



Deep Learning–Based Clustering of OCT Images for Biomarker Discovery in Age-Related Macular Degeneration (PINNACLE Study Report 4)

Robbie Holland, MEng,¹ Rebecca Kaye, MD,² Ahmed M. Hagag, MD,^{3,4} Oliver Leingang, PhD,⁵ Thomas R.P. Taylor, MD,² Hrvoje Bogunović, PhD,^{3,6} Ursula Schmidt-Erfurth, MD,⁵ Hendrik P.N. Scholl, MD,^{7,8} Daniel Rueckert, PhD,^{1,9} Andrew J. Lotery, MD,² Sobha Sivaprasad, MD,^{3,4} Martin J. Menten, PhD^{1,9}

Purpose: We introduce a deep learning–based biomarker proposal system for the purpose of accelerating biomarker discovery in age-related macular degeneration (AMD).

Design: Retrospective analysis of a large data set of retinal OCT images.

Participants: A total of 3456 adults aged between 51 and 102 years whose OCT images were collected under the PINNACLE project.

Methods: Our system proposes candidates for novel AMD imaging biomarkers in OCT. It works by first training a neural network using self-supervised contrastive learning to discover, without any clinical annotations, features relating to both known and unknown AMD biomarkers present in 46 496 retinal OCT images. To interpret the learned biomarkers, we partition the images into 30 subsets, termed clusters, that contain similar features. We conduct 2 parallel 1.5-hour semistructured interviews with 2 independent teams of retinal specialists to assign descriptions in clinical language to each cluster. Descriptions of clusters achieving consensus can potentially inform new biomarker candidates.

Main Outcome Measures: We checked if each cluster showed clear features comprehensible to retinal specialists, if they related to AMD, and how many described established biomarkers used in grading systems as opposed to recently proposed or potentially new biomarkers. We also compared their prognostic value for late-stage wet and dry AMD against an established clinical grading system and a demographic baseline model.

Results: Overall, both teams independently identified clearly distinct characteristics in 27 of 30 clusters, of which 23 were related to AMD. Seven were recognized as known biomarkers used in established grading systems, and 16 depicted biomarker combinations or subtypes that are either not yet used in grading systems, were only recently proposed, or were unknown. Clusters separated incomplete from complete retinal atrophy, intraretinal from subretinal fluid, and thick from thin choroids, and, in simulation, outperformed clinically used grading systems in prognostic value.

Conclusions: Using self-supervised deep learning, we were able to automatically propose AMD biomarkers going beyond the set used in clinically established grading systems. Without any clinical annotations, contrastive learning discovered subtle differences between fine-grained biomarkers. Ultimately, we envision that equipping clinicians with discovery-oriented deep learning tools can accelerate the discovery of novel prognostic biomarkers.

Financial Disclosure(s): Proprietary or commercial disclosure may be found in the Footnotes and Disclosures at the end of this article. *Ophthalmology Science* 2024;4:100543 © 2024 by the American Academy of Ophthalmology. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Age-related macular degeneration (AMD) is the leading cause of irreversible vision loss in the developed world and is projected to affect nearly 300 million people by 2040.¹ Clinical grading systems^{2–5} drive the screening, diagnosis, and monitoring of AMD using a small set of known imaging biomarkers in retinal OCT. The early and intermediate AMD stages are characterized by increasing sizes of drusen, which are lipidic subretinal deposits that often precede conversion to late AMD. It is currently impossible to

reliably predict if and when a patient will progress to late AMD, which is classified as either macular neovascularization (MNV), identified by macular fluid, or geographic atrophy, indicated by retinal and retinal pigment epithelium thinning and hypertransmission of OCT signal in the choroid. The small number of known biomarkers, which permit only a coarse stratification of patients, is limiting the prognostic ability of current grading systems. There is a clear and unmet need for new imaging biomarkers that

better describe patients' physiology and provide improved risk stratification.

Identifying new biomarkers, or subtypes of known ones, is a very challenging task. A new understanding of the pathophysiology of a specific disease can lead to the discovery of new biomarkers.⁶ More commonly, biomarkers are proposed by clinicians making anecdotal observations in small cohorts. Using this ad hoc approach, clinicians have recently proposed double-layer sign (DLS),⁷ vitreomacular interface abnormalities,⁸ and intraretinal hyperreflective foci⁹ as biomarkers for AMD. Before entering clinical practice, candidate biomarkers must undergo extensive testing of their predictive value, reproducibility, and utility.^{10,11} Because drusen were proposed and integrated into clinical grading systems 2 decades ago, no new biomarkers for early stage disease have been fully established.¹² The reliance on anecdotal observations, coupled with the onerous process of biomarker discovery, is a major constraint limiting the rate of biomarker discovery in AMD.

To expedite the initial stage of biomarker discovery, there has been an increased push to harness data-driven approaches.¹³ However, retinal specialists have not been able to translate the growing amount of retinal data into increased rates of biomarker discovery. Systematically searching for potentially new disease subtypes in large populations of patients is a tedious and resource-intensive task. Identifying new subtypes involves finding subgroups of images or patients that share some newly proposed characteristic. However, dividing a data set of only 10 000 retinal images into as many as 30 subgroups would require making at least 300 000 simultaneous pairwise comparisons. The limited capacity in visual working memory^{14,15} makes manual identification and exhaustive analysis of common patterns and biomarkers on this scale practically infeasible for human experts.

Deep learning has shown potential for automating visual tasks at scale in medical data sets.¹⁶ Traditional supervised approaches train networks to predict if any known biomarkers are present in the image. This both constrains networks to train only on the subset of labeled images and limits their ability to discover new biomarkers outside the set known to the clinical annotators. The self-supervised learning paradigm^{17–22} surpasses these limitations by training networks using unlabeled images. To train networks without supervision, researchers designed self-supervised tasks where networks learn features through self-discovery with minimal human input. Examples of these tasks include image reconstruction²³ and contrastive predictive encoding.¹⁹ After training, these learned features can be interpreted and potentially inform clinicians of new biomarkers that were present, but previously unseen, in the data.

Prior applications of self-supervised learning for biomarker discovery mistakenly incentivized networks to learn features that are uncorrelated with disease. Image reconstruction²⁴ and generative adversarial tasks²⁵ prioritize learning features unrelated to imaging biomarkers, such as position and orientation of the retina in the image, and global image brightness and contrast. These features dominate the learning process and obfuscate subtler and

more fine-grained indications of new biomarkers. Recent advances in self-supervised learning by contrastive methods address this issue by training networks to explicitly ignore these irrelevant features.^{20–22} As a result, they extract more robust and semantically relevant image features. Despite their potential for discovering new imaging features related to disease, contrastive methods have received little to no attention for biomarker discovery.

In this work, we develop the first contrastive-learning-based proposal system for AMD imaging biomarkers. By pretraining a network with contrastive learning, we were able to automatically extract self-supervised features from 46 496 OCT images of 6236 eyes from 3456 patients without any clinical supervision. We then automatically identified 30 distinct groups of images, termed clusters, that share similar features. Two teams of retinal specialists independently reviewed the clusters, assigning clinical interpretations to each. Finally, we report clusters characterized by unique sets of features relating to known, suspected, and unknown AMD imaging biomarkers.

Methods

The workflow used in this study is illustrated in [Figure 1](#). We begin this section by introducing the OCT data set and the set of known biomarkers. We then describe the first stage of the automatic discovery process, by which we learn imaging features from this data set using self-supervised contrastive learning. In the second stage, we automatically identify subgroups, termed clusters, of images with shared features. Treating those shared features as possible biomarker candidates, we next describe our strategy to interpret them with direct involvement from ophthalmologists. Finally, we outline our full procedure for evaluating our method.

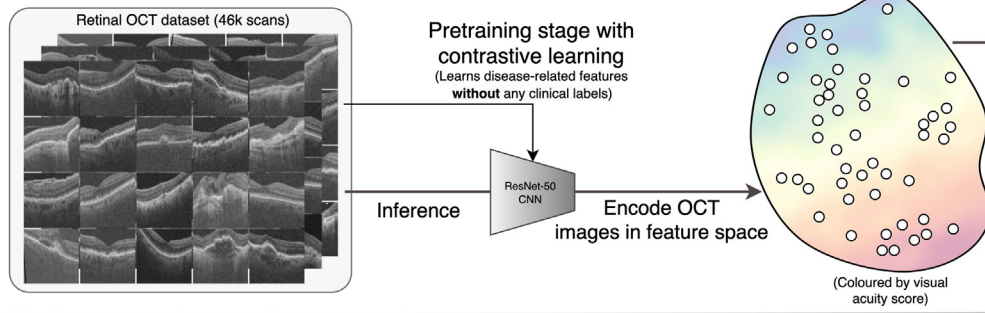
OCT Data Set

Experiments were conducted using a data set of 46 496 OCT images of 6236 eyes from 3456 patients acquired over an 8-year period at the Southampton Eye Unit and collected by the PINNACLE consortium.²⁶ The PINNACLE study ([ClinicalTrials.gov](#) NCT04269304) protocol was approved in the United Kingdom by the East Midlands–Leicester Central Research Ethics Committee (ref. 19/EM/0163) and further by the institutional review boards of all involved institutions. It adheres to the principles of Good Clinical Practice and is in accordance with the Declaration of Helsinki. Informed consent was obtained from the participants after an explanation of the nature and possible consequences of the study. All images were acquired using Topcon 3D OCT scanners (Topcon Corporation). From each, we extracted a 2D mediolateral slice centered on the fovea, adjusted it to a size of 208×256 pixels, and set the pixel size to $7.0 \times 23.4 \mu\text{m}^2$, representing half the median resolution. Visual acuity logMAR scores, which measured the patient's functional quality of vision, were measured at 34 093 visits. These were converted to letter scores, indicating how many letters the patient can read on a chart (from 5–95).

Current grading system annotations are as follows: a subset of 7730 images from 1031 eyes from 1001 patients were labeled by the retinal specialists using the established AMD grading protocols that derive from known imaging biomarkers.

Early/intermediate AMD was characterized by drusen at least $63 \mu\text{m}$ in diameter. We also recorded MNV, complete retinal pigment epithelial and outer retinal atrophy (cRORA; atrophy of

1. Learning AMD features



2. Clustering and proposing biomarkers

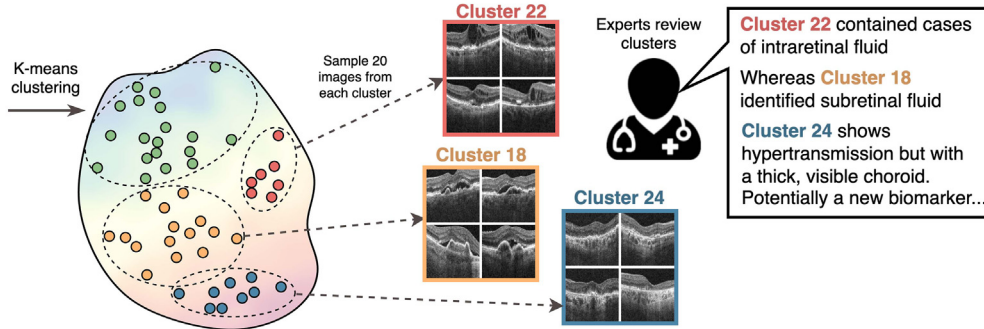


Figure 1. In this study, we design a biomarker proposal system based on contrastive learning. After self-supervised pretraining, we have clustered images that share similar sets of features. Finally, 2 teams of retinal specialists independently identify the characteristic features of each cluster that potentially relate to new biomarkers. AMD = age-related macular degeneration.

width $\geq 250 \mu\text{m}$ and $< 1000 \mu\text{m}$), larger cRORA (atrophy of width $\geq 1000 \mu\text{m}$),²⁷ and healthy cases with no visible pathological biomarkers.

Self-Supervised Contrastive Learning

We used contrastive learning to train a self-supervised feature extractor f , which will form the basis of our biomarker proposal system. Contrastive learning works by creating 2 augmented versions of each image using a predefined set of transformations (Fig 2). By training f to extract a similar set of features from each augmented version, f learns to discern imaging features that are invariant to the transformations. By intentionally using transformations that do not alter the presentation of structural imaging biomarkers, we aimed to increase both the amount and subtlety of learned features that relate to potentially new biomarkers. To this end, we varied global image brightness and contrast, image rotation and aspect ratio, horizontal symmetry, and a randomly sized and located crop. The latter requires f to relate features that cooccur across the OCT images, such as hypertransmission (in the choroidal region) and retinal thinning. The exact parameters used for these transformations are tailored for retinal OCT by Holland et al.²⁸

We choose the contrastive loss from Bootstrap Your Own Latent,²⁸ which trains f to minimize the distance in feature space between the 2 transformed image versions. We parameterize f by a ResNet50 (4x) network²⁹ and train for 120 000 steps using the Adam optimizer with a momentum of 0.9, weight decay of 1.5×10^{-6} , and a learning rate of 5×10^{-4} . Once training was complete, we first removed the final linear layer of f before extracting 2048 unsupervised features from each labeled image. This so-called feature space, or vector database, compactly encoded each image’s most salient features for further analysis.

Clustering and Feature Attribution

We employed a cluster-based approach to identify subgroups of self-supervised features learned by f . These clusters are later reviewed by clinicians, who give clinical interpretations to their characteristic sets of features. To perform the clustering, we used k -means to assign each image to 1 of 30 distinct clusters. We then used 3 annotation strategies to enhance the interpretability of the clusters before review. First, we reordered clusters by their median visual acuity, so that higher cluster numbers indicated a more degraded quality of functional vision. We referred to our clusters using the schema C1 (best vision) up to C30 (worst vision). Second, we correlated each cluster with a disease stage from the current grading system. To do this, we found the probability that any given image from each cluster was annotated with each of the known biomarkers. Third, we highlighted the features in each image that are most representative of its assigned cluster (Fig 2). To this end, we fit a single, linear layer that mapped from the feature space to cluster assignments. Then, to highlight the cluster-specific image features, we generated an attribution map by applying GradCAM³⁰ between the image and the cluster label using the final 2 convolutional layers of f .

Interpreting Clusters with Retinal Specialists

Until this stage, the entire automated discovery process had not used any clinical annotations or manual input. We then aimed to understand and interpret the features that defined and distinguished each cluster. To this end, we recruited 2 independent teams of retinal specialists, each consisting of a senior and a junior ophthalmologist, and conducted 2 parallel 1.5-hour semistructured interviews. For each cluster, we first showed the specialists 10 randomly drawn images and their associated attribution maps,

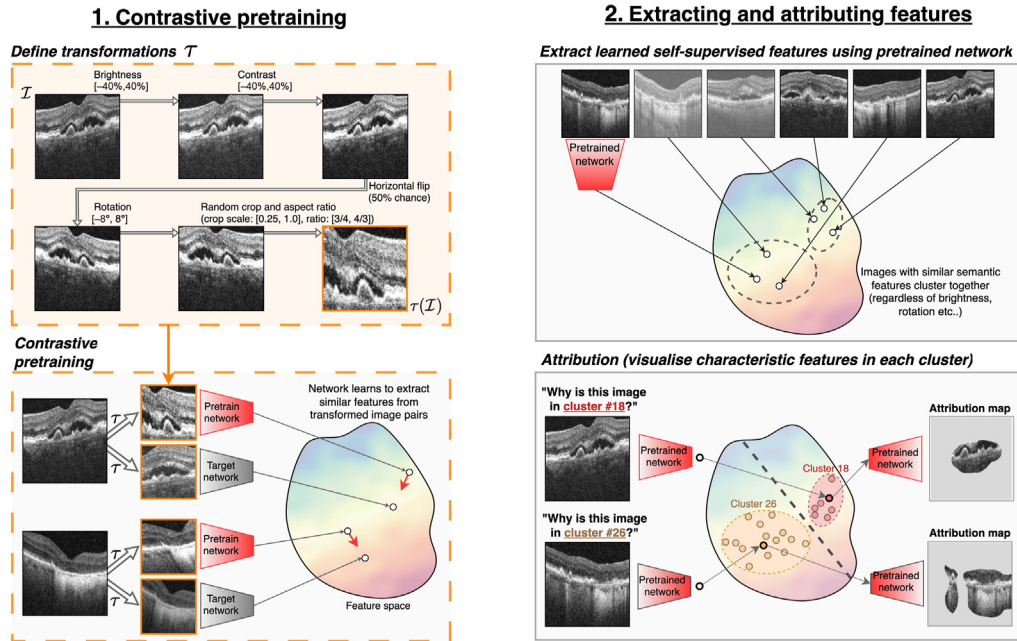


Figure 2. The fully automated backbone of our biomarker proposal system consists of 2 stages. First, self-supervised contrastive learning trains models to identify biomarkers and other image features without any clinical annotations. To do so, it trains networks to ignore a specified set of known invariant image features, defined by the set of contrastive transformations, amplifying the signal of any biomarker-related features. Second, we extract self-supervised image features and cluster images with similar features. After this, we compute attribution maps that highlight the cluster-specific features in each image to assist interpretation by retinal specialists.

ensuring all images were derived from different patients. By comparing multiple images from the same cluster, the retinal specialists could visually extract, distill, and articulate any image features that represented the majority. While interpreting the images, the specialists were also permitted to inspect other clusters that correlated with the same disease grading. This allowed them to better identify image features that are not currently differentiated by established grading systems but were, by definition, distinguishable by their self-supervised features. This could enable the detection of potentially new biomarker subtypes that are currently conflated in each indiscriminate disease stage.

After reviewing the 10 images, each team suggested up to 3 common features using clinical terms, such as large drusen or poor image quality, for each cluster and ranked them according to their prevalence among the images. If there is no identifiable feature describing the majority of images, the cluster is labeled “heterogeneous.” To validate these descriptions, a further set of 10 unseen images and attribution maps were subsequently revealed. The suggested features were only taken as the team’s proposed descriptions if they also described the majority of the unseen images. Finally, if there is semantic consensus between the 2 independent teams’ descriptions, we reported this as the overall description of the cluster. Cluster descriptions may reveal potentially new biomarkers, providing the basis for proposing novel biomarker candidates.

Evaluation

Our main evaluation assessed the ability of our system to propose novel biomarkers. To assess this, we determined how often both expert teams identified the same features in each of the 30 clusters. Next, we recorded the number of clusters where the identifiable features related to biomarkers, as a measure of the ability of contrastive learning to ignore disease-unrelated image features. Finally, we observed how many of those biomarkers were already

included in existing grading systems, as opposed to those that have been proposed by clinical researchers, but are not yet included in grading systems, or were previously unknown. These statistics are our predominant measure of the effectiveness of our biomarker proposal system.

In addition to being discernible in images, biomarkers should provide some degree of risk stratification and prognostic value for late AMD. We compared the predictive value of our clusters with the known biomarkers and a demographic baseline that uses only the patient’s age and sex. To make predictions using our clusters, each image was represented by a vector of size 30 that encodes its similarity to each cluster. This vector was then used by a Lasso linear regression model to estimate the risk. We quantify each system’s prognostic capability by calculating the mean absolute error in years to forecast the time until a patient’s conversion to MNV or cRORA. In addition, we predicted the current visual acuity. Finally, we included a fully supervised learning baseline by fitting a linear support vector regression model directly to the feature space to demonstrate the performance gap between our interpretable cluster-based approach and black-box models. Each experiment used 10-fold cross-validation on random 80/20 partitions while ensuring a patient-wise split. We repeated the entire experiment, from clustering to regression, using 7 random seeds and reported means and standard deviations.

Results

Quality of Clusters

Both teams of retinal specialists independently reported semantically identical descriptions in 27 out of 30 clusters (Fig 3) and remarked on the consistency of image features

AI-derived clusters with expert-derived descriptions

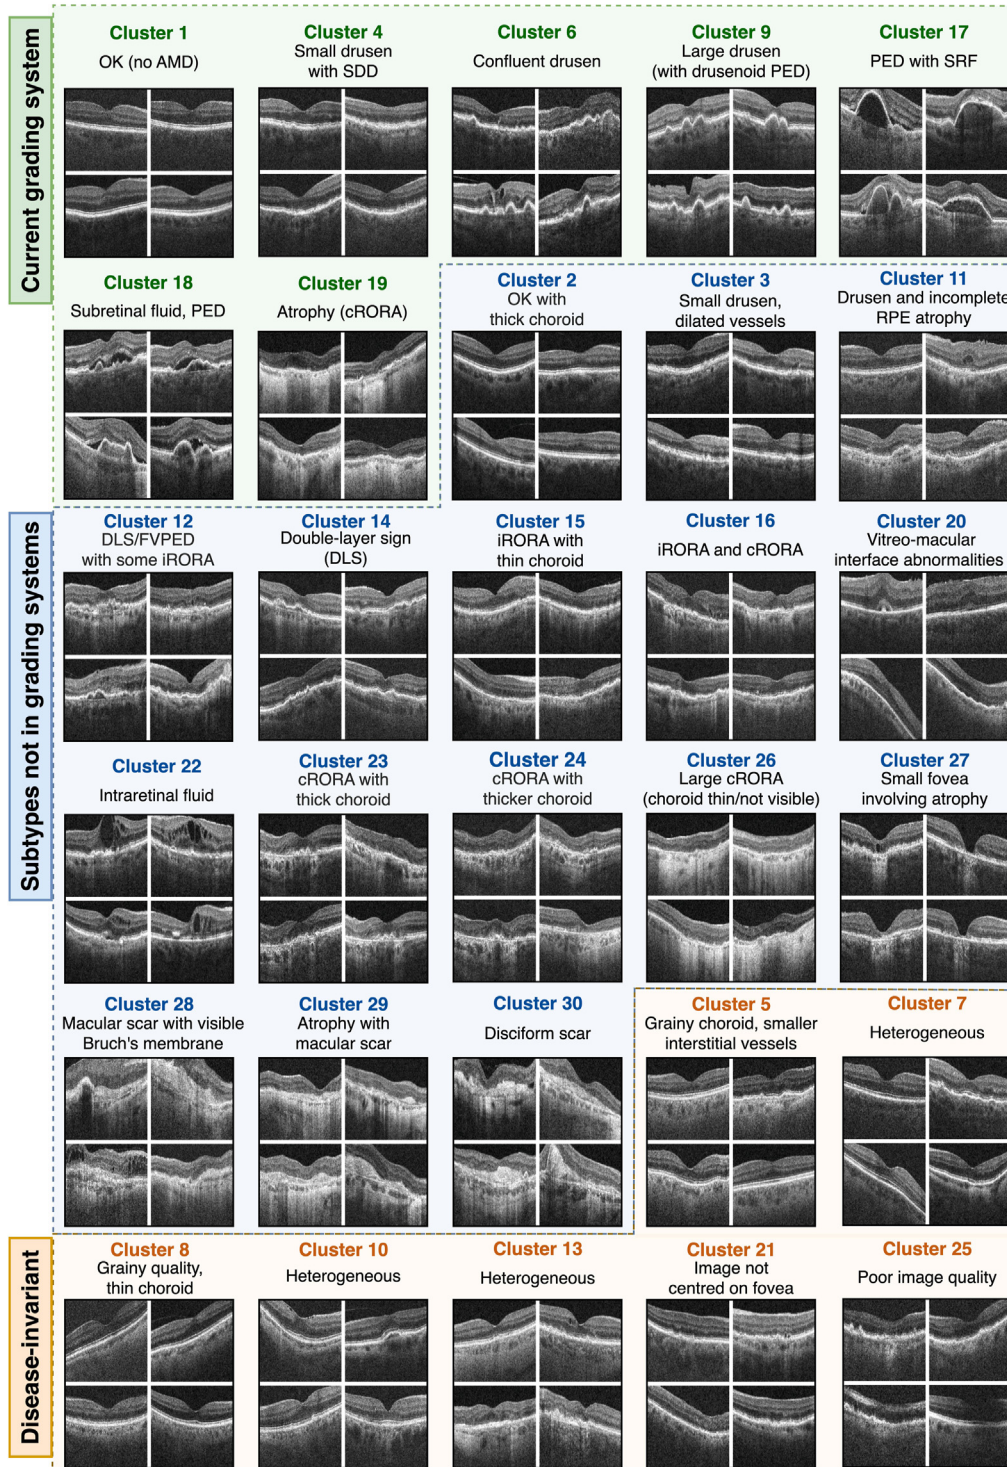


Figure 3. Each cluster with its description derived independently by 2 teams of retinal specialists. In each, we show 4 representative images from different patients. Our proposal system identified these clusters without any human supervision or prior knowledge of known biomarkers through a process of self-discovery. Out of 30 clusters, 23 were related to AMD, of which 16 made subtle distinctions between fine-grained biomarkers that were either unknown to retinal specialists or not included in existing clinical grading systems. AMD = age-related macular degeneration; cRORA = complete retinal pigment epithelial and outer retinal atrophy; DLS = double-layer sign; FVPED = fibrovascular pigment epithelial detachment; iRORA = incomplete retinal pigment epithelial and outer retinal atrophy; PED = pigment epithelial detachment; SDD = subretinal drusenoid deposits; SRF = subretinal fluid.

Table 1. The Number of Images Contained in Each Cluster, in Addition to the Number of Unique Patients the Images Originate from

Cluster	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
#Images	310	262	400	370	234	249	166	166	198	197	395	268	205	233	435
#Patients	196	173	166	174	156	102	114	103	69	121	170	110	126	69	174
Ratio	1.6	1.5	2.4	2.1	1.5	2.4	1.5	1.6	2.9	1.6	2.3	2.4	1.6	3.4	2.5
Cluster	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
#Images	326	187	152	226	123	218	155	303	396	136	233	152	206	211	235
#Patients	135	55	75	148	73	133	82	114	139	93	75	46	88	67	55
Ratio	2.4	3.4	2.0	1.5	1.7	1.6	1.9	2.7	2.8	1.5	3.1	3.3	2.3	3.1	4.3

In Most Clusters, the Average Patient Contributed No More Than 3 Images. Clusters C29 and C30, Which Had Higher Ratios, Related to Scarring and Largely Document Patients Being Monitored Posttreatment.

observed within each. The 2 teams disagreed, or found no clearly identifiable features, in only 3 clusters (C7, C10, C13) deemed heterogeneous. Notably, 23 clusters were found to capture different biomarkers related to AMD. We have approximately divided these into a group of 7 clusters representing biomarkers already in current grading systems and a group of 16 that exhibit potentially new subtypes. In a minority of clusters (C5, C8, C21, C25), the common features were instead related to image or scan quality. Importantly, there were no clusters for image brightness, rotation, and position of the retina in the image, or any feature explicitly included in the set of contrastive transformations. The high proportion of clusters depicting identifiable imaging biomarkers confirms our hypothesis that contrastive learning is able to extract high-level semantic features and ignore many confounding image features.

Each cluster contained at least 120 images, originating from at least 46 unique patients (see Table 1). This indicated that the clusters captured biomarkers that arose independently among multiple patients with AMD, rather than features that were specific to only a few outlier patients. This is also shown in Figure 4 where, in most clusters, images from 10 randomly drawn patients presented consistent features related to AMD that were distinct from other clusters.

This was reinforced by GradCAM attribution maps that, for example, highlighted drusen in C9, hypertransmission in C26, and subretinal hyperreflective material (disciform scars) in C30.

Clusters Describe Known Biomarkers

Both teams of retinal specialists found clusters that distinguished healthy retinas (C1), small drusen with subretinal drusen deposits (C4), confluent drusen (some post anti-VEGF treatment) (C6), and large drusen (C9). Drusen size is an established predictor for the risk of progression to late AMD.^{31,32} Moreover, they identified clusters described as subretinal fluid (C18) and pigment epithelial detachment (C17), which are both indicative of MNV. Finally, images in cluster C19 clearly exhibit atrophy or cRORA. These clusters represent the known biomarkers used in clinical grading systems that currently map the progression of early AMD and its wet and dry variants in late AMD.

Clusters Contain Biomarkers Not Yet in Established Grading Systems

The majority of the clusters distinguished cases that would fall under the same category in established grading systems. As shown in Figure 5, our clusters create a more fine-grained decomposition of the more broadly defined known biomarkers. Moreover, correlating clusters with known disease stages also highlighted a subgroup (C20 with vitreomacular interface abnormalities) that had been mislabeled as healthy. Our clusters differentiated intraretinal fluid (C22) from subretinal fluid (C18). These are typically conflated under neovascular AMD,³³ despite the known association between intraretinal fluid and degraded visual acuity,^{34–36} which was reflected in our own measurements where average visual acuity for C22, 58.7 letters (95% confidence interval [55.9, 61.4]), was worse than in C18, 68.2 letters (95% confidence interval [66.7, 69.8]) (Fig 6).

Other clusters described biomarkers that have only been tentatively linked to the progression of AMD such as DLS (C14), which is associated with subclinical type 1 macular neovascularization,⁷ and vitreomacular interface abnormalities (C20) including epiretinal membrane, which has been shown to increase the growth rate of geographic atrophy.³⁷ Moreover, our clusters reflected the distinction in the severity of atrophy between "incomplete retinal pigment epithelial and outer retinal atrophy" (iRORA) (C12, C15) and cRORA (C24, C26) that was only recently proposed in 2018.²⁷ Clusters also differentiated incomplete retinal pigment epithelial atrophy (C11) from iRORA and cRORA. One cluster was characterized by small fovea involving atrophy and macular holes (C27) and had the worst visual acuity of any cluster not depicting scarring. Fovea-centered atrophy has been shown to have the most deleterious effects on visual acuity.³⁸ The subtlety of these distinctions supports our hypothesis that features learned by contrastive learning can extract fine-grained biomarkers.

Clusters Describe Potentially New Biomarkers

Some clusters were differentiated by the presence of a thick, visible choroid. For example, in contrast to C1, C2 contained healthy retinas but with a thick, permeable choroid. Similarly, C16 showed iRORA with a thick choroid in contrast to C15 exhibiting iRORA with a thin

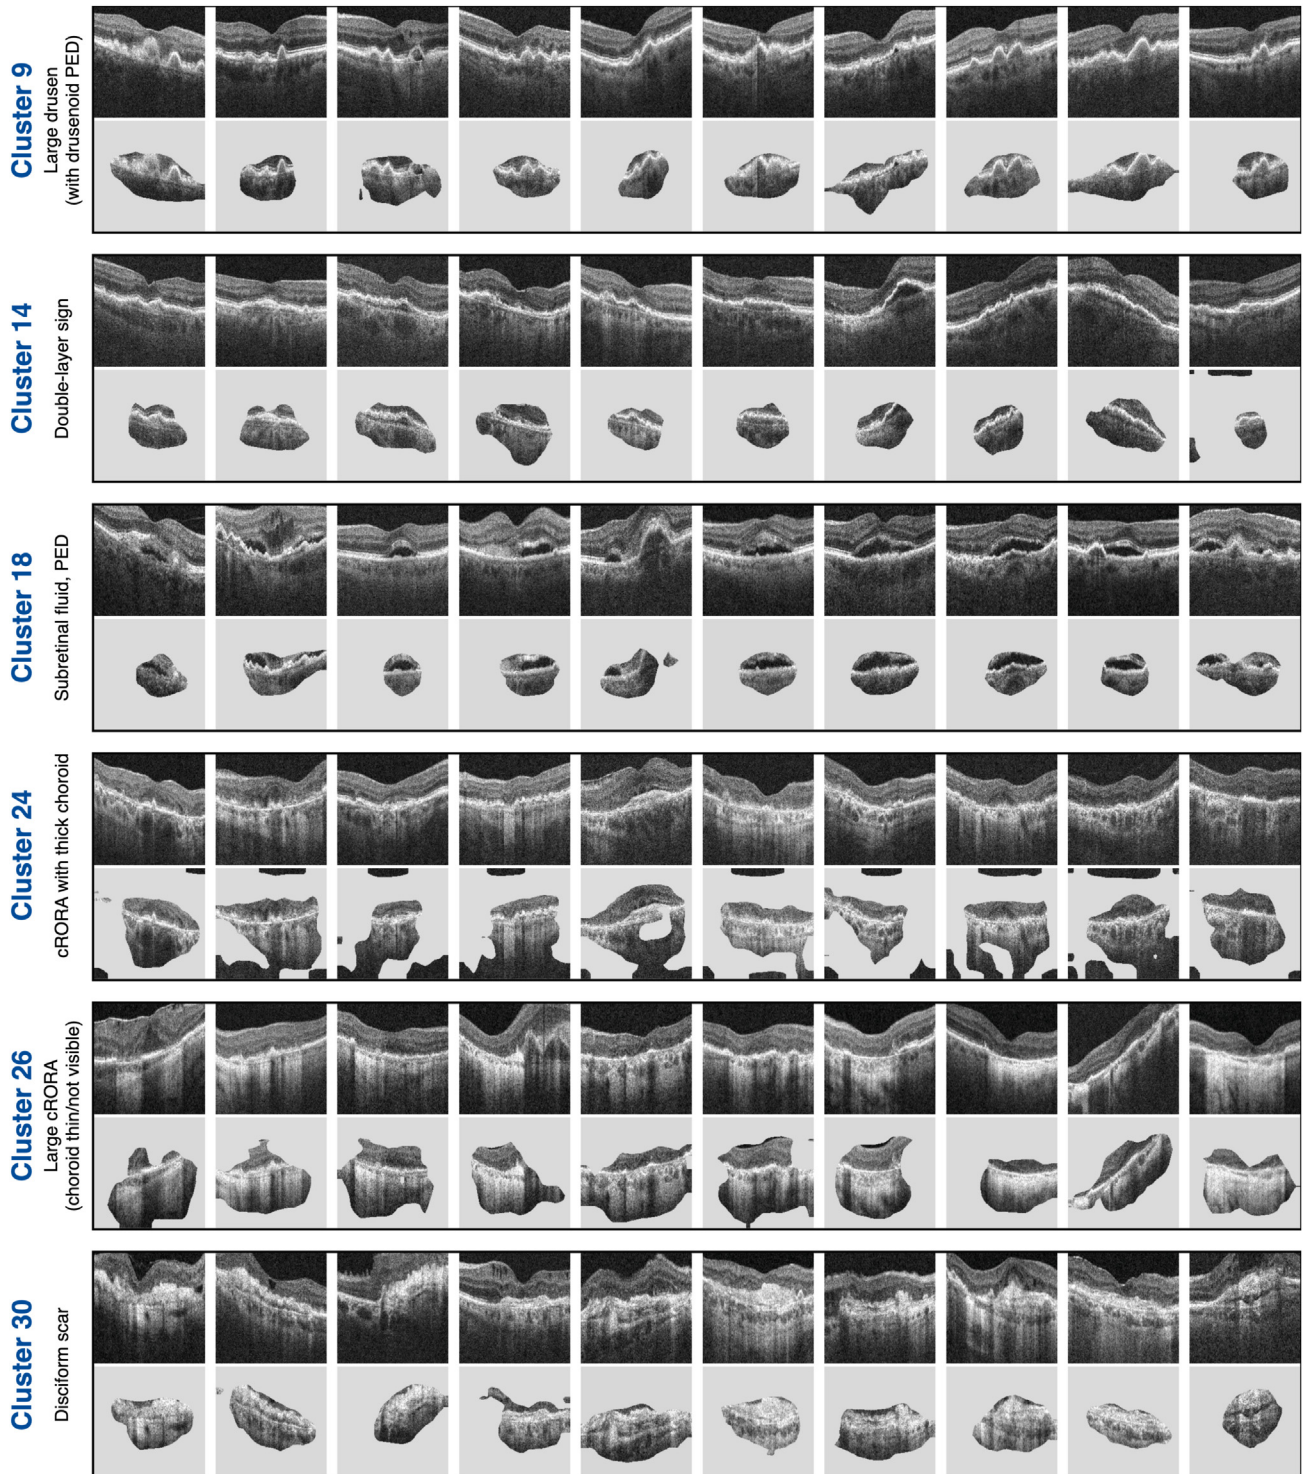


Figure 4. Images from 10 randomly drawn patients from 6 clusters, as shown to retinal specialists during the cluster interpretation interviews. The clusters were largely homogeneous and had identifiable features that described the majority of the images (written to the left). This was reinforced by cluster-specific attribution maps (below each image) that indicate that a consistent set of self-supervised features define each cluster. cRORA = complete retinal pigment epithelial and outer retinal atrophy; PED = pigment epithelial detachment.

choroid. Both teams of retinal specialists were most interested in C24, which showed cRORA but with a thick, visible choroid. This was unexpected, as usually atrophy is accompanied by a thin choroid as in C19 and

C26. Poor choroidal perfusion of the macula is implicated in the pathogenesis of AMD,³⁹ which may indicate that patients in clusters with thinner choroids are at higher risk. Another novel combination includes DLS and

Probability per cluster that each OCT image was labelled with each known biomarker

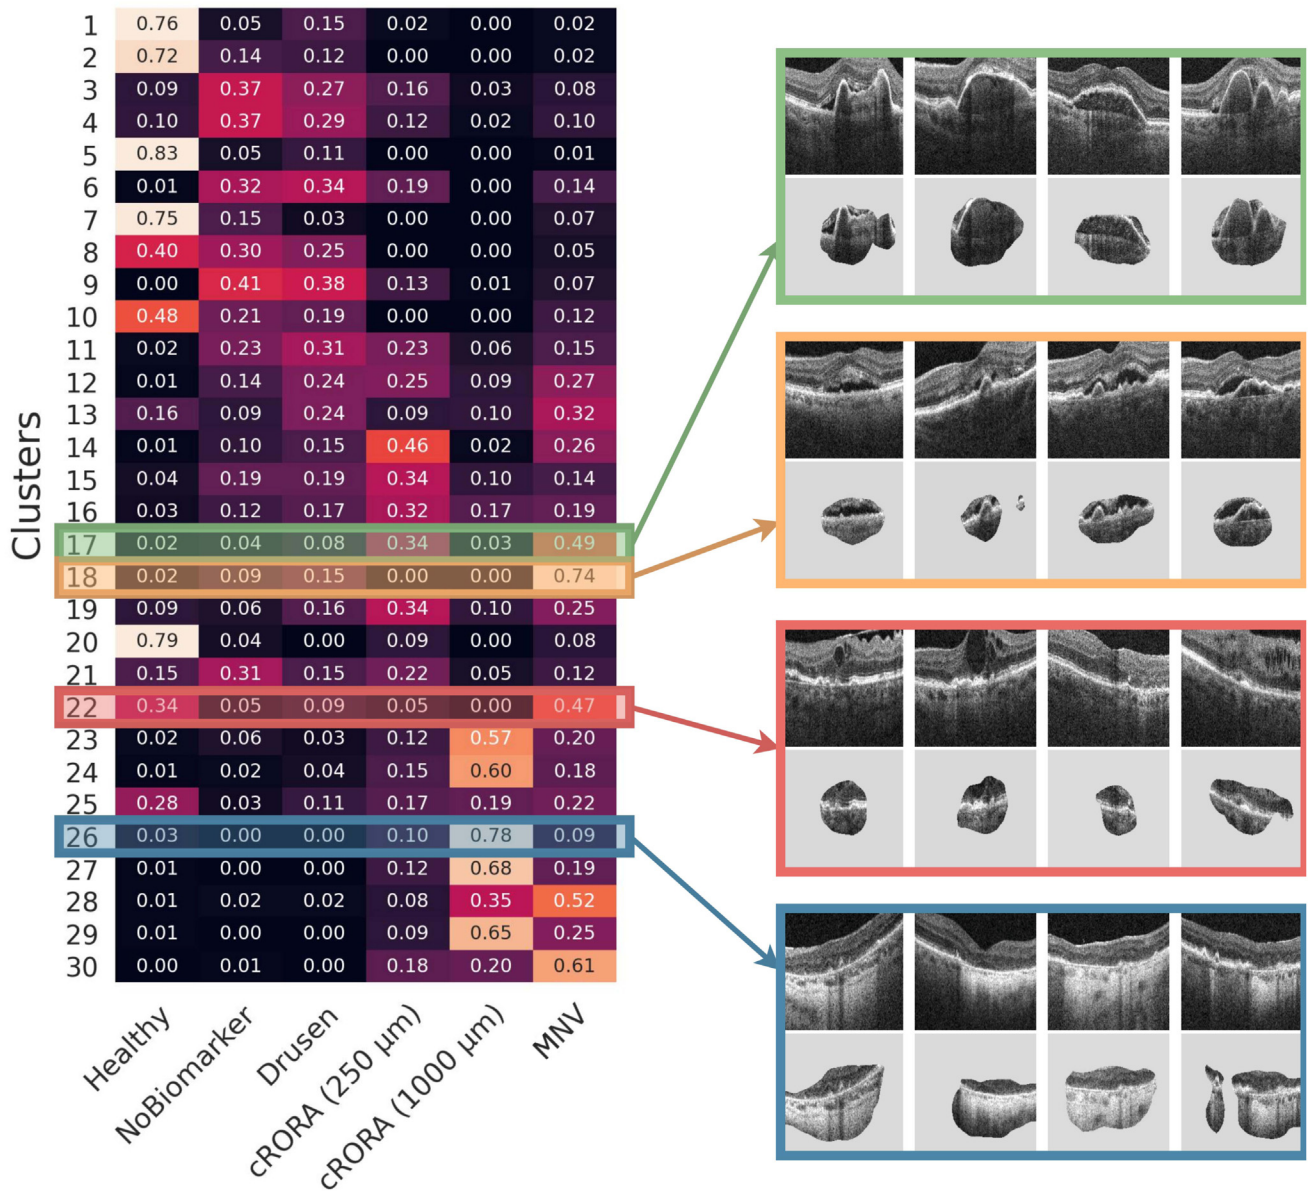


Figure 5. Clusters were correlated with known biomarker annotations and disease stages using conditional probability. By comparing clusters that are indistinguishable from current grading systems but are, by definition, distinguishable by their self-supervised features, we hope to identify new biomarker subtypes that are currently conflated in each indiscriminate disease stage. cRORA = complete retinal pigment epithelial and outer retinal atrophy; MNV = macular neovascularization.

fibrovascular pigment epithelial detachment (C12). These findings evidence the capability of our system to propose biomarker subtypes that had not yet been considered.

Risk Stratification with Clusters

By simulating a grading system using cluster membership, we found that the fine-grained biomarker-related distinctions made by our clusters were more prognostic than the broad

disease stages in the clinical grading system. This was evidenced by reduced mean absolute error in predicting time until conversion to late wet AMD and late dry AMD (Table 2). Notably, the current grading system (defined earlier) did not stratify patient risk much better than the patient’s age and sex. The fully supervised black-box model performed the best, but this approach makes uninterpretable predictions. We found our clusters were especially effective at predicting current visual acuity (11.5 vs. 18.4 mean absolute error). This performance is enabled by

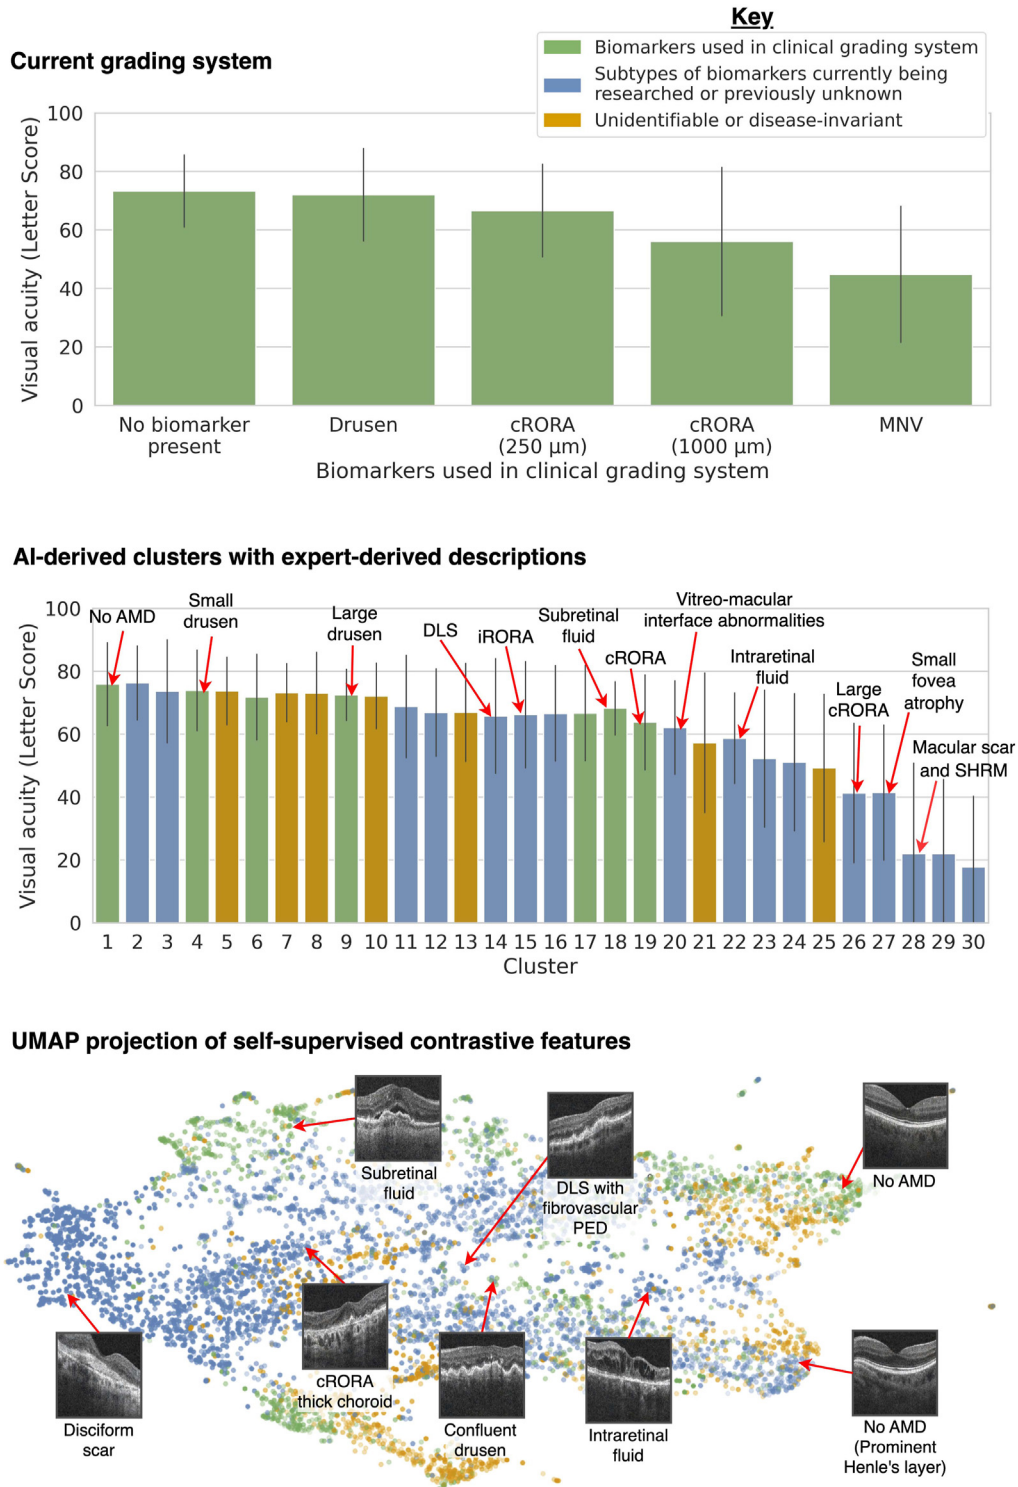


Figure 6. Comparing the stratification of visual acuity (in letter score) by each stage of the clinical grading system cluster (top) to our clusters (middle). Each bar represents the average visual acuity with error bars for standard deviation. Our interpretable clusters provide a significantly improved stratification of degradation to visual acuity compared with the set of known biomarkers. We also use a UMAP projection (bottom) to depict where retinal images reside in the self-supervised contrastive feature space used to create the clusters. AMD = age-related macular degeneration; cRORA = complete retinal pigment epithelial and outer retinal atrophy; DLS = double-layer sign; iRORA = incomplete retinal pigment epithelial and outer retinal atrophy; MNV = macular neovascularization; UMAP = uniform manifold approximation and projection.

Table 2. Clusters Outperformed the Established Clinical Grading Systems for AMD in Predicting Future Conversion to Disease, Shown by Reduced MAE in Years for Late AMD, MNV, and cRORA

System	Time to Late AMD (MAE Yrs) ↓	Time to MNV (MAE Yrs) ↓	Time to cRORA (MAE Yrs) ↓	Current Visual Acuity (MAE Letters) ↓
Demographic	0.76 ± 0.01	0.82 ± 0.01	0.70 ± 0.03	19.1 ± 0.35
Current grading system	0.76 ± 0.01	0.82 ± 0.01	0.69 ± 0.04	18.4 ± 0.40
Clusters	0.75 ± 0.01	0.78 ± 0.02	0.63 ± 0.05	11.5 ± 0.25
Fully supervised	0.71 ± 0.02	0.73 ± 0.01	0.61 ± 0.03	10.0 ± 0.20

We Also Find Clusters Substantially Reduced MAE in Letter Scores for Predicting Visual Acuity. Fully Supervised Machine Learning Outperforms Our Clusters but Is Uninterpretable and Cannot Form Better Grading Systems.

AMD = age-related macular degeneration; cRORA = complete retinal pigment epithelial and outer retinal atrophy; MAE = reduced mean average error; MNV = macular neovascularization.

the use of fovea-centered images that effectively capture the macula. Moreover, we find clusters provide more granular stratification of the degradation in visual acuity (Fig 6). These results reinforce that our clusters are related to disease progression.

Discussion

In this study, we introduced a contrastive learning-based approach to biomarker discovery in AMD. We hypothesized that contrastive learning could reveal new biomarkers by focusing deep learning models on potentially new biomarkers. Through a process of automated self-discovery, our system rediscovered known biomarkers and proposed candidates for potentially new ones. Without using any clinical annotations, we automatically grouped a data set of over 46 000 retinal OCT images into 30 distinct clusters. In only two 1.5-hour interviews, 2 teams of retinal specialists empowered with our tool built a granular taxonomy of the AMD progression in a data set. The teams independently verified that clusters modeled fine-grained distinctions between various subtypes of AMD. They remarked both on the consistency in shared features between images of the same cluster and on the subtlety of the different biomarkers captured by different clusters. In addition to finding tentative biomarkers that are not yet in current grading systems, retinal specialists were surprised to find clusters depicting nuanced and granular subtypes of biomarkers that had not yet been identified in clinical research. These subtypes included new combinations such as thick choroid with hypertransmission, DLS seen with a fibrovascular pigment epithelial detachment, and small fovea involving atrophy. Moreover, in quantitative benchmarks, our clusters provided greater risk stratification and prognostic value than the current clinical grading system. Our results strongly suggest that biomarker proposals powered by contrastive learning can accelerate the initial idea-generation stage of biomarker discovery.

Our approach offers a data-driven and scalable alternative to traditionally anecdotal biomarker proposals. It provides a fast, effective, and inexpensive preliminary review of any medical data set, which can subsequently be refined using fewer resources. It can also enable the analysis of the

compound effect on progression risk posed by enumerable biomarker combinations and subtypes. Our method, from pretraining the feature extractor, to clustering and finally clinical interpretation, is not specific to any disease, anatomical region, or scanning instrument. However, if a user has images acquired using a different domain, such as a change in OCT instrument or pathology of interest, they should first retrain the feature extractor on data from their specific domain before identifying new biomarker proposals. By integrating our tool into existing biomarker discovery pipelines, advances in deep learning can improve patient care without deploying a single black-box neural network to the clinic.

Although our clusters were found to be largely homogeneous, there were instances where not every image within a cluster shared the majority feature. We emphasize that clusters should be treated as biomarker proposals requiring further refinement by clinicians, rather than as definitive biomarkers. They are intended as preliminary suggestions that provide a starting point for further investigation. Consequently, our quantitative study only indicates that a more nuanced set of biomarkers may enhance risk stratification for AMD. This also complicates measuring definitive quantities of our clusters, such as the minimum width of atrophy (such as cRORA > 1000 μ m). However, all the aforementioned issues can be addressed by manually refining each cluster.

Moreover, clusters are not guaranteed to be characterized by biomarker-related features. Of the 30 clusters, 4 related to invariant image features, such as scan location, are difficult to model in the set of contrastive transformations and may require a more sophisticated technique to explicitly remove from the set of self-supervised features. For now, these clusters can easily be removed during the review stage and could even be filtered from the data set as a preprocessing step (for example, removing C21 would omit acquisitions not centered on the fovea). However, in another 3 clusters, the retinal specialists were unable to identify any majority feature. This could either indicate the existence of catch-all heterogeneous clusters of miscellaneous images or that both teams of retinal specialists missed potentially new biomarker subtypes. As 2 of these clusters had no obvious signs of AMD, it may also indicate that identifying early stage biomarkers, which produce more subtle structural

deformities, can be more challenging than detecting more pronounced late-stage ones. Improving the chance of detecting new early stage biomarkers may require filtering late-stage patients from the process. A more definitive limitation of our system is that by using 2D fovea-centered B-scans, we potentially miss pathologies in other regions of the retina. This potentially missed pathology may explain some of the error in predicting visual acuity, which was 10 letters using fully supervised learning, equating to 2 lines in a logMAR chart. To incorporate biomarkers outside the macula and enhance the prognostic value of our clusters, we plan to repeat this study using the full 3D OCT volume.

In this work, we introduced a deep learning-based biomarker proposal system powered by self-supervised contrastive learning, capable of identifying biomarker can-

didates in large medical image data sets with minimal human input. Our system, without any prior knowledge of AMD, rediscovered all the known biomarkers. Crucially, it also proposed biomarker candidates that had either only been raised in a small number of reports or had not been yet identified by retinal specialists, to the best of our knowledge. Ultimately, we envision deep learning-based self-discovery systems can improve patient outcomes by accelerating the rate of biomarker research into this poorly understood disease.

Acknowledgments

The PINNACLE study is funded by a Wellcome Trust Collaborative Award (ref. 210572/Z/18/Z). The authors would like to acknowledge Toby Prevost for his contribution to the acquisition of funding by the Wellcome Trust.

Footnotes and Disclosures

Originally received: November 6, 2023.

Final revision: April 26, 2024.

Accepted: April 26, 2024.

Available online: May 31, 2024. Manuscript no. XOPS-D-23-00285R2.

¹ BioMedIA, Department of Computing, Imperial College London, London, United Kingdom.

² Clinical and Experimental Sciences, Faculty of Medicine, University of Southampton, Southampton, United Kingdom.

³ Institute of Ophthalmology, University College London, London, United Kingdom.

⁴ Moorfields Eye Unit, National Institute for Health Research, London, United Kingdom.

⁵ Laboratory for Ophthalmic Image Analysis, Department of Ophthalmology and Optometry, Medical University of Vienna, Vienna, Austria.

⁶ Christian Doppler Laboratory for Artificial Intelligence in Retina, Christian Doppler Forschungsgesellschaft, Vienna, Austria.

⁷ Institute of Molecular and Clinical Ophthalmology Basel, Basel, Switzerland.

⁸ Department of Ophthalmology, University of Basel, Basel, Switzerland.

⁹ Institute for AI and Informatics in Medicine, School of Computation, Information and Technology, School of Medicine and Health, Technical University Munich, Munich, Germany.

Meeting Presentation: An earlier version of this material has been presented at the 2023 ARVO Annual Meeting, held in New Orleans, LA, USA.

Disclosure(s):

All authors have completed and submitted the ICMJE disclosures form.

The author(s) have made the following disclosure(s):

R.H.: Funding Support — Wellcome Trust Funding via Collaborative Award, “Deciphering AMD by deep phenotyping and machine learning (PINNACLE),” ref. 210572/Z/18/Z.

R.K.: Funding Support — Wellcome Trust Funding via Collaborative Award, “Deciphering AMD by deep phenotyping and machine learning (PINNACLE),” ref. 210572/Z/18/Z.

A.M.H.: Support — Wellcome Trust Funding via Collaborative Award, “Deciphering AMD by deep phenotyping and machine learning (PINNACLE),” ref. 210572/Z/18/Z; UK Biobank (Resources under Application Number 45477); Other financial or non-financial interests — Boehringer Ingelheim (employment).

O.L.: Support — Wellcome Trust Funding via Collaborative Award, “Deciphering AMD by deep phenotyping and machine learning

(PINNACLE)” ref. 210572/Z/18/Z; UK Biobank (Resources under Application Number 45477).

T.R.P.T.: Research support — Medical expertise, interpretation of results, manuscript reviewing and editing.

H.B.: Grants or contracts — Apellis, Heidelberg Engineering; Payment or honoraria for lectures — Bayer, Apellis.

U.S-E.: Support — Wellcome Trust Funding via Collaborative Award, “Deciphering AMD by deep phenotyping and machine learning (PINNACLE),” ref. 210572/Z/18/Z; UK Biobank (Resources under Application Number 45477); Grants or contracts — Genentech, Roche, Kodiak, Apellis, Novartis, RetInSight; Royalties or licenses — RetInSight; Consulting fees — Apellis.

H.P.N.S.: Grants or contracts — the Swiss National Science Foundation (Project funding: “Developing novel outcomes for clinical trials in Stargardt disease using structure/function relationship and deep learning” #310030_201165), National Center of Competence in Research Molecular Systems Engineering (NCCR MSE: Molecular Systems Engineering (phase II)) #51NF40-182895), Wellcome Trust (PINNACLE study), Foundation Fighting Blindness Clinical Research Institute (ProgStar study), Co-director of the Institute of Molecular and Clinical Ophthalmology Basel (IOB) (Constituted as a non-profit foundation and receives funding from the University of Basel, the University Hospital Basel, Novartis, and the government of Basel-Stadt); Consulting fees — Boehringer Ingelheim Pharma GmbH & Co (Member of the Scientific Advisory Board), Droia NV (Member of the Scientific Advisory Board), Janssen Research & Development, LLC (Johnson & Johnson) (Member of the Scientific Advisory Board), Okuvision GmbH (Member of the Scientific Advisory Board), ReVision Therapeutics Inc. (Member of the Scientific Advisory Board), Saliogen Therapeutics Inc. (Member of the Scientific Advisory Board), Alnylam Pharmaceuticals Inc. (Consultant), Gerson Lehrman Group Inc. (Consultant), Guidepoint Global, LLC (Consultant), Tenpoint Therapeutics (Consultant); Patents planned, issued or pending — Methods to assess cone availability in patients with vision impairments (U.S.Ser. No. 63/305,616, 2022), Human Cone Photoreceptor Optogenetic Constructs (U.S. Ser. No. 63/231,406, 2021); Participation on a Data Safety Monitoring Board or Advisory Board — Belite Bio (DRAGON trial, NCT05244304; LBS-008-CT02, NCT05266014), F. Hoffmann-La Roche Ltd (VELODROME trial, NCT04657289, DIAGRID trial, NCT05126966; HUTONG trial), ViGeneron (protocol number VG901-2021-A), Novo Nordisk (FOCUS trial, NCT03811561).

D.R.: Support — Wellcome Trust Funding via Collaborative Award, “Deciphering AMD by deep phenotyping and machine learning (PINNACLE),” ref. 210572/Z/18/Z.

A.J.L.: Support — Wellcome Trust Funding via Collaborative Award, “Deciphering AMD by deep phenotyping and machine learning (PINNACLE),” ref. 210572/Z/18/Z; Consulting fees — Bayer, Apellis, Eyebio, Roche; Support for attending meetings and/or travel — Roche; Leadership or fiduciary role in other board, society, committee or advocacy group — Royal College of Ophthalmologists.

S.S.: Support — Novartis (Research grant to institution); Grants or contracts — Boehringer Ingelheim, Bayer (Research grant to the institution); Consulting fees — Roche, AbbVie, Apellis, Bayer, Biogen, Boehringer Ingelheim, Novartis, Janssen Pharmaceuticals, Optos, Ocular Therapeutix, OcuTerra, Sanofi, Eyepoint, Amgen, Stealth Biotherapeutics; Payment or honoraria for lectures, presentations, speakers bureaus, manuscript writing or educational events — Boehringer Ingelheim, Bayer; Support for attending meetings and/or travel — Bayer, Roche; Participation on a Data Safety Monitoring Board or Advisory Board — Bayer, Nova Nordisk; Leadership or fiduciary role in other board, society, committee or advocacy group — Trustee Macular Society, Chair RCOphth Scientific Committee, Editor in Chief, EYE; Stock or stock options — Eyebio.

M.J.M.: Support — Wellcome Trust Funding via Collaborative Award, “Deciphering AMD by deep phenotyping and machine learning (PINNACLE),” ref. 210572/Z/18/Z.

Funded by the Wellcome Trust (London, United Kingdom) in the scope of the Wellcome Trust Collaborative Award, “Deciphering AMD by deep phenotyping and machine learning” Ref. 210572/Z/18/Z. The funding organization had no role in the design or conduct of this research.

HUMAN SUBJECTS: Humans subjects were included in this study. The PINNACLE study ([ClinicalTrials.gov](https://clinicaltrials.gov/ct2/show/study/NCT04269304) NCT04269304) protocol was approved in the United Kingdom by the East Midlands–Leicester Central

Research Ethics Committee (ref. 19/EM/0163) and further by the institutional review boards of all involved institutions. It adheres to the principles of Good Clinical Practice and is in accordance with the Declaration of Helsinki. Informed consent was obtained from the participants after an explanation of the nature and possible consequences of the study.

No animal subjects were used in this study.

Author Contributions:

Conception and design: Holland, Rueckert, Menten

Data collection: Holland, Kaye, Hagag, Leingang, Bogunović, Schmidt-Erfurth, Scholl, Rueckert, Lotery, Sivaprasad

Analysis and interpretation: Holland, Kaye, Hagag, Taylor, Lotery, Sivaprasad, Menten

Obtained funding: Schmidt-Erfurth, Scholl, Rueckert, Lotery, Sivaprasad

Overall responsibility: Holland, Bogunović, Rueckert, Lotery, Sivaprasad, Menten

Abbreviations and Acronyms:

AMD = age-related macular degeneration; **cRORA** = complete retinal pigment epithelial and outer retinal atrophy; **DLS** = double-layer sign; **iRORA** = incomplete retinal pigment epithelial and outer retinal atrophy; **MNV** = macular neovascularization.

Keywords:

AMD, Contrastive learning, Biomarker discovery, Deep learning, OCT.

Correspondence:

Robbie Holland, MEng, Imperial College London, South Kensington Campus, SW7 2AZ, London, United Kingdom. E-mail: robert.holland15@imperial.ac.uk

References

1. Wong WL, Su X, Li X, et al. Global prevalence of age-related macular degeneration and disease burden projection for 2020 and 2040: a systematic review and meta-analysis. *Lancet Glob Health*. 2014;2:e106–e116.
2. Bird AC, Bressler NM, Bressler SB, et al. An international classification and grading system for age-related maculopathy and age-related macular degeneration. The International ARM Epidemiological Study Group. *Surv Ophthalmol*. 1995;39:367–374.
3. Klein R, Meuer SM, Myers CE, et al. Harmonizing the classification of age-related macular degeneration in the three-continent AMD consortium. *Ophthalmic Epidemiol*. 2014;21:14–23.
4. Ferris FL, Davis MD, Clemons TE, et al. A simplified severity scale for age-related macular degeneration: AREDS Report No. 18. *Arch Ophthalmol*. 2005;123:1570–1574.
5. Ferris III FL, Wilkinson CP, Bird A, et al. Clinical classification of age-related macular degeneration. *Ophthalmology*. 2013;120:844–851.
6. McDermott JE, Wang J, Mitchell H, et al. Challenges in biomarker discovery: combining expert insights with statistical analysis of complex omics data. *Expert Opin Med Diagn*. 2013;7:37–51.
7. Shi Y, Motulsky EH, Goldhardt R, et al. Predictive value of the OCT double-layer sign for identifying subclinical neovascularization in age-related macular degeneration. *Ophthalmol Retina*. 2019;3:211–219.
8. Robison CD, Krebs I, Binder S, et al. Vitreomacular adhesion in active and end-stage age-related macular degeneration. *Am J Ophthalmol*. 2009;148:79–82.e2.
9. Fragiotta S, Abdolrahimzadeh S, Dolz-Marco R, et al. Significance of hyperreflective foci as an optical coherence tomography biomarker in retinal diseases: characterization and clinical implications. *J Ophthalmol*. 2021;2021:6096017.
10. Abramson RG, Burton KR, Yu J-PJ, et al. Methods and challenges in quantitative imaging biomarker development. *Acad Radiol*. 2015;22:25–32.
11. Kessler LG, Barnhart HX, Buckler AJ, et al. The emerging science of quantitative imaging biomarkers terminology and definitions for scientific studies and regulatory submissions. *Stat Methods Med Res*. 2015;24:9–26.
12. Age-Related Eye Disease Study Research Group. The Age-Related Eye Disease Study system for classifying age-related macular degeneration from stereoscopic color fundus photographs: the Age-Related Eye Disease Study Report Number 6: the Age-Related Eye Disease Study Report No. 6. *Am J Ophthalmol*. 2001;132:668–681.
13. Alyass A, Turcotte M, Meyre D. From big data analysis to personalized medicine for all: challenges and opportunities. *BMC Med Genomics*. 2015;8:33.
14. Cowan N. The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav Brain Sci*. 2001;24:87–185.
15. Miller GA. The magical number seven plus or minus two: some limits on our capacity for processing information. *Psychol Rev*. 1956;63:81–97.
16. Schmidt-Erfurth U, Sadeghipour A, Gerendas BS, et al. Artificial intelligence in retina. *Prog Retin Eye Res*. 2018;67:1–29.

17. Gidaris S, Singh P, Komodakis N. Unsupervised representation learning by predicting image rotations. *arXiv*. 2018;1803:07728.
18. Zhang R, Isola P, Efros AA. Colorful image colorization in European conference on computer vision. In: Leibe B, Matas J, Sebe N, Welling M, eds. *Computer Vision – ECCV 2016*. Cham: Springer; 2016:9907. Lecture Notes in Computer Science.
19. Oord AV, Li Y, Vinyals O. Representation learning with contrastive predictive coding. *arXiv*. 2018;1807:03748.
20. Chen T, Kornblith S, Norouzi M, Hinton G. A simple framework for contrastive learning of visual representations. In: *Proceedings of the 37th International Conference on Machine Learning*. Vienna, Austria: PMLR; 2020:119.
21. He K, Fan H, Wu Y, Xie S, Girshick R. Momentum contrast for unsupervised visual representation learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020:9729–9738.
22. Grill JB, Strub F, Alché F, et al. Bootstrap your own latent-a new approach to self-supervised learning. *Adv Neural Inf Process Syst*. 2020;33:21271–21284.
23. Kingma DP, Welling M. Auto-encoding variational Bayes. *arXiv*. 2013;1312:6114.
24. Seeböck P, Waldstein SM, Klimescha S, et al. Unsupervised identification of disease marker candidates in retinal OCT imaging data. *IEEE Trans Med Imaging*. 2018;38:1037–1047.
25. Schlegl T, Seeböck P, Waldstein SM, et al. f-AnoGAN: Fast unsupervised anomaly detection with generative adversarial networks. *Med Image Anal*. 2019;54:30–44.
26. Sutton J, Menten MJ, Riedl S, et al. Developing and validating a multivariable prediction model which predicts progression of intermediate to late age-related macular degeneration-the PINNACLE trial protocol. *Eye (Lond)*. 2023;37:1275–1283.
27. Sadda SR, Guymer R, Holz FG, et al. Consensus definition for atrophy associated with age-related macular degeneration on OCT: classification of atrophy Report 3. *Ophthalmology*. 2018;125:537–548.
28. Holland R, Leingang O, Bogunović H, et al. Metadata-enhanced contrastive learning from retinal optical coherence tomography images. *arXiv*. 2022;2208:02529.
29. 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*. Las Vegas, NV: IEEE; 2016:770–778.
30. Selvaraju RR, Cogswell M, Das A, et al. Grad-Cam: visual explanations from deep networks via gradient-based localization. *Proceedings of the IEEE International Conference on Computer Vision*. Venice, Italy: IEEE; 2017:618–626.
31. Abdelfattah NS, Zhang H, Boyer DS, et al. Drusen volume as a predictor of disease progression in patients with late age-related macular degeneration in the fellow eye. *Invest Ophthalmol Vis Sci*. 2016;57:1839–1846.
32. Schlanitz FG, Baumann B, Kundi M, et al. Drusen volume development over time and its relevance to the course of age-related macular degeneration. *Br J Ophthalmol*. 2017;101:198–203.
33. Mitchell P, Liew G, Gopinath B, Wong TY. Age-related macular degeneration. *Lancet*. 2018;392:1147–1159.
34. Jaffe GJ, Martin DF, Toth CA, et al. Macular morphology and visual acuity in the comparison of age-related macular degeneration treatments trials. *Ophthalmology*. 2013;120:1860–1870.
35. Sharma S, Toth CA, Daniel E, et al. Macular morphology and visual acuity in the second year of the comparison of age-related macular degeneration treatments trials. *Ophthalmology*. 2016;123:865–875.
36. Kaiser PK, Wyckoff CC, Singh RP, et al. Retinal fluid and thickness as measures of disease activity in neovascular age-related macular degeneration. *Retina*. 2021;41:1579–1586.
37. Grunwald JE, Pistilli M, Ying G-S, et al. Growth of geographic atrophy in the comparison of age-related macular degeneration treatments trials. *Ophthalmology*. 2015;122:809–816.
38. Bagheri S, Koulouri I, Saad L, Vavvas DG. Foveal vs. total geographic atrophy as a predictor of visual acuity in AMD. *Invest Ophthalmol Vis Sci*. 2017;58:3400–3400.
39. Mori K, Gehlbach PL, Ito YN, Yoneya S. Decreased arterial dye-filling and venous dilation in the macular choroid associated with age-related macular degeneration. *Retina*. 2005;25:430–437.