Predicting Glaucoma Progression with a Variable-Time Neural Controlled Differential Equation

Sai Mulagan, MS[1]; Sophia Y. Wang, MD, MS[2]; Cody Carroll, PhD[1]

[1]University of San Francisco, San Francisco, CA ; [2]Byers Eye Institute, Stanford University, Palo Alto, CA

Introduction

- Glaucoma causes irreversible vision loss through progressive optic nerve damage. Predicting future visual field (VF) deterioration is critical for timely intervention.
- VF tests measure mean deviation (MD), a summary metric of overall visual field sensitivity (in decibels, dB), recorded at irregular intervals.
- We adapt the Neural Controlled Differential Equation (Neural CDE)[1,4] framework to predict MD at any future time from a short baseline window of irregularly-sampled VF tests.

Methods

Cohort

Starting cohort
31,913 VFs
5,465 eyes
3,082 patients

→ **After filtering**
≥2 baseline + ≥1 future VF
2,272 eyes
1,391 patients

→ **Train**
1,814 eyes · 5,263 target VFs

→ **Test**
458 eyes · 1,406 target VFs

80/20 patient-level split

Evaluation subgroups
First-year slope < −1 dB/yr = fast
Fast: 597 (26.3%) | **Slow: 1,675 (73.7%)**

1-year	3-year	5-year
179 test eyes	232 test eyes	140 test eyes

Model

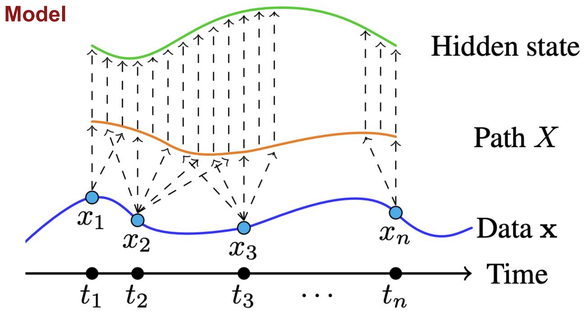

- **Input.** Each eye contributes an irregular sequence of (time, MD) pairs from the first-year baseline window. The sequence is interpolated into a continuous path X(t) using Hermite cubic splines (torchcde[1]).
- **Architecture.** A Neural CDE integrates a learned vector field along the input path:

  dz(t) = f(z(t)) · dX(t)

- f is a 3-layer MLP (hidden 128, tanh). z is initialized from X(0) via linear projection, integrated with RK4 (step 0.5). A readout MLP maps [z(T); target time t] to predicted MD. **The target time is an input**, so one trained network predicts at any future time.
- **Training.** 50,033 parameters · 48 hidden channels · Adam (lr 1e-3, wd 1e-4) · MSE loss · grad clip 1.0. Early stopping (patience 25) on a 15% inner-validation split.

Evaluation

- **Primary (variable-time)**: All 1,406 held-out future VFs from 458 test eyes, evaluated at their actual irregular time points.
- **Secondary (per-interval):** Performance broken down by year-long intervals. For each test eye, the closest VF within ±180 days of 1, 3, and 5 years.
- **Subgroup analysis:** Stratified by progression speed (first-year MD slope < −1 dB/yr defines fast progressors; 26.3% of eyes).
- **Metrics:** R², RMSE, MAE with 95% bootstrap CIs (1,000 resamples).

Results

Neural CDE: Actual vs Predicted MD

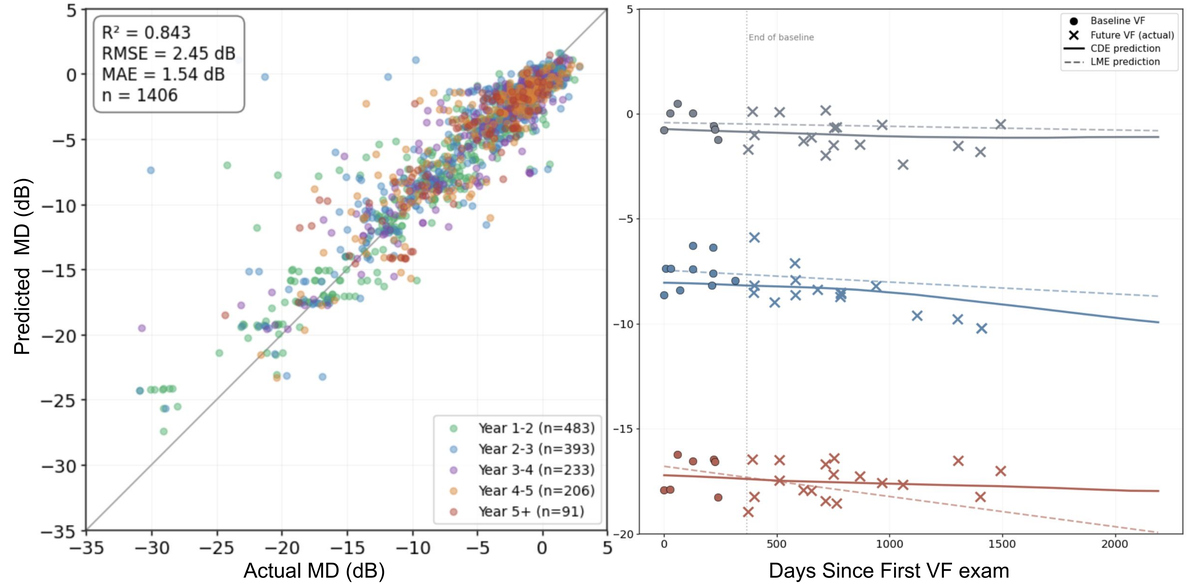

Performance by Year Interval (with 95% Bootstrap CI) CDE vs FPCA

Interval	RMSE Fast	RMSE Slow	R² Fast	R² Slow	RMSE Fast	RMSE Slow	R² Fast	R² Slow
Year 1–2	2.92	1.77	0.872	0.925	3.771	3.274	0.844	0.716
Year 2–3	3.03	2.89	0.772	0.762	3.804	3.224	0.847	0.700
Year 3–4	1.84	2.29	0.874	0.823	4.193	3.080	0.836	0.710
Year 4–5	2.69	2.16	0.707	0.805	4.452	3.118	0.846	0.717
Year 5–6	2.55	2.53	0.788	0.753	4.257	3.156	0.681	0.660
Overall	**2.77**	**2.32**	**0.834**	**0.845**	**3.665**	**3.091**	**0.946**	**0.728**

Variable-Time Evaluation (all 1,406 test VF predictions)

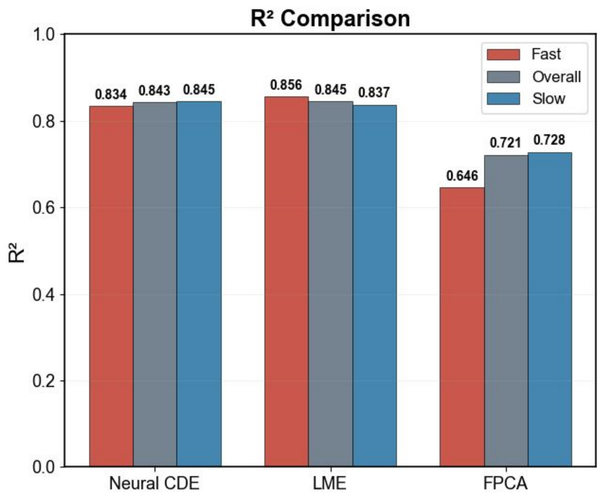

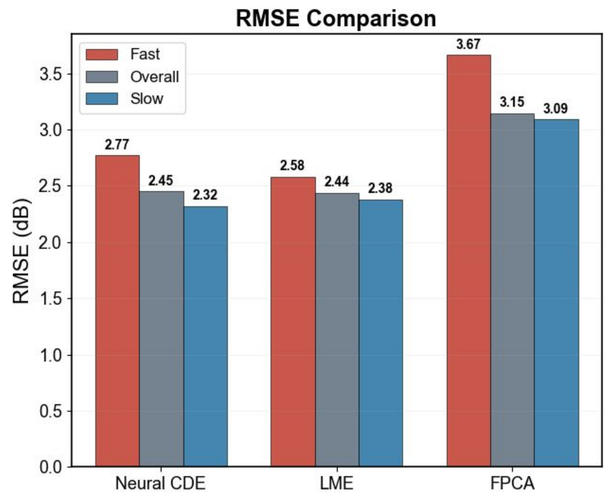

Baselines

- Linear mixed effects model (LME): population-level linear model with per-eye random intercept and slope, predicted via best linear unbiased prediction (BLUP).
- Functional principal component analysis (FPCA)²: prior published work on a similar Stanford glaucoma cohort, which additionally incorporated retinal nerve fiber layer (RNFL) thickness and electronic health record (EHR) features
- Right panel: example trajectories for two fast and two slow progressors. CDE (solid) captures nonlinear decline that LME (dashed) cannot.

Conclusions and Discussion

- **The Neural CDE matches LME (R² 0.843 vs 0.845) and outperforms FPCA² on fast progressors (R² 0.834 vs 0.646) without any EHR or imaging features.**
- **Lower RMSE than FPCA across all time intervals for both fast and slow progressors.**
- **Stable performance from 1 to 6 years out, enabling long-horizon forecasting from one year of baseline data.**
- The CDE handles irregular VF timing natively, no imputation or fixed-grid resampling required.
- Outperforms FPCA² at every interval despite fewer input features, suggesting the continuous-time formulation captures progression dynamics more effectively than functional principal components.
- Separate fast/slow CDE models did not improve over the single model the CDE adapts trajectory shape from the input path without explicit subgroup stratification.
- Extending the baseline from 1 to 2 years improves R² only modestly (0.843 → 0.859), making the 1-year model more practical for early clinical use.

Limitations

- Single-institution retrospective dataset; external validation pending.
- Median baseline sequence length is 2 VFs, limiting how much trajectory shape the CDE can extract from the input.
- The FPCA² comparison uses a different iteration of the Stanford glaucoma cohort with different inclusion criteria. Notably, FPCA can accept eyes with as few as 1 baseline VF, whereas the CDE requires ≥2, and the FPCA model additionally incorporated retinal nerve fiber layer (RNFL) thickness and electronic health record (EHR) features as inputs.

Future Work

- Multimodal extension: add EHR or ophthalmic imaging features as additional CDE input channels.
- Pointwise (52-point) VF prediction with the same variable-time framework.
- Prospective external validation across institutions.

References

[1] Kidger, P. et al. Neural Controlled Differential Equations for Irregular Time Series. *NeurIPS* 2020.

[2] Wang, S.Y. et al. Functional principal component analysis of glaucomatous visual field progression. *Ophthalmology* (prior published work).

[3] Garway-Heath, D.F. et al. UK Glaucoma Treatment Study (UKGTS). *Lancet* 2015.

[4] Chen, R.T.Q. et al. Neural Ordinary Differential Equations. *NeurIPS* 2018.

Contact: mulagan@stanford.edu

Byers Eye Institute · Stanford University