 patients. This study adheres to the Declaration of Helsinki.

Evaluation Setup. For each dataset, we uniformly sample 3,000 visual field pairs and categorize them into three subgroups based on Hodapp-Parrish-Anders-on (HPA) mean deviation (MD): Mild ($\text{MD} > -6$ dB), Moderate ($-12 \text{ dB} < \text{MD} < -6$ dB), and Severe ($\text{MD} < -12$ dB). We specify miscoverage rates $\alpha = 0.25, 0.3$, and 0.4 , splitting each subgroup into calibration and test sets in a 4:1 ratio. For all the experiments, we set $q = 0.9 \sim 0.95$, and $\delta = 0.1$. To demonstrate the robustness of our method, we do not fix the random seed.

Evaluation Metrics. We measured the accuracy of our diffusion model using Mean Absolute Error (MAE) in decibels (dB). For evaluating ABCI, we compute the empirical coverage and the average interval size on the test set. A valid method obtains coverage error less than α , while smaller interval size is preferred.


Archetypal Analysis and Diffusion Model. We use 17 precomputed visual field archetypes [8], verified to align with OHTS clinical classifications [12]. Our diffusion model is trained on UWHVF for 20 epochs (batch size 64, learning rate 10^{-4}) and on GRN for 10 epochs (batch size 5, learning rate 10^{-5}), following [22].

5 Results

On the UWHVF dataset, our horizon-conditioned diffusion model achieved a 3.20 dB MAE, reducing error by over 50% compared to the 6.78 dB baseline [22]. On the GRN dataset, which is $20\times$ larger, our model’s 3.71 dB MAE is comparable to the 3.64 dB baseline. This highlights the effectiveness of horizon conditioning for smaller datasets. Next, we assess ABCI on these fixed diffusion models.

Table 1. Quantitative Results on UWHVF Dataset.

Stage	α	ABCI				im2im-uq		PUQ	
		β	Cvrg Risk	# Comp.	Pixel Itvl Size	Cvrg Risk	Pixel Itvl Size	Cvrg Risk	Pixel Itvl Size
Mild	0.25	0.1	0.157	6.53	0.199	0.192	0.338	<u>0.437</u>	0.033
	0.3	0.1	0.198	6.44	0.170	0.223	0.312	<u>0.460</u>	0.028
	0.4	0.1	0.309	7.53	0.137	0.344	0.263	<u>0.493</u>	0.023
Moderate	0.25	0.14	0.183	4.17	0.222	0.173	0.357	<u>0.358</u>	0.032
	0.3	0.14	0.239	4.27	0.192	0.208	0.331	<u>0.369</u>	0.029
	0.4	0.14	0.287	4.08	0.148	0.307	0.286	0.390	0.023
Severe	0.25	0.155	0.193	4.33	0.254	0.179	0.344	0.203	0.058
	0.3	0.155	0.217	4.52	0.245	0.224	0.313	0.212	0.049
	0.4	0.155	0.349	4.51	0.181	0.325	0.267	0.233	0.037

**Fig. 2.** Qualitative results on the UWHVF dataset. On the left, we present the ground truth visual field, the bounds predicted by ABCI, and the selected main visual patterns. On the right, we show the bounds produced by the baseline im2im-uq.

5.1 Results on UWHVF Dataset

We first evaluate ABCI on the UWHVF dataset, with quantitative results presented in Table 1. ABCI effectively controls coverage risk, maintaining it below α across all settings. While im2im-uq meets the coverage requirement, its intervals are overly conservative due to its limited ability to capture visual loss structures. In contrast, PUQ fails to guarantee archetypal contribution coverage for all mild cases and most moderate cases.


In Fig. 2, we visualize the two methods that meet the coverage requirement. ABCI effectively captures visual patterns, making the bounds more informative. Additionally, the main uncertainty components reflect potential differences between the diffusion model prediction and the ground-truth progression. In contrast, im2im-uq produces scattered bounds that lack interpretability.

5.2 Results on GRN Dataset

Next, we evaluate ABCI on the GRN dataset in Table 2. Both ABCI and im2im-uq achieve coverage guarantees, while PUQ fails. Similarly, ABCI generally pro-

Table 2. Quantitative Results on GRN Dataset.

Stage	α	ABCI				im2im-uq		PUQ	
		β	Cvrg Risk	# Comp.	Pixel Itvl Size	Cvrg Risk	Pixel Itvl Size	Cvrg Risk	Pixel Itvl Size
Mild	0.25	0.1	0.179	5.88	0.278	0.192	0.571	0.429	0.059
	0.3	0.1	0.196	5.75	0.244	0.223	0.538	0.463	0.049
	0.4	0.1	0.357	5.92	0.169	0.344	0.396	0.495	0.041
Moderate	0.25	0.15	0.169	3.91	0.258	0.173	0.458	0.297	0.158
	0.3	0.15	0.232	4.77	0.292	0.208	0.424	0.319	0.138
	0.4	0.15	0.308	3.88	0.214	0.307	0.201	0.343	0.108
Severe	0.25	0.18	0.210	2.84	0.171	0.179	0.712	0.205	0.176
	0.3	0.18	0.245	2.85	0.159	0.224	0.679	0.225	0.149
	0.4	0.18	0.203	2.71	0.164	0.325	0.186	0.248	0.119

**Fig. 3.** Qualitative results on the GRN dataset. On the left, we present the ground truth visual field, the bounds predicted by ABCI, and the selected main visual patterns. On the right, we show the bounds produced by the baseline im2im-uq.

duces tighter intervals than im2im-uq. As shown in Fig. 3, im2im-uq’s bounds lack interpretability. Moreover, these bounds must be clipped to $(0, 1)$ to ensure reasonable values, highlighting its tendency to generate overly biased estimates.

5.3 Cross-Dataset Experiments

The experiments thus far validate ABCI’s effectiveness on two VF datasets. However, real-world clinical settings may involve diverse patient populations, leading to distribution shifts that violate conformal inference’s exchangeability assumption. To evaluate ABCI’s practical value, we conduct a cross-dataset stress test. Due to UWHVF’s limited size, its diffusion model fails to generalize. Instead, we use the GRN model for calibration on GRN and testing on UW, and vice versa, as shown in Table 3. Results demonstrate ABCI’s strong empirical robustness, maintaining coverage error at or below the target α in most cases.

Table 3. ABCI cross-dataset coverage risks under different settings.

Set/ α /Stage		Mild			Moderate			Severe		
Calibration	Test	0.25	0.3	0.4	0.25	0.3	0.4	0.25	0.3	0.4
GRN	UWHVF	0.246	0.276	0.407	0.121	0.142	0.220	0.203	0.238	0.315
UWHVF	GRN	0.089	0.135	0.212	0.148	0.209	0.319	0.134	0.181	0.313

6 Conclusion

We propose ABCI, a clinically interpretable and reliable approach to visual field progression prediction by leveraging diffusion models and conformal risk control. Unlike existing work, our method explicitly quantifies uncertainties in archetypal vision loss patterns, providing interpretable support for clinical decision-making. Evaluations on UWHVF and GRN datasets show ABCI consistently obtains desired coverage while capturing meaningful vision loss patterns, underscoring its robustness and practical value in computer-aided clinical intervention.

Acknowledgments. This study was funded in part by the National Institute of Health (NIH) (grant number R01 EY037101-02). This work was also supported in part by a gift from AWS AI to Penn Engineering’s ASSET Center for Trustworthy AI, and also funded by Research to Prevent Blindness.

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

References

1. Angelopoulos, A.N., Bates, S.: A gentle introduction to conformal prediction and distribution-free uncertainty quantification. arXiv preprint arXiv:2107.07511 (2021)
2. Angelopoulos, A.N., Bates, S., Fisch, A., Lei, L., Schuster, T.: Conformal risk control. arXiv preprint arXiv:2208.02814 (2022)
3. Angelopoulos, A.N., Kohli, A.P., Bates, S., Jordan, M.I., Malik, J., Alshaabi, T., Upadhyayula, S., Romano, Y.: Image-to-image regression with distribution-free uncertainty quantification and applications in imaging. In: Proceedings of the 39th International Conference on Machine Learning. pp. 717–730. PMLR (2022)
4. Belhasin, O., Romano, Y., Freedman, D., Rivlin, E., Elad, M.: Principal uncertainty quantification with spatial correlation for image restoration problems. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023)
5. Berchuck, S.I., Mukherjee, S., Medeiros, F.A.: Estimating rates of progression and predicting future visual fields in glaucoma using a deep variational autoencoder. Scientific Reports **9**(1), 18113 (2019)
6. Cutler, A., Breiman, L.: Archetypal analysis. Technometrics **36**(4), 338–347 (1994)
7. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in neural information processing systems **34**, 8780–8794 (2021)
8. Elze, T., Pasquale, L.R., Shen, L.Q., Chen, T.C., Wiggs, J.L., Bex, P.J.: Patterns of functional vision loss in glaucoma determined with archetypal analysis. Journal of The Royal Society Interface **12**(103), 20141118 (2015)

9. Eslami, M., Kim, J.A., Zhang, M., Boland, M.V., Wang, M., Chang, D.S., Elze, T.: Visual field prediction: evaluating the clinical relevance of deep learning models. *Ophthalmology Science* **3**(1), 100222 (2023)
10. Fujino, Y., Murata, H., Mayama, C., Asaoka, R.: Applying “lasso” regression to predict future visual field progression in glaucoma patients. *Investigative ophthalmology & visual science* **56**(4), 2334–2339 (2015)
11. Horwitz, E., Hoshen, Y.: Confusion: Confidence intervals for diffusion models. arXiv preprint arXiv:2211.09795 (2022)
12. Keltner, J.L., Johnson, C.A., Cello, K.E., Edwards, M.A., Bandermann, S.E., Kass, M.A., Gordon, M.O., Group, O.H.T.S., et al.: Classification of visual field abnormalities in the ocular hypertension treatment study. *Archives of Ophthalmology* **121**(5), 643–650 (2003)
13. Koenker, R., Bassett Jr, G.: Regression quantiles. *Econometrica: journal of the Econometric Society* pp. 33–50 (1978)
14. Lee, S.S.Y., Mackey, D.A.: Glaucoma–risk factors and current challenges in the diagnosis of a leading cause of visual impairment. *Maturitas* **163**, 15–22 (2022)
15. Lei, J., G’Sell, M., Rinaldo, A., Tibshirani, R.J., Wasserman, L.: Distribution-free predictive inference for regression. *Journal of the American Statistical Association* **113**(523), 1094–1111 (2018)
16. Liu, J., Sood, S., Razavian, N., Chen, D., Yousefi, S., Elze, T., De Moraes, G., Boland, M.V., Wellik, S.R., Pasquale, L.R., et al.: Predicting progression of glaucoma based on visual field tests. *Investigative Ophthalmology & Visual Science* **64**(8), 335–335 (2023)
17. Montesano, G., Chen, A., Lu, R., Lee, C.S., Lee, A.Y.: Uwhvf: a real-world, open source dataset of perimetry tests from the humphrey field analyzer at the university of washington. *Translational Vision Science & Technology* **11**(1), 2–2 (2022)
18. Nouri-Mahdavi, K., Hoffman, D., Gaasterland, D., Caprioli, J.: Prediction of visual field progression in glaucoma. *Investigative ophthalmology & visual science* **45**(12), 4346–4351 (2004)
19. Park, K., Kim, J., Lee, J.: Visual field prediction using recurrent neural network. *Scientific reports* **9**(1), 8385 (2019)
20. 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)
21. Tham, Y.C., Li, X., Wong, T.Y., Quigley, H.A., Aung, T., Cheng, C.Y.: Global prevalence of glaucoma and projections of glaucoma burden through 2040: a systematic review and meta-analysis. *Ophthalmology* **121**(11), 2081–2090 (2014)
22. Tian, Y., Zang, M., Sharma, A., Gu, S.Z., Leshno, A., Thakoor, K.A.: Glaucoma progression detection and humphrey visual field prediction using discriminative and generative vision transformers. In: *International Workshop on Ophthalmic Medical Image Analysis*. pp. 62–71. Springer (2023)
23. Vovk, V., Gammerman, A., Shafer, G.: *Algorithmic Learning in a Random World*, vol. 29. Springer, Berlin, Germany (2005)
24. Wen, J.C., Lee, C.S., Keane, P.A., Xiao, S., Rokem, A.S., Chen, P.P., Wu, Y., Lee, A.Y.: Forecasting future humphrey visual fields using deep learning. *PloS one* **14**(4), e0214875 (2019)