THEORETICAL ARTICLE

A proposed theoretical framework for retinal biomarkers

Ian James Callum MacCormick[1] | **Bo Zhang[2]** | **Daniel Hill[1,3]** |
Maria Francesca Cordeiro[3,4,5] | **Dylan S. Small[2]**

[1]Centre for Inflammation Research, University of Edinburgh, The Queen's Medical Research Institute, Edinburgh BioQuarter, Edinburgh, UK

[2]Wharton Department of Statistics and Data Science, University of Pennsylvania, Philadelphia, Pennsylvania, USA

[3]UCL Institute of Ophthalmology, London, UK

[4]Western Eye Hospital Imperial College Healthcare NHS Trust, London, UK

[5]ICORG, Imperial College, London, UK

Correspondence
Ian MacCormick, Centre for Inflammation Research, University of Edinburgh, The Queen's Medical Research Institute, Edinburgh BioQuarter, 47 Little France Crescent, Edinburgh, EH16 4TJ, UK.
E-mail: imaccorm@ed.ac.uk

Abstract

Objective: Propose a theoretical framework for retinal biomarkers of Alzheimer's disease (AD).

Background: The retina and brain share important biological features that are relevant to AD. Developing retinal biomarkers of AD is a strategic priority but as yet none have been validated for clinical use. Part of the reason may be that fundamental inferential assumptions have been overlooked. Failing to recognize these assumptions will disadvantage biomarker discovery and validation, but incorporating them into analyses could facilitate translation.

New theory: The biological assumption that a disease causes analogous effects in the brain and retina can be expressed within a Bayesian network. This allows inferences about abstract theory and individual events, and provides an opportunity to falsify the foundational hypothesis of retina–brain analogy. Graphical representation of the relationships between variables simplifies comparison between studies and facilitates judgements about whether key assumptions are valid given the current state of knowledge.

Major challenges: The framework provides a visual approach to retinal biomarkers and may help to rationalize analysis of future studies. It suggests possible reasons for inconsistent results in existing literature on AD biomarkers.

Linkage to other theories: The framework can be modified to describe alternative theories of retinal biomarker biology, such as retrograde degeneration resulting from brain disease, and can incorporate confounding factors such as co-existent glaucoma or macular degeneration. Parallels with analogue confirmation theory and surrogate marker validation suggest strengths and weaknesses of the framework that can be anticipated when developing analysis plans.

KEYWORDS
Alzheimer's disease, Bayesian, biomarker, brain disease, cerebral malaria, retina, surrogate, validation

Highlights

- Retinal biomarkers hold great promise for Alzheimer's disease (AD), but none are currently used clinically.
- Assumptions about the cause of retinal and brain changes are often overlooked, and this may disadvantage biomarker discovery and validation.
- We present a new approach to retinal biomarkers that describes cause and effect graphically in a Bayesian network.
- We show how this allows a more complete assessment of how well a biomarker might reflect the brain, and how data from right and left eyes can be used to rule out poor biomarker candidates.

1 | OBJECTIVE

We propose a theoretical framework for retinal biomarkers of Alzheimer's disease (AD) and other brain diseases. It is based on the concept that disease causes analogous effects in the brain and retina, and the framework expresses this mathematically within a Bayesian network. The intention is to leverage this biological assumption to make better inferences, and to facilitate critique of this assumption in light of empirical data. In proposing this framework, we aim to promote new thinking about analytical approaches to retinal biomarkers of AD, and stimulate discussion about how best to use data to discover and validate them. The framework is demonstrated using abstract theory and empirical data to answer clinical questions and evaluate the assumption of analogy between the retina and brain.

2 | BACKGROUND

AD is a serious problem for individuals and their families. It is also a growing challenge to society as the number of cases is increasing each year.[1] Treatments are likely to be most effective early in the course of disease and biomarkers are needed to identify preclinical cases, prognosticate outcome, and predict or measure treatment effect.[2]

Existing reference-standard biomarkers for AD involve assessment of brain amyloid beta (Aβ) protein using positron emission tomography (PET) or assays of cerebrospinal fluid (CSF). Unfortunately, these methods are relatively expensive, difficult, or invasive and so are impractical at the population scale necessary to tackle the growing prevalence of AD.[2]

In this context, the prospect of retinal biomarkers is appealing for two reasons. First, the retina is *like* the brain in ways that are relevant to AD pathogenesis. Second, the retina is *unlike* the brain because it is accessible to high-resolution, non-invasive, and repeated measurements.[3] The retina is the only visible part of the central nervous system (CNS), and this generates hope that visible retinal manifestations of AD reflect similar changes in the brain, and that measuring these will allow us to overcome major challenges in diagnosis, stratification, and monitoring. This has motivated significant efforts to explore associations between the retina and brain in AD (e.g., Snyder et al.,[3] Koronya et al.,[4]), and the development of retinal biomarkers was recently the subject of a special meeting convened by the Alzheimer's Association.[3]

Despite several associations between retinal and cerebral changes in AD, as yet there are no clinically validated retinal biomarkers. Several reasons have been identified[2] including the need for methodological standardization and within-subject longitudinal data. Taking retinal thickness as an example, there has historically been significant variation in optical coherence tomography (OCT) equipment between studies, and study design has been overwhelmingly cross-sectional rather than longitudinal.[6]

However, we suggest there is an additional problem, which is that traditional analytical approaches to prospective retinal biomarkers are not based on an explicit theoretical framework. As a result, they do not account for the nature of the inference being made from the retina to the brain on the basis of common biology, which is an *analogy*.[7] Analogical arguments are common in science, and use known similarities between two objects to support the conclusion that further similarity exists (see Bartha[7] for an extensive review of analogical reasoning).

This has important implications with practical consequences. For example, analyses usually take it for granted that associations between the retina and brain exist because the variable in each site is caused by the same disease process. Indeed, the utility of studying the retina is often justified by a description of salient retinal and brain biology altered by AD.[2,3] An important corollary follows from such descriptions: this biological account does not just involve the retina and brain, but also a causal influence from AD acting symmetrically on these two regions of the CNS. On this account the causal structure between these three factors (retina, brain, causal disease process) is as important as the factors themselves. Similar concepts apply to the validation of surrogate endpoints in clinical trials.[8,9]

Consequently, while associations between measurements of the retina and brain are important, unadjusted associations between a putative retinal proxy and a brain variable (and still less comparisons of unpaired retinal and brain data between case and control groups) cannot show, in principle, whether the retina faithfully indicates the

FIGURE 1 Graphical representation of analogy, adapted from Dardashti et al.[11] Yellow and blue indicate target and source domains, respectively. Green indicates a node common to both domains. A, Independent domains are represented by parallel graphs. A brain disease causes effects on the brain, and a parallel retinal disease causes effects on the retina. The absence of a connection between sides encodes the assumption that any observed association between retina and brain is purely coincidental. There is no known reason to interpret such associations as the result of a deeper commonality. B, Domains have similar features because of a common cause acting on retina and brain. For example, Alzheimer's disease causing amyloid beta (Aβ) protein accumulation in the brain and retina (brain effect, retina effect), measured by positron emission tomography and curcumin fluorescence. Brain Aβ in turn leads to downstream effects (brain symptoms*). Associations between retina and brain domains can be legitimately interpreted as analogous if the graph connects biologically similar entities via a common cause. Formulation as a directed acyclic graph makes these assumptions explicit and estimates can be done using a Bayesian network or structural equation model. *NB the disease node cannot be the same as the symptoms node. For example, a measurement of cognition could only be used once in the graph—either to define disease severity or define outcome

state of the brain in response to a disease. On the other hand, the traditional practice of matching or adjusting for one or more additional variables actually imposes a causal structure on associations between the retina and brain that may not be realistic,[10] is often unacknowledged, and is different from study to study. Clinicians may be unwittingly misled about a biomarker if these assumptions are both obscure and invalid, and in the past, this type of error has led to serious harm.[8]

A further consequence of this causal structure is that within a traditional regression model the association between a perfect retinal biomarker and related brain variable will disappear if stratified by a marker of disease severity (e.g., cognition in AD) because they will be rendered conditionally independent. A similar effect may occur if biomarkers are studied in restricted patient groups in which disease is very mild or very advanced. This may lead to inappropriate rejection of excellent biomarkers.

It is clear that similarities between the retina and brain create the potential for powerful retinal biomarkers of AD. However, neglecting the structure of relationships among AD, brain, and retina will certainly disadvantage attempts to discover and validate them.

3 | PROPOSED NEW ANALYTICAL FRAMEWORK

We propose a theoretical framework for retinal biomarkers that incorporates this biological context within a Bayesian network. It is similar to analogical reasoning from physical models[11] and has parallels with statistical validation of surrogate endpoints in randomized controlled trials.[9]

Analogies have a source domain that is accessible, and a target domain that one would like to know about. Inference is from source to target.[7] In the case of AD the target is the brain, which cannot be directly observed, and the source is the retina, in which potential biomarkers can be measured. For example, curcumin-labelled Aβ plaques have been imaged in vivo in the retinas of AD patients and animal models, and these might reflect similar plaques in the brain.[4,12]

However, it is possible—even common—for spurious associations to exist between domains. In such cases it would be wrong to infer any deeper meaning from the association because, regardless of statistical significance, it is merely an epiphenomenon. This situation is illustrated in Figure 1A. In contrast, a valid analogy must have biologically

plausible horizontal and vertical associations. There must not only be a horizontal association between domains (e.g., eye and brain), but this association must exist between things that are reckoned to be similar entities with respect to relevant biology, and the domains must also be linked by a plausible common premise or cause that has similar effects on each side. For example, in AD Aβ plaques occur in both the retina and brain, and a correlation between these would suggest a relevant horizontal association between domains. But inferring from this that retinal Aβ is clinically *analogous* to brain Aβ requires the assumption that both plausibly result from the same AD disease process acting on different parts of the CNS to a similar degree (Figure 1B). These ideas are discussed under the terms *isomorphism* and *universality argument* by Dardashti et al.[11] and Bartha.[7]

Recognizing that these assumptions are necessary means that they can be stated explicitly and laid open to scientific debate. It also allows them to be expressed in a Bayesian network.

A Bayesian network is a way of displaying variables and their causal relationships to each other. It consists of nodes joined by edges. Nodes represent variables and edges represent conditional dependencies. The combination of nodes and edges constitutes a directed acyclic graph, which shows how probabilities of parent nodes contribute to the probabilities of daughter nodes. It allows researchers to encode causal relationships between variables based on their understanding of the system being modelled. See Greenland et al.,[10] Textor et al.,[13] and Pearl[14] for a further introduction to directed acyclic graphs.

Such a network can be used to update confidence in abstract theory (see Appendix A in supporting information), and address clinical problems,[15] such as:

1) Estimating brain disease: Suppose a patient is positive for retinal and perhaps other biomarkers, what is the chance she has analogous brain pathology?
2) Predicting outcome: Suppose a patient is positive for a retinal biomarker. What is the chance he will have some clinical outcome?
3) Testing biomarker assumptions: Is the analogy between the retina and brain likely to be true for a given biomarker?

We illustrate our approach to these questions using data from children with cerebral malaria[16] and a putative retinal proxy of brain Aβ protein.[17] We discuss testable implications that can indicate when the assumption of a retina–brain analogy is unlikely to be correct, and suggest how the analogy hypothesis can be falsified using data from left and right eyes—which are often much easier to collect than direct measurements of brain disease.

3.1 | Estimating brain disease and predicting outcome

Paediatric cerebral malaria (CM) is relevant because it is paradigmatic of the retina reflecting the brain. CM is an acute hematological parasitic syndrome with high mortality and a very strong biological link between retinal and brain pathology—to the extent that the presence of malarial retinopathy is a formal diagnostic criterion for the most strict defi-

nition of the syndrome.[18,19] The mechanisms linking parasitemia to death in CM are obscure but blood–brain barrier (BBB) breakdown is strongly suspected. Unfortunately, this cannot be measured directly in sub-Saharan Africa, where CM typically occurs. However, because there are pre-existing biological reasons to suppose that retinal and brain manifestations of CM are similar in nature and extent, we can use retinal data—and assumptions about the analogy between retina and brain coded into our graph and probability table—to estimate values for unmeasured nodes. These assumptions are represented in Figure 2 by total body malaria biomass (disease D) being a parent node to BBB breakdown (B) and retinal leakage (RL); and in the accompanying probability table by stating that the conditional probability of B given D (which we cannot measure directly) is *similar* to the conditional probability of RL given D (Table 1). Malaria biomass is approximated by plasma *Plasmodium falciparum* histidine-rich protein 2 (PfHRP2).

How similar must these conditional probabilities be? For the sake of simplicity, we might hypothesize that the retina and brain are affected identically by malaria. In this case the child nodes (B and RL) are logically equivalent with respect to the parent node (D),[21] and therefore the conditional probabilities are exactly the same. However, clinical science is rarely so clear-cut. Less than perfect values may still be informative, and allow for realistic measurement error, or additional unmeasured causal factors. The similarity of the conditional probabilities relating each domain to the parent node can be specified using two sensitivity parameters, $\Gamma 1$ and $\Gamma 0$, which are based on the odds ratio of conditional probabilities between source and target domains (Appendix B in supporting information). $\Gamma 1 = \Gamma 0 = 1$ if the conditional probabilities are identical. Deviation from 1 in either parameter indicates a certain amount of asymmetry, and values that are very far from 1 or markedly inconsistent between studies would suggest the hypothesis of analogy may be completely untenable for a given retinal proxy or marker of disease status. This is similar in principle to the Relative Effect statistic describing surrogate endpoints.[22] For illustration, Figure 2 outlines the idealistic assumption that $\Gamma 1 = \Gamma 0 = 1$. With this we can answer clinical questions, for example:

1) What is the probability of BBB breakdown if retinal leakage is present in either eye? This is 10%, compared to 9% if retinal leakage is not present.
2) What is the probability of death if retinal leakage is observed? This is 12%, assuming that we specify a probability of death given brain leakage of 47%, and death without brain leakage of 8% (based on the frequency of death in the presence of radiological brain swelling;[16] see Appendix C in supporting information for calculations).

3.2 | Testing biomarker assumptions

It is important to note that the divergence from equivalence can be tested if empirical data are obtained for previously unobserved nodes in the graph. Analogical confirmation has been criticized in cosmological settings in which observation of the target domain is

Directed acyclic graph and observed frequencies

Observed probabilities

Calculated probabilities assuming eye and brain are affected identically, compared with measured frequencies

Chance of high malaria biomass (D)

P(D=1)	0.75

Chance of retinal leak given D

| P(RL=1|D=1) | 0.11 |
| P(RL=1|D=0) | 0.06 |

Chance of BBB breakdown (B) given D

P(B|D)	Unmeasured

Chance of outcome (Eb), given B

P(Eb|B)	Unmeasured

Calculated
Assuming the chance of BBB breakdown, given disease, is the same as the chance of retinal leak given disease

P(B|D) = P(RL|D)	0.11
P(B)	0.09
P(B=1|RL=1)	0.1
P(B=1|RL=0)	0.09

Assume that 47% of patients with BBB breakdown die, and 8% without BBB breakdown die

P(Eb=1|B=1)	0.47
P(Eb=1|RL=1)	0.12

		Retinal leakage (RL)	
		0	1
Malaria biomass (D)	0	31	2
	1	90	11
Clinical outcome (Eb)	0	108	7
	1	13	6

FIGURE 2 In this graph the disease process (D) acting on retina and brain may plausibly be captured by a blood marker of total body malaria parasite biomass (PfHRP2). Blood–brain barrier breakdown (B) and visible retinal leakage in the worst affected eye (RL) are postulated to be the result of high levels of malaria parasite, indicated by arrows from the parent to child nodes (B←D→RL). In this example, "high" parasite biomass is arbitrarily defined as the top 75% of the sample (HRP2 >4000 units). Blood–brain barrier breakdown causes death, the clinical outcome (Eb). In addition, blood–brain barrier breakdown, though unseen, is hypothesized to be similar in nature and extent to visible retinal leakage. This assumption is not represented in the graph but can be encoded by forcing the conditional probabilities of brain leakage given disease and retinal leakage given disease to be equal, or at least similar (i.e., P[B|D] ∼ P[RL|D]).

TABLE 1 Values of probabilities, and differences between model estimates and empirically measured frequencies, can be used to critique assumptions inherent in the model. Observed frequencies are from.[16]

	Model estimate	Observed frequency	Comments
P(B)	0.09	0.16	
P(B = 1|D = 1) P(B = 1|D = 0)	0.11 0.06	0.16 0.15	Observed values are similar but small, suggesting that HRP2 does not have a substantial effect on brain swelling in this end of the disease spectrum.
P(RL|D = 1) P(RL|D = 0)	n/a n/a	0.11 0.06	As above, values are small, suggesting that HRP2 does not have a substantial effect on retinal leakage in this sample with severe disease.
P(B = 1|RL = 1) P(B = 1|RL = 0)	0.10 0.09	0.54 0.12	The large difference between the model estimate and observed frequency suggests a problem with the model assumption that the association between B and RL results from PfHRP2 as the common cause. This interpretation is consistent with residual association between B and RL after controlling for D.
P(Eb = 1|RL = 1)	0.12*	0.46	The large difference between estimated and observed frequency suggests that one or more unmeasured components of the estimate are wrong. Since our specification of P(Eb|B) is accurate, P(B|D) must be faulty. This is consistent with the negligible effect of HRP2 on brain swelling and retinal leak (B←D→RL) in rows 2 and 3.

*Based on pre-specified value of 0.47 for P(Eb = 1|B = 1) and 0.08 for P(Eb = 1|B = 0). These are measured frequencies of death given the presence or absence of radiological brain swelling.[20]

Abbreviations: B, blood–brain barrier breakdown; D, high malaria biomass; Eb, clinical outcome; P, probability; RL, retinal leakage; n/a, not applicable.

not possible, because arguably this may assume the very conclusion of analogy that is under debate (see Crowther et al.[23] and further discussion in Evans and Thébault[24]). However, our application to retinal biomarkers is different. Biomarkers are valuable because they are applicable to a large population when the reference standard test is restricted for some reason, but some measurement of brain disease is usually possible even if only on a relatively small sample (e.g., $A\beta$ PET in AD). These valuable data allow associations to be estimated, such as between a retinal biomarker and brain reference standard (with and without adjustment for disease status[9,22]), and between retinal marker and clinical endpoints. However, investigators can also use these data to estimate $\Gamma 1$ and $\Gamma 0$, and use the distance between calculated and prespecified values of Γ to test the assumption of analogy implicit in their biological framework. Comparing values estimated from the graphical model to empirical values from a population sample may suggest limits to how much information the retinal biomarker can give in particular settings, and also point to problems with biological assumptions inherent in the graphical model.

Taking CM as an example, we can use empirical data on radiological brain swelling in a subset of CM patients to represent BBB breakdown. This suggests a measured $\Gamma 1$ of 1.5 (95% confidence interval [CI] 0.88–2.93) and $\Gamma 0$ of 2.4 (0.51–18.33). We can interpret this in the context of a strong observed association between brain swelling and death (odds ratio [OR] 14.0[20]), retinal leakage and death (OR 13.2),[16] and retinal leakage and brain swelling (OR 4.8).[16] These are consistent with a graph structure implying that the retina and brain ought to be affected similarly by this disease.

If we accept the graph structure as correct, comparisons of model estimates and observed frequencies (Table 1) allow us to critique the identification of the nodes—here chiefly raising questions about whether the parent node should be PfHRP2. The literature indicates that PfHRP2 has significant prognostic power in mild and moderate malaria.[25,26] Our data here suggest that in CM, which is at the very severe end of the disease spectrum, PfHRP2 does not adequately explain the pathogenesis of brain swelling and that other pathological processes may be involved. In addition, although Γ is not significantly different from 1, the width of confidence interval indicates the sample is too small to judge this accurately. Divergence from 1 would be consistent with the existence of additional types of retinal leakage, and with brain swelling from other causes such as cytotoxic edema. These point estimates are consistent with the causal pathway between parasite biomass and death being more complex than: parasite biomass → cerebral hemorrhage → brain swelling → death.[14] This illustrates how our framework allows otherwise unarticulated biological assumptions to be specified to make predictions, and perhaps more importantly, critique assumptions in light of additional data.

However, testing the assumption of analogy between the retina and brain need not wait until data are obtained from the brain. Using our framework, this assumption can also be assessed using data from right and left eyes, which can be collected easily. This is because, considered subregions of the CNS, the left and right eyes are far more similar to each other than either eye is to the brain.

It follows that a putative retinal marker of a brain disease such as CM or AD should, at the very least, provide highly similar results for right and left eyes. In contrast, data indicating that right and left eyes are not analogous (e.g., $\Gamma 1$ and $\Gamma 0$ are far from 1, or inconsistent between studies) would therefore suggest that a biomarker is unlikely to accurately reflect the brain. Analysis of right and left eye with respect to pathogenesis cannot prove an analogous impact of that disease on the eye and brain, but it could potentially provide useful evidence *against* this hypothesis. Falsifying this hypothesis may be disappointing, but is extremely useful because it can direct attention to more fruitful biomarkers.

We illustrate this with paired data from a study of retinal hyperspectral imaging in AD, which reported measurements from right and left eyes.[17] This allows us to assess the assumption that each eye is affected analogously by AD for this particular biomarker (Figure 3). These data suggest that there is some symmetry between the probability of one eye being positive given AD, and the fellow eye being positive, given AD. The 95% CI for $\Gamma 1$ and $\Gamma 0$ both contain 1 although they are wide, consistent with a small sample size ($\Gamma 1 = 3.29$ [1 to ∞]; $\Gamma 0 = 0.46$ [0.1–1.35]). This is in the context of an insignificant association between eyes ($P = .17$, Wilcoxon signed rank test). Our analysis therefore suggests that the sample ($n = 25$ subjects) is too small to judge whether hyperspectral imaging gives an analogous signal from each eye within subjects, and so arguably the study is also underpowered to make conclusions about retinal hyperspectral imaging as a biomarker of AD effects on the brain. Comparing large focal leakage in left and right eyes with respect to PfHRP2 gives some context to the scale of Γ that one might expect. Here $\Gamma 1 = 1.17$ (0.72–1.91) and $\Gamma 0 = 1.11$ (0.44–3.54), $n = 214$ subjects with data from both eyes.[16]

4 | MAJOR CHALLENGES ADDRESSED BY THE FRAMEWORK

This approach has several advantages over traditional strategies for analyzing retinal biomarkers, which often allude to common biology between the retina and brain but do not use this explicitly to inform inferences. In our framework connections assumed to exist between variables are clearly visible and can be scrutinized by peers, and testable implications of assumptions can be used to falsify hypotheses about the analogy between the brain and retina for a given disease in a particular population. This sort of approach could help to rationalize the discovery and validation of retinal biomarkers within ongoing or future cohort studies. Although it does not suggest a new biomarker candidate, the framework offers a new perspective to the conventional thinking about biomarkers expressed by the widespread use of hypothesis tests or multivariate models,[6] and highlights analytical issues that may help to explain puzzling results. It therefore addresses a critical challenge confronting AD.[3,27] We have illustrated the approach using a Bayesian analysis of binary variables, but suggest the principles can be applied to other types of data through use of surrogate endpoint statistics.

Directed acyclic graph and observed frequencies

Conditional probabilities

Assessment of analogy

		Control (D=0)	Case (D=1)
Fellow eye	0	12	5
	1	7	10
Study eye	0	15	2
	1	4	13

P(D)	= 30/38 = 0.79
P(SE=1|D=1)	= 13/15 = 0.87
P(FE=1|D=1)	= 10/15 = 0.67
P(SE=0|D=1)	= 2/15 = 0.13
P(FE=0|D=1)	= 5/15 = 0.33
P(SE=1|D=0)	= 4/19 = 0.21
P(FE=1|D=0)	= 7/19 = 0.37
P(SE=0|D=0)	= 15/19 = 0.79
P(FE=0|D=0)	= 12/19 = 0.63

	$\frac{P(SE=1|D=0)/1-P(SE=1|D=0)}{P(FE=1|D=0)/1-P(FE=1|D=0)}$
Γ0	$\frac{0.21/1-0.21}{0.37/1-0.37} = \frac{0.27}{0.59} = $ **0.46**
	In controls (D=0) the study eye is less likely to have the biomarker than the fellow eye
	$\frac{P(SE=1|D=1)/1-P(SE=1|D=1)}{P(FE=1|D=1)/1-P(FE=1|D=1)}$
Γ1	$\frac{0.87/1-0.87}{0.67/1-0.67} = \frac{6.69}{2.03} = $ **3.29**
	In cases (D=1) the study eye is more likely to have the biomarker than the fellow eye

FIGURE 3   Frequency of positive retinal hyperspectral imaging in Alzheimer's disease (AD) patients and controls (Figure 6A in Hadoux et al.[17]) can be used to calculate Γ1 and Γ0. Disease status is defined by amyloid beta burden on positron emission tomography imaging. The directed acyclic graph illustrates the idea that in this case, AD acts on both eyes to cause changes in hyperspectral imaging signal. Results allow the chance of positive imaging in either eye to be calculated, given case or control status. These are conditional probabilities. Here, FE = 1 indicates "fellow eye" is positive; SE = 0 that "study eye" is negative, etc. Γ is calculated as a ratio of conditional probabilities. Values far from 1, or that are inconsistent between studies, would suggest that the retinal measure from one eye did not reflect the fellow eye, and arguably may therefore be even less likely to accurately reflect the brain

Further validation and refinement are necessary to transform the ideas presented here into a practical research tool. We hope this article will stimulate debate about retinal biomarker assessment and ultimately lead to an improved standard of analysis. In the meantime, we suggest that authors reporting results of retinal biomarkers do the following:

1) Analyze variables within subjects (paired data), rather than comparing separate groups (unpaired data).
2) Describe sample selection, or any stratification within the sample, to indicate where their subjects sit within the spectrum of disease severity.
3) Include a directed acyclic graph to illustrate the proposed connections among retinal, brain, and disease variables; and explain why the retinal and brain markers in question are biologically likely to represent similar manifestations of the same disease. Graphs can be drawn easily in DAGitty (http://dagitty.net/).[13]
4) Provide a comparison of biomarker status between eyes given disease (e.g., Γ0) and given the absence of disease (e.g., Γ1) with confidence intervals. This is similar to the relative effect statistic for clinical trials.[22]
5) Report the unadjusted association between biomarker signal from left and right eyes, and also the association adjusted for disease severity. This would help indicate if the assigned parent node did

indeed render the child nodes conditionally independent, as would be expected in the case of analogous downstream effects (cf. the adjusted association statistic[22]).

5 | LINKS TO OTHER THEORIES

We have described this approach principally in terms of a mechanistic theory in which changes in the retina and brain result from a common cause (brain ← disease → retina). This is a prominent theory in AD research,[2,3] but an alternative is that retinal changes result from brain changes (disease → brain → retina). This may be relevant to some prospective retinal biomarkers of AD, particularly measures of retinal thickness, which could result from retrograde degeneration of the optic nerve.[6] These theoretical assumptions could also be described by a modified Bayesian network. While we do not go into detail here, measurement of the brain variable would make disease and retina variables conditionally independent, and the analogy assumption would not apply.

Another important consideration is the role of alternative causes of retinal changes, such as age-related macular degeneration and glaucoma.[2,3] Like AD, these have strong associations with age and so may represent confounders of a causal effect on the retina. These possibilities can be included by adding additional nodes and edges to the

directed acyclic graph. Similarly, measurement error can be included in a Bayesian network by adding a node representing the measured variable as a daughter of the node for the "true" random variable.

Parallels between our framework and literature on analogical reasoning and surrogacy suggest consistency with existing theories.[8,11,22] These parallels also highlight potential pitfalls. Our framework depends on assumptions about a common cause with proportionate effects on source and target domains. It is therefore crucial to select a robust measure of disease severity acting on both retina and brain (cf. universality), and to select markers of retina and brain changes that have legitimate biological parallels (cf. isomorphism).[11] An imperfect measurement of causal disease status could lead to paradoxical results, such as if the marker of AD disease somehow describes aspects of pathogenesis that act more severely on the brain than the retina, or vice versa.[9]

As well as completely capturing disease effect, the ideal causal disease variable should be randomized to prevent confounding of relationships: disease→brain, and disease→retina.[9] Experimental randomization is not possible in observational studies, but an instrumental variable may provide some of these advantages.[28,29]

The best use of currently validated biomarkers such as PET or CSF Aβ protein should be carefully considered, because these could be thought of either as measurements defining AD disease or as markers of the effect of AD pathogenesis. They cannot fill both roles at the same time.

Finally, use of paired data is essential to avoid a lack of transitivity, which can obscure associations between retina and brain when the value of a marker in each location is consistently high or low for a group on average but not on an individual level [8,9].

Considering these points, it should be noted that many also apply to traditional statistical approaches to biomarkers as well.[9] In practice hypothesis tests or multiple regression are often treated as if they are impartial; but in reality, they effectively encode many causal assumptions.[10] A key advantage of a graphical model is the explicit display of these assumptions, allowing discussion of their propriety. Furthermore, as well as formulating it mathematically, a graphical model allows the analogy assumption to be tested in light of empirical data about downstream effects. This goes far beyond narrative descriptions of the eye as a window to the brain—it allows the transparency of this metaphorical window to be judged on a case-by-case basis. As a result, perhaps we will be able to see further through it.

GLOSSARY OF SOME TERMS RELEVANT TO BAYESIAN NETWORKS (SEE Greenland et al.[10] FOR DETAILS)

Node: A node represents a variable in a Bayesian network. Variables can be observed or unobserved (latent), numeric, or Boolean (true/false).

Edge: An edge connects two nodes. Edges have a direction—upstream nodes are ancestors or parents and downstream nodes are descendants or children. A variable X is an ancestor or cause of

another variable Y if there is a directed path of arrows leading out of X into Y. Edges represent conditional dependencies between nodes. Unconnected nodes are conditionally independent of each other.

Conditional dependency: Refers to the probability of an event occurring, given that another event has already occurred. P(X) means "the probability of X," and P(X|Y) means "the probability of X given that Y has occurred."

Probability function: Gives the probability that an outcome will occur, for every possible outcome. In a Bayesian network each node has a probability function. Values are dictated by parent nodes, but can also be assigned according to expert opinion or empirical measurements. A substantial inconsistency in the probability function derived logically from the network and empirical measurements would suggest an error in measurement or graph structure.

CONFLICTS OF INTEREST

All authors declare no conflicts of interest. Author disclosures are available in the supporting information.

REFERENCES

1. Brookmeyer R, Abdalla N, Kawas CH, Corrada MM. Forecasting the prevalence of preclinical and clinical Alzheimer's disease in the United States. *Alzheimers Dementia*. 2018;14:121-129. https://doi.org/10.1016/j.jalz.2017.10.009
2. Alber J, Goldfarb D, Thompson LI, et al. Developing retinal biomarkers for the earliest stages of Alzheimer's disease: what we know, what we don't, and how to move forward. *Alzheimers Dementia*. 2020;16:229-243. https://doi.org/10.1002/alz.12006
3. Snyder PJ, Alber J, Alt C, et al. Retinal imaging in Alzheimer's and neurodegenerative diseases. *Alzheimers Dementia*. 2021;17:103-111. https://doi.org/10.1002/alz.12179
4. Koronyo Y, Biggs D, Barron E, et al. Retinal amyloid pathology and proof-of-concept imaging trial in Alzheimer's disease. *JCI Insight*. 2017;2:e93621. https://doi.org/10.1172/jci.insight.93621
5. Cordeiro MF, Normando EM, Cardoso MJ, et al. Real-time imaging of single neuronal cell apoptosis in patients with glaucoma. *Brain*. 2017;140:1757-1767. https://doi.org/10.1093/brain/awx088
6. Chan VTT, Sun Z, Tang S, et al. Spectral-domain OCT measurements in Alzheimer's disease: a systematic review and meta-analysis. *Ophthalmology*. 2019;126:497-510. https://doi.org/10.1016/j.ophtha.2018.08.009
7. Bartha P, Analogy and Analogical Reasoning. In: Zalta EN, ed. *The Stanford Encyclopedia of Philosophy*. Spring; 2019. https://plato.stanford.edu/archives/sum2022/entries/reasoning-analogy/
8. VanderWeele TJ. Surrogate measures and consistent surrogates. *Biometrics*. 2013;69:561-565. https://doi.org/10.1111/biom.12071
9. MacCormick IJC, Czanner G, Faragher B. Developing retinal biomarkers of neurological disease: an analytical perspective. *Biomark Med*. 2015;9:691-701. https://doi.org/10.2217/bmm.15.17
10. Greenland S, Pearl J, Robins JM. Causal diagrams for epidemiologic research. *Epidemiology*. 1999;10:37-48.

11. Dardashti R, Hartmann S, Thébault K, Winsberg E. Hawking radiation and analogue experiments: a Bayesian analysis. *Stud. Hist. Philos. Sci. B.* 2019;67:1-11. https://doi.org/10.1016/j.shpsb.2019.04.004

12. Sidiqi A, Wahl D, Lee S, et al. In vivo retinal fluorescence imaging with curcumin in an Alzheimer mouse model. *Front Neurosci.* 2020;14:1–13. https://doi.org/10.3389/fnins.2020.00713

13. Textor J, van der Zander B, Gilthorpe MS, Liśkiewicz M, Ellison GT. Robust causal inference using directed acyclic graphs: the R package 'dagitty'. *Int J Epidemiol.* 2016;45:1887-1894. https://doi.org/10.1093/ije/dyw341

14. Pearl J. *The Book of Why: The New Science of Cause and Effect Judea Pearl and Dana Mackenzie.* Allen Lane; 2018.

15. Lauritzen SL, Spiegelhalter DJ. Local computations with probabilities on graphical structures and their application to expert systems. *J R Stat Soc (Series B).* 1988;50:157-224.

16. MacCormick IJC, Barrera V, Beare NAV, et al. How does blood-retinal barrier breakdown relate to death and disability in pediatric cerebral malaria? *J Infect Dis.* 2022;225:1070-1080. https://doi.org/10.1093/infdis/jiaa541

17. Hadoux X, Hui F, Lim JKH, et al. Non-invasive in vivo hyperspectral imaging of the retina for potential biomarker use in Alzheimer's disease. *Nat Commun.* 2019;10:4227. https://doi.org/10.1038/s41467-019-12242-1

18. Severe malaria. *Trop Med Int Health.* 2014;19:7-131. https://doi.org/10.1111/tmi.12313_2

19. MacCormick IJ, Beare NA, Taylor TE, et al. Cerebral malaria in children: using the retina to study the brain. *Brain.* 2014;137:2119-2142.

20. Seydel KB, Kampondeni SD, Valim C, et al. Brain Swelling and death in children with cerebral malaria. *N Engl J Med.* 2015;372:1126-1137. https://doi.org/10.1056/NEJMoa1400116

21. Talbott W. Bayesian epistemology. In: Zalta EN, ed. *The Stanford Encyclopedia of Philosophy.* Winter; 2016.

22. Buyse M, Molenberghs G. Criteria for the validation of surrogate endpoints in randomized experiments. *Biometrics.* 1998;54:1014-1029.

23. Crowther K, Linnemann NS, Wüthrich C. What we cannot learn from analogue experiments. *Synthese.* 2019:1-26. https://doi.org/10.1007/s11229-019-02190-0

24. Evans PW, Thébault KPY. On the limits of experimental knowledge. *Philos Trans A Math Phys Eng Sci.* 2020;378:20190235. https://doi.org/10.1098/rsta.2019.0235

25. Fox LL, Taylor TE, Pensulo P, et al. Histidine-rich protein 2 plasma levels predict progression to cerebral malaria in malawian children with plasmodium falciparum infection. *J Infect Dis.* 2013;208:500-503. https://doi.org/10.1093/infdis/jit176

26. Seydel KB, Fox LL, Glover SJ, et al. Plasma concentrations of parasite histidine-rich protein 2 distinguish between retinopathy-positive and retinopathy-negative cerebral malaria in Malawian children. *J Infect Dis.* 2012;206:309-318.

27. Khachaturian AS, Hayden KM, Mielke MM, et al. New thinking about thinking, part two. Theoretical articles for Alzheimer's & Dementia. *Alzheimers Dementia.* 2018;14:703-706. https://doi.org/10.1016/j.jalz.2018.05.002

28. Zhang H, Qin J, Berndt SI, et al. On Mendelian randomization analysis of case-control study. *Biometrics.* 2020;76:380-91. https://doi.org/10.1111/biom.13166

29. Howey R, Shin S-Y, Relton C, Smith GD, Cordell HJ. Bayesian network analysis incorporating genetic anchors complements conventional Mendelian randomization approaches for exploratory analysis of causal relationships in complex data. *PLos Genet.* 2020;16:e1008198. https://doi.org/10.1371/journal.pgen.1008198

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.