tvst

Author Response to Letter

Author Response: Using Provenance Dispersion to Triage Uncertain Neurosymbolic OCT Diagnoses

We thank Bair for the careful reading of our work and for highlighting provenance-aware uncertainty as an important translational direction for neurosymbolic OCT systems.[1,2] We agree that, in clinical settings, understanding when evidence is borderline, internally conflicted, or outside the intended diagnostic frame can be as actionable as identifying the most likely diagnosis.

In the present framework, a practical view of competing hypotheses is already available through the pathology-level probability outputs reported for all diagnostic classes.[1,2] Our symbolic layer propagates the convolutional neural network's soft sign likelihoods (without binarization) into continuous diagnostic probabilities, and the primary label is obtained by selecting the highest-supported pathology while retaining the full probability profile for interpretation.[2] This probability profile can be inspected to understand how strongly alternative diagnoses are supported and to identify cases in which multiple pathologies receive similar support.

With respect to provenance, our use of $k = 1$ should not be interpreted as yielding a single hard decision or a single explanation for only one condition. Rather, the model computes probabilities for every pathology in parallel, and the top-1 proof path is selected separately for each pathology as the trace supporting that pathology's score.[1,3,4] Consequently, potentially contrasting diagnostic hypotheses can already be appreciated by examining the distribution of pathology probabilities (as illustrated in Figure 3).[2] In practice, these competing supports could also be summarized into an overall uncertainty indicator, for example, by aggregating dispersion across pathology-level probabilities, without changing the deployed decision rule.

We also acknowledge the real-world relevance of coexisting disease patterns. In our study, we did not exclude all images with multiple concurrent pathological features; Table 1 indicates that overlap was retained in the analyzed data.[2] The exclusions were limited to concurrent retinal diseases outside the diagnostic scope of the study that could confound the target label (e.g., diabetic retinopathy in an AMD case), because those processes were not represented by the sign set and symbolic rules evaluated here.[2]

Bair's proposal to incorporate alternative proof paths and conflict indexes (e.g., proof-margin or entropy-based summaries) provides a valuable roadmap for triage-oriented extensions.[1,3,4] We view several of these directions as natural next steps for the framework. In particular, we will consider extending provenance beyond $k = 1$ (i.e., incorporating $k > 1$ proof paths) alongside an expansion of the rule set, so that alternative lines of symbolic support can be summarized into explicit dispersion- or conflict-based uncertainty measures. In parallel, we anticipate that ambiguity can be reduced by strengthening spatial reasoning, for example, by incorporating retinal layer segmentation to encode location-dependent cues, which may help distinguish findings such as subretinal versus intraretinal fluid based on anatomical position rather than appearance alone. Finally, we will consider whether the symbolic layer could be complemented with an additional "pathological but not otherwise specified" category to better accommodate presentations that are abnormal yet fall outside the current sign/rule taxonomy.

Taken together, we appreciate Bair's emphasis on moving from interpretability toward operational trustworthiness, and we will consider these extensions, multi-path provenance, enriched spatial reasoning, and broader coverage for out-of-taxonomy abnormalities, in future developments of the approach.

Aleksandar Miladinović[1], Alessandro Biscontin[1], Miloš Ajčević[2], Simone Kresevic[2], Agostino Accardo[2], Daniele Tognetto[3], and Leandro Inferrera[3]

[1] Institute for Maternal and Child Health IRCCS "Burlo Garofolo," Trieste, Italy
[2] Department of Engineering and Architecture, University of Trieste, Trieste, Italy
[3] Department of Medicine, Surgery and Health Sciences, Ophthalmology Clinic, University of Trieste, Trieste, Italy
e-mail: aleksandar.miladinovic@burlo.trieste.it.

tvst.arvojournals.org | ISSN: 2164-2591

References

1. Bair H. Using provenance dispersion to triage uncertain neurosymbolic OCT diagnoses [letter]. *Transl Vis Sci Technol*. 2026;15(3):9.
2. Miladinović A, Biscontin A, Ajčević M, et al. Neurosymbolic AI framework for explainable retinal disease classification from OCT images. *Transl Vis Sci Technol*. 2026;15(1):6.
3. Li Z, Huang J, Naik M. Scallop: a language for neurosymbolic programming. *Proc ACM Programming Lang*. 2023;7(PLDI):1463–1487.
4. Li Z, Huang J, Liu J, Naik M. Neurosymbolic programming in scallop: principles and practice. *Foundations Trends Programming Lang*. 2024;8:118–249.

translational vision science & technology