Neuroethics (2023) 16:25
https://doi.org/10.1007/s12152-023-09531-4

ORIGINAL PAPER

Brain age Prediction and the Challenge of Biological Concepts of Aging

Jan-Hendrik Heinrichs

Received: 23 February 2023 / Accepted: 11 September 2023 / Published online: 23 September 2023

Abstract Brain age prediction is a relatively new tool in neuro-medicine and the neurosciences. In research and clinical practice, it finds multiple use as a marker for biological age, for general health status of the brain and as an indicator for several brain-based disorders. Its utility in all these tasks depends on detecting outliers and thus *failing* to correctly predict chronological age. The indicative value of brain age prediction is generated by the gap between a brain's chronological age and the predicted age, the brain age gap (BAG). This article shows how the clinical and research use of brain age prediction tacitly pathologizes the states that it is sensitive to. It will be argued that the tacit character of this transformation conceals the need for its explicit justification.

Keywords Brain Age Prediction · Aging · Biomarker · Pathologisation

J.-H. Heinrichs (✉)
Forschungszentrum Jülich, Institute for Neuroscience and Medicine 7, Brain and Behaviour, Jülich, Germany
e-mail: J.Heinrichs@fz-juelich.de

J.-H. Heinrichs
Faculty of Arts and Humanities, Rheinisch-Westfälische Technische Hochschule Aachen, Aachen, Germany

Introduction

Brain age prediction is advertised as a novel and powerful tool for neurology, psychiatry and beyond. The promise associated with this technique is to gain clinically relevant information about patients' health status from the comparison of their brain-imagery with that of others at the same chronological age. Admittedly, mere comparison of brain imagery to a chronological age-corrected standard is not novel [for earlier comparisons see for example [1]. What is novel about the contemporary technique is the scale at which it is carried out, the scale of data used for comparison and the number of criteria of comparison [2, 3].

Brain age prediction uses a machine learning system trained on imaging data and age information of a training group to predict the age of individuals based on their imaging data. The main goal of generating an age prediction is to calculate the so-called brain age gap (*BAG*) or brain age delta, the difference between predicted and chronological age.

This rough sketch can be filled out by many different procedures employing different training groups, imaging modalities, machine learning approaches etc. [4, 5]. These differences play a major role for the direction of future development, as will be discussed below.

The fact that one of the most useful outputs of brain age prediction is the *brain age gap* might at first look slightly paradoxical. This is because brain age prediction – counter to what the name might lead

one to expect – is mostly used to do something different than predicting chronological age. It differentiates between and draws the relation between a set of biological states – the brain age – and chronological age. In order to see why the apparent paradox that the gap between true chronological age and predicted age is most informative isn't really that paradoxical, one needs to track how and for what purpose these two 'ages' a being related to each other.

The use of brain age prediction implicitly presupposes a difference between chronological and biological aging. It does so by virtue of training a model to predict chronological aging from information about the biological state of individual brains. It further presupposes that there is a *typical* trajectory of biological aging that occurs during chronological aging. It does so again by presupposing that information about the biological state of a brain is suited to predict its chronological age. If there were no typical trajectory (or trajectories), no typical pattern(s) of biological change during chronological ageing, such a prediction wouldn't be possible.

Under which circumstances a brain age prediction model works best, depends – among other things – on its intended purpose. For predicting the chronological age of a person outside the training group – which given the costs and availability of alternatives is an uncommon but still possible purpose – it works best if it generates minimal predictive error across all populations.

For this particular purpose, a nearly non-normative model is most useful. With 'non-normative' I refer to a model based on a training set including participants with very different biological states and properties including different health status. Such a model would – counter to common practice – not be limited to training data of healthy volunteers. Used for the purpose of chronological age prediction, a non-normative model is not intended to generate different prediction error for subjects conforming and subjects not conforming to some normative standard such as health.

The latter, however, is the more common aim in the use of brain age prediction. It is usually used to identify brains which differ in their trajectory of biological change over time from some norm, typically health. For this particular purpose it is useful and common to train a model on a dataset including only healthy persons, thus generating a normative model.

In this case, a brain age prediction model works best if it generates minimal error for the training group and individuals which meet the norm implicit in the model, and if it generates a different error for individuals not meeting this norm, i.e. for outliers in the biological trajectory.

As a contrast compare this to counting the rings of a tree. This too is a technique that predicts the chronological age of an organism, but in general there is no such thing as TAG, tree age gap, from which we would expect meaningful information about the health or general biological state of the tree. If brain age prediction worked as well as tree age prediction, it would be useless for most current purposes.

A positive brain age gap indicates what is often called premature or accelerated aging, while negative numbers can be read as delayed aging [2]. Yet, it is not clear whether *BAG* represents any circumscribable single phenomenon or set of phenomena, much less whether these are best described as *aging*. More research on the underlying neurobiological properties detected by brain age prediction will be required to differentiate which phenomena *BAG* is in fact sensitive too. But this kind of research will not discover whether 'aging' is a fitting terminological choice for any of these phenomena, rather it will decide – as shown in the following – that it is not.

The uses of *BAG*

Whether and how *BAG* is best seen as indicating aging depends on what practical purposes it is used for. There are several suggestions for such uses, in particular as a marker of general brain *health*, for early detection, prognosis, or differential diagnosis of brain-based *disorders*, and for the prediction of *treatment* outcome [2]. Tellingly, all these practical goals, especially the former ones, are oriented towards the identification of pathological states and processes and, where possible, at successively finding means for intervention.

General Marker of Health

As a marker of general brain health the brain age gap can indicate an increased risk of mortality [6]. In addition, it is a valuable addition to the diagnostic toolkit in neurology – and possibly neuro-psychiatry – insofar as it can indicate several different brain

related disorders without yet providing a differential diagnosis or a specific prediction of any of them.

Despite its use, BAG is neither the only nor the most cost-efficient way of predicting increased mortality risk. And while *BAG* is correlated with several brain-related disorders, it is not the best marker for predicting any specific such disorder, especially given the large number of confounds, such as educational status [7] obesity [8], blood pressure, smoking habits, dietary effects etc. [9]. Nevertheless, the function as a general marker for brain health, or brain maintenance [10], is thought to appeal to a lay person's understanding of aging as a generalized deterioration of biological structure and function [11].

Whether or not transparent to patients, *BAG* has its clinical use as a general marker. As such it refers to a generalized state of the organism's brain, as many other medical markers do. In this it differs from other uses of *BAG* which, as discussed below, identify brain age with *specific* biological causes of reduced functioning. This generality – might one say: lack of specificity – is a clinical advantage for general predictions of mortality, and thus for engaging in further investigations into the probable causes of the patient's higher mortality risk. In this use, *BAG* might improve upon established screening tools [12], however, given the costs of MRI-based markers, its potential for wider use in health screening depends amongst others on the availability of alternative diagnostic tools, its population specific diagnostic value and the availability of consecutive therapeutic interventions.

Specific predictor or diagnostic marker

BAGs generality is a disadvantage for differential diagnoses. For this purpose, modifications of the unitary measure *BAG* might become necessary. Thus, once *BAG* is employed for more specific predictive or even diagnostic purposes, the need for and direction of further research becomes apparent. When used to predict or diagnose the onset of brain-based disorders, one main goal is to improve the specificity of the marker. But for several procedures of measuring and calculating brain age gap, different brain based disorders seem to result in indistinguishable values thus limiting specificity of the tool [2]. If *BAG* predicts the onset of some brain-based disorder without specifying which one, its utility remains quite limited. If, however, *BAG* or a derivative thereof were suited to predict or even diagnose specific brain-based disorders its utility would be significantly higher, especially given that differential diagnosis of some such disorders can take extremely long with the current tools.

Consequently, this is one area that attracts significant research effort. For example, several research groups try to compare the different sensitivities of different imaging modalities. Niu and colleagues combined T1 weighted structural MRI, diffusion MRI and resting state fMRI to improve age prediction [13]. Cole used six imaging modalities present in the UK Biobank dataset [14]. It is possible to employ such a combinatorial approach not to simply improve age prediction but to identify underlying neurobiological phenomena more precisely. Rokicki and colleagues investigated not only whether multimodal measurement can improve the accuracy of brain age predictions, but they also analyzed whether different modalities are suited for generating more specific brain age gap measures for different brain-based disorders [15].

Differentiating MRI-modalities is only one way of improving the specificity of *BAG* measures. Another way to approach this task is to vary what exactly to measure during data acquisition, for example whether to take a region or voxel-based approach, to include whole brain measures or specific areas only, which areas to include. Cole for example reported that in his multimodal approach [14] only 34 phenotypes of the original 1079 variables were informative about age. Leonardsen and colleagues used a transfer learning approach, selecting 64 features in their deep learning-based brain-age prediction model for predicting brain-based disorders. They report relatively high predictive values for Alzheimer's Disease, Multiple Sclerosis, Mild Cognitive Impairment, and psychotic diagnoses [16].

There is currently a trend to use brain age prediction models – via transfer learning and other techniques – for the purpose of detecting common brain-based disorders. This is one obvious next step given that large, specialized datasets and thus trained models for detecting specific neurological disorders are not always available – and might not be for some time to come [17]. However, this is only *one* obvious next step. One alternative – nearly as obvious – will be to use brain age prediction models to identify neurobiological contributions to aging which currently are not considered to be pathological.

Here is one example from current research. Bashyam and colleagues describe how models with a moderate fit for age prediction – i.e., with a moderate as opposed to high or low mean absolute error in age prediction – are superior in predicting known brain disorders. They are superior because they take into account phenotypes that are affected by other parameters than age, i.e., pathological changes such as "small vessel ischaemic disease, amyloid plaques, and tau neurofibrillary tangles" [17]. Tight fitting models – i.e., models with low mean absolute error in age prediction – avoid such patterns, making them worse at detecting known disorders. Thus, tight fitting models are better suited to identify the neurobiological manifestations of aging beyond what is currently considered a disorder.

So, one might think that depending on imaging modality and model in use, one can deliver on the core promises of brain age prediction, namely predict mortality with a general marker of brain health, predict or diagnose known age-related disorders, and identify non-disorder changes in the aging brain, i.e., its biological age.

On the limits of *BAG* as a specific marker and as a marker for *age*

Contrary to this impression, I want to claim that *BAG* is currently severely limited in its clinical utility for specific diagnoses and, what is more, develops its own paradoxical dynamic when identifying non-disorder changes in the aging brain.

The limits of *BAG* as a specific marker

The first claim seems to already be accepted by many scientists in the field. Brain age gap is not alone in this: "Most bio markers [of aging] including telomere length lack specificity regarding the mechanisms of aging processes" [18]. Different scientists draw different conclusions from this fact, however. While some turn towards specific markers for known brain-based disorders, others try to modify *BAG* into such a marker as seen above in the work of Leonardsen or Rokicki, who investigate whether some results of BAG measures are specific for known brain-based disorders. Leonardsen et al. analysed the components of their predictive model and generated "increased predictive value for MS, AD and MCI when using the internal representations of the SFCN-reg [regression Simple Fully Convolutional Network] underlying the brain age predictions" [16]. Rokicki et al. started with the observation that "BAG has been shown to be more accurate in predicting the conversion of mild cognitive impairment (MCI) to AD compared to T1-based MRI features such as cortical thickness and regional brain volumes (Gaser & Franke, 2013) and potentially can be used as a *biomarker for early dementia risk screening* (Wang et al., 2019)" [15]. They investigated whether there are "distinct deviations in patients with psychiatric and neurological disorders" [15], which they claim to have found together with different age prediction performance: "prediction performance in the patient populations varied across the different disorders. T1w-based BAGs were *most robust in classifying AD, and also to some extent in SZ and BD*" [15].

As mentioned above, there are good reasons to take this path if better markers are not yet available and their development might depend on assembling larger datasets for a given disorder.

A serious limit for the latter approach is probably the tension between the inclusive requirements of *BAG* and the exclusive focus on disease markers. A marker for a specific disease should typically be both sensitive and specific, that is it should react to the presence of the disease in question and *only* to its presence. *BAG* on the other hand needs to include the effect of several different processes, pathological as well as non-pathological. It must account for a range of related differences beyond those caused by a single disorder. Thus, even if, which is highly improbable, one specific disease would cause a distinctive difference in *BAG* – say for example Early Onset Alzheimer's causes a *BAG* of x years – other phenomena will overlap with this effect and the person's overall *BAG* will most likely not be x. As mentioned above, BAG knows several confounds [9]. This methodological issue basically forced researchers to decide what they want their specific procedure to do, or rather what they want it to do well and relinquish the inclusivity of *BAG*. This is why Baecker et al. in their overview and introduction to the method conclude in their paragraph on *Differential diagnosis of brain-based disorders* "that an abnormal brain age cannot be a stand-alone measure of diagnosis, as it lacks specificity, especially in light of the inherently large neuroanatomical heterogeneity in the general

population" [2]. Rather they see its potential in providing complementary information for other diagnostic biomarkers and methods.

The limits of *BAG* as a (stable) marker for age

My second claim raises a more substantial issue, namely that the use of the biomarker *BAG* results in a paradox. Very briefly the paradox is this: by using brain *age* gap as a clinical marker, one tacitly takes the biological states this marker indicates to be *pathological*. This claim consists of two parts. First, the judgment that an *individual* person's brain appears older than that of an age controlled reference class, one takes the biological states of this person's brain to be pathological. The second part is slightly more complicated. Brain age predictors compare a person's imaging results to those of many other individuals of known age. They seem to have an implicit model of the typical, i.e. statistically normal measures in a given age controlled reference class. Identifying the biological states which the predictor takes to be indicative for advanced age *across individuals* is a first step towards taking these states and their causes – not age *tout court* but age promoting processes (see below) – as potentially pathological.

In both parts, two additional steps are required for full pathologization: One is correlating these changes with possible functional decline and the other is taking them to be amenable to intervention. We'll see in the following, that both steps follow naturally from taking the first ones just outlined.

Before I detail the steps of pathologization of age, a short observation on the expected effects is in order, because they complete the paradox that surrounds the use of brain age prediction: To take biological states to be pathological implies endorsing means for diagnosis and therapeutic or preventative intervention and a commitment to eliminate them as corollaries of chronological age. Successful interventions into processes originally associated with aging will *ceteris paribus* reduce the prevalence of the latter. Once prevalence declines so does the regular association between the process in question and chronological age. In effect, a biomarker specific for this process will no longer be useful as indicators of age or age-related diseases. Admittedly, this effect is to be expected in the very long run only, given the state of the art of interventions in aging processes and the expectable speed of their dissemination.

Nevertheless, this paradox has interesting practical repercussions, which are caused by the character of this particular case of pathologization. It happens within a certain research paradigm and not in the context of a dispute between disciplines or even further social influence groups. And even within the research paradigm it occurs more or less tacitly via the use of specific, highly useful tools, namely biomarkers for age, such as Brain Age Gap. In this, this particular case of pathologization stands in contrast with others. This case of pathologization does not result from medicalization – the active extension of the definitional sovereignty of medicine – which would spark open social dispute and therefore public deliberation. Nor is it a result of a deliberative process about the status as pathological within the discipline or across disciplines. Due to the research paradigm-internal and tacit character of this pathologization of aging, the social and interdisciplinary discourse about whether certain states of persons belong to the field of interpretation and intervention of medicine is omitted – just as tacitly. This process, however, is what confers social justification on the social decision to consider a state or type of states a valid target for medical interventions. Even if this process should come to the same conclusion – which is possible but not foreseeable at present – its absence generates a lack of justification.

The paradox effect just outlined is shared with other biomarkers of aging, it finds its origin in their reliance on a biological concept of aging. Biological conceptions of age necessarily draw a contrast to chronological age. The intuitive and plausible core idea behind every biological conception of age is that a person can be biologically older or younger than her chronological age. This is not just a common-sensical intuition expressed in describing some persons as aging well and others as aging badly. As a scientific concept it allows us to describe how biological processes of change with chronological age occur faster in some organisms than in others. Once it is accepted that biological age is different from chronological age, it isn't possible anymore to identify an organism's age by calendrical methods. Instead, a measuring method of some biological state of the organism is required. Which state exactly should be measured depends on the specifics of the current theory of

biological aging [19, but see 20]. Thus, understanding aging as biological aging requires using biomarkers of aging and these only make sense against this conceptual background.

Now a marker of aging will typically not be sensitive to the whole of the aging processes or its effects. Rather individual biomarkers will be sensitive to specific processes or effects of aging. That is the reason why there are several different biomarkers from diverse biological subdisciplines and techniques [21]. The brain age gap (or other organ age gaps) is just one such marker. What such a marker is sensitive to can also be called a mechanism of aging – or a set of such mechanisms and their effects in the case of *BAG*.

This idea has recently been cast into a science-based philosophical theory of aging by Maël Lemoine [22]. His definition makes it obvious why biological age deviates from chronological age:

> "The aging of an organism is a process resulting from the combination of mechanisms limiting its lifespan ("promotive") and mechanisms modulating their effects ("protective"). The balance between the effects of these two types determines the rate of aging" [22].

If aging results from promotive and protective processes and these can occur in different organisms and environments to different degrees, the rate of aging varies from organism to organism and is not fixed by chronological age. In principle it is even possible that protective mechanisms overcompensate promotive mechanisms for some time, which would result in a reduction of biological age. In fact, there are several observations of a reverse of BAG values associated with special circumstances [23] or interventions [24]. One advantage of Lemoine's reconstruction of the processes of aging is that it is informed by current science but neutral enough with regard to scientific theories of aging. It is compatible with there being just one cause of aging, which result in a plurality of different mechanisms, as well as with there being several separate causes, which in various constellations result in different mechanisms [20].

One contingent fact about aging in real world species is that promotive mechanisms of aging predominantly result in some form of structural damage, functional decline, or depletion and protective mechanisms typically merely counteract promotive mechanisms or repair the damage they do. One can imagine species in which promotive mechanisms of aging result in resilience and functional increase – in fact, fantasy authors have imagined such creatures – and it would still make sense to describe their development as aging and investigate their rate of biological aging. Real world cases of biological aging, however, share two core characteristics: depending on the specific organism and its environment they can (1) progress faster (or slower) in comparison to other organisms of the same species and (2) they result in structural damage, functional decline, or depletion of the organism. That sound suspiciously like another type of process with similar development path: pathological processes. Depending on the definition of 'pathological', processes are considered as pathological only if (2) they result in structural damage, functional decline, or depletion as (1) compared to the biological state of other organisms in a reference class [25].

To support the case that using a biological concept of aging and measuring age with biomarkers of aging takes their referents to be pathological, it helps compare the scientific trajectory of the identification of mechanisms of aging. Lemoine points out a telling analogy:

> "This situation probably reflects a sort of backtracking process similar to that observed in the investigation of diseases: a disease is first defined as a collection of symptoms, then generally as the underlying imbalance responsible for those symptoms, then as the local deficiency causing the imbalance, and finally based on the etiology of the disease" [22].

The claim is that the investigation of mechanisms of aging shows close parallels to the investigations of diseases. It starts from the observation of macroscopic changes in an organism and its behavior, then proceeds to identify microscopic changes and their causes. Lemoine's description is so accurate that it tempts the reader to overlook one minor inaccuracy. The inaccuracy is contained in the claim that the process of identifying mechanisms of aging is *similar* to the investigation of diseases. I want to claim that it isn't similar, it is exactly the same.

What could the differences between these two processes of investigation be? One answer could be based on different methodologies. There could be different initial observations starting these investigations, there could be different methods of collecting

and analyzing data. Admittedly, biogerontology, which is probably the dominant sub-discipline studying aging, while a young field, is already encompassing several different strands of investigation, including evolutionary biology, ecological sciences, and physiology as well as molecular biology. It is hardly possible to compare these strands of research in and beyond the biogerontological field. There is, however, no separate biogerontological physiology or biogerontological molecular biology, but a set of shared approaches, methods, and not least, theories across fields of investigation.

Another answer could refer to differences in the subject under investigation. It would beg the question to argue that one investigates pathological processes and the other non-pathological ones. Here are some more informative possible suggestions: a) investigations of diseases but not of mechanisms of aging seek out extrinsic causes of mortality, b) investigations of diseases but not of mechanisms of aging seek out effects independent of chronological age, c) investigations of diseases but not of mechanisms of aging seek out causes of mortality (or loss of function and damage) occurring in some but not all members of a taxon.

The first two suggestions seem to be ruled out by actual research into purely internal diseases such as some cancers, auto-immune-diseases etc. as well as by investigations into age related diseases ranging from arthritis to Alzheimer's. The third suggestion needs to be taken more seriously. Indeed, investigations into diseases typically target processes that do not affect every member of a taxon. But this is neither a universal rule, nor would it be a necessary one if it happened to be true at the moment. Especially with parasitic diseases, it is quite likely that cases of full endemic contamination occur. Matthewson and Griffiths for example speculate that there might well be no cockatoos without feather lice [26]. Even if there were not, that would not rule out the kind of investigation described above starting with a set of symptoms and ending with the etiology. Caplan points out that universality is not distinctive of aging, referring to gum disease, which is universal in humans, but still clearly investigated and treated as a disease [27]. A process (or state) can still be investigated as any other disease even if it is universal.

To summarize the argument up to this point: The progress of different mechanisms of aging relative to the chronological age of the organism is what biological conceptions of age refer to and what markers of aging measure. Biological aging is typically associated with structural damage, functional decline, depletion. Mechanisms of aging are detected and investigated in the same way as pathological processes are.

Where does that leave markers of aging in general and brain age prediction in particular? They are scientifically and practically on par with investigations into pathological processes resulting in structural damage, functional decline, and depletion. This equivalence of investigations into mechanisms of aging and into pathological states is responsible for the tacit character of the pathologization of aging. It happens through the use of standard tools and methods of scientific investigation.

The further claim introduced above was that the use of the marker 'brain age gap' amounts to taking the biological state it indicates to be (potentially) pathological. As mentioned above, this claim has two parts. First of all, biological states of an individual are taken to be pathological in case this person has a positive value in brain age gap. Zero and negative values in BAG would point to average or above average state of the biological mechanisms in question – this is why BAG points to *potentially* pathological states, not pathological states *sans phrase*. And clearly, just as with other medical markers, not every positive BAG value would be taken to indicate pathology, but levels raised beyond a threshold to be fixed by validation studies. The reason why an individual's positive BAG is taken to indicate pathological states should be clear by now. The states in question are investigated and seen as potential target of intervention in the same way as other pathological states, they are associated with structural damage, functional decline, depletion just as other pathological states are, and they negatively deviate from the statistical norm in the person's reference group.

But not just an *individual's* positive BAG is taken to indicate potentially pathological states. The same is true for the states which brain age prediction is sensitive to *in general*. This can be shown in two ways. First by simple aggregation. Individual positive BAG values indicate pathological states; each component sensitive to a single biological measure in brain age prediction can result in a positive BAG value while holding all other components equal. Thus, each

component used in brain age prediction can indicate a pathological state, i.e., indicates a potentially (if resulting in positive BAG) pathological state.

Second, and more importantly, brain age prediction models generate an implicit representation of the typical measures of biological states in a given age controlled reference class. The computation of a BAG-value is not sensitive to all variables present in the data-set but takes into account a specific subset of these states, which can be identified by analyzing the model *ex post*.

It is not just likely, it is a pressing research hypothesis that some, if not most, of the variables used in the computation of BAG values refer to biological features which changes with chronological age and is associated with damage of decline. Thus, the brain age prediction model identifies candidates for biological states the change of which most likely results in functional loss, damage or decline.

In addition, there is no principled distinction into '*pathological* biological mechanisms associated with age and with damage or decline' and '*non-pathological* mechanisms associated with age and with damage or decline' or into mechanisms licensing the development and use of interventions and mechanisms not doing so. On the contrary, they are investigated and taken to be valid targets of intervention alike. Thus, brain age prediction identifies variables referring to biological states the change of which is likely to be associated with structural damage, functional decline, depletion, i.e. they refer to potentially pathological states.

What is more: Once one has clarified that a specific biological mechanism leads to structural damage, functional decline, depletion, then the decision against developing or using interventions in this mechanism seems to require justification. This applies not only to the decision whether to intervene at all, but also to what threshold of intervention one sets. That is, if a mechanism is identified as causing damage, there arises a need to justify the decision not to stop it as long as the damage is only as great as it is for other people of the same age for whom one also does not intervene. A threshold of pathology set relative to a similar affected population becomes questionable. Thus, the use of brain age prediction seems not only to result at in a potential pathologization of some states it is sensitive to. It's use also makes it very difficult to still accept mechanisms as normal aging which are identified via brain age prediction and consecutive investigations as possible targets of the development and use if interventions.

One might object that the mere identification of *potentially* pathological mechanisms is not sufficient to fully pathologize them. A state or process is considered to be pathological only, if it crosses a certain threshold, and the setting of this threshold is a deliberate process within and often across disciplines. This objection, however, affects the pathologization of individual BAG values only. Admittedly, individual elevated BAG will only be considered to indicate pathology if it crosses a certain threshold. Even here, one should, however, keep in mind that as opposed to previous practice, this is a case of pathologization of states previously accepted as regrettable but normal. What previously counted as *aging badly* is understood as pathological.

However, the objection that processes need to cross a threshold in order to count as pathological does not affect the claim that the use of brain age prediction identifies processes associated with structural damage, functional decline, depletion across individuals and thereby pathologizes these processes. At least it does not, if one does not accept what I have explicitly rejected above, namely that one can set a threshold relative to the similarly affected population. If an independent threshold of pathology is chosen based on e.g. functional criteria, the effects of mechanisms of aging are guaranteed to cross it, it is just a question of time.

A similar claim has been made by Arthur Caplan, who writes "Why shouldn't we treat aging as a disease? Scientists are beginning to do exactly that as the mechanisms and causes of aging become clearer and begin to appear open to manipulation and alteration" [27]. I should add that it is not only when manipulation and alteration is possible, but when diagnosis becomes possible and manipulations and alterations become a valid research goal that the mechanisms come to be treated as pathological [cf. 28]. In a similar vein, Bjørn Hofmann points out this specific role of markers in his analysis of the role of technology in defining diseases: "Thus, the technological constituted signs and markers are basic to the demarcation of disease. They define disease entities and are applied to recognise disease in the particular case, and as such provide a technological semiology of disease" [29]. Even if currently manipulation and

alterations of mechanisms of aging are restricted to animal models and thus not available for the human case, with the invention of markers for biological age they have become a valid research target which draws substantial investment of scientific expertise and resources.

This is not to make the stronger claim that aging *is* pathological or that mechanisms of aging *are* diseases. While that might be the case, the present article is intended to defend the less ambitious claim that the use of biological concepts and biomarkers of aging amounts to treating biological aging as potentially pathological and mechanisms of aging as potential diseases. So much seems to be accepted even by authors who otherwise argue against understanding age as pathological, such as Schramme, who explicitly concedes: "Admittedly, aging might be avoidable in the future. If and when this is possible, it might become a disease" [30]. And the fact that the question whether aging is pathological is not in fact answered but a positive answer is presupposed in scientific and possibly clinical practice raises the issue of a lack of explicit justification of the practice.

Here is one related difference one could make in the scientific engagement with mechanisms of aging and pathological processes. One could *decide* that while both are investigated in the same way, research into manipulations and alterations is a valid research goal in one and not the other case. This is indeed a suggestion regularly made in moral theorizing [see 31] to name just one example of a common strategy], if not so much in philosophy of science.

This type of argument can proceed from a purely normative foundation claiming that we should not interfere with aging processes for some practical reasons. Mostly such practical reasons refer to the value of cultural techniques of coping with aging. Take one of the philosophically most sound examples of this argument: Thomas Schramme rightly points out that there are several valuable individual and cultural mechanisms of coping with aging. There is individual development of skills in several tasks:

> "[...] old people are much better, especially more efficient, at performing some tasks than younger people because they develop different abilities with age, while admittedly losing some others. If we eliminate senescence, these adaptive mechanisms would be lost" [30].

And there are cultural provisions changing the role of individuals while still allowing for full social inclusion and cooperation:

> "[…] people of a certain age are not expected to work for money. The rationale for having a welfare system and pensions is obviously related to the frailties of old age" [30].

While I'm sympathetic to valuing such cultural achievements, I think they not only need to prove their value when confronted with new background conditions such as interventions into the aging process. They also need to be comparatively superior to the combination of new background conditions and alternative cultural techniques. It is at least not obvious that the coping mechanisms described by Schramme – and others discussed in literature – are still comparatively valuable under conditions of different aging processes.

An alternative way to support the decision to investigate disease but not mechanisms of aging tries to refer to some factual properties of aging. One prominent strand of arguments categorizes mechanisms of aging or aging itself as a natural process and therefore not a valid target for investigating interventions, while diseases are somehow not natural and therefore valid targets of intervention.[1]

This type of argument immediately leads us astray because it appeals directly to the metaphysics of diseases and aging without recourse to the epistemological task of distinguishing between those. Thus, this

[1] To be frank, it remains completely unclear what "natural" is intended to mean in arguments of that type. As discussed above it cannot mean "universal", it cannot reasonably mean "inevitable", it is probably not intended as an opposite to "cultural" nor to "supernatural". I'll therefore not follow up on claims of naturalness but turn to the next reasonable alternative. This alternative has been suggested by some authors in wondering whether aging but not disease has an evolutionary purpose. A suggestion in this direction has been made by August Weismann in the nineteenth century [32]. He proposed the idea that aging has the purpose of weeding out elderly and less fit individuals from a population. However, Weismann rejected this idea in his own later writing [33]. It clearly relies on faulty evolutionary thinking, such as the idea of selection on the level of the whole population and the idea of a purpose in the process of evolution [34]. The same seems to be true for more contemporary references to evolutionary theories in supporting some kind of purpose of aging as amply discussed by Tim Lewens or Alan Buchanan [35, 36].

type of argument simply isn't available when one tries to identify the difference between *investigations* into pathological processes and processes of aging.

Either way, the core problem of the argument is simply that is does not provide us with any way of distinguishing investigations. If one were to use any marker of aging and identified a process that leads to damage, depletion of functional decline, we do not know whether it is a mechanism of aging or a pathological process, and the suggestion that only one of them is a valid target for research and intervention does not tell us which one.

In a nutshell, this means that *BAG* is neither a sufficiently specific marker for established brain-based disorder nor a marker for age as distinguishable from disease. It stands at the beginning of a scientific trajectory which investigates the effects and – later on – the causes and processes of so-called mechanisms of aging in the same manner as pathological states and processes are being investigated. It therefore has great potential to identify mechanisms of aging suited for intervention and thus tacitly transfer them into the category of age-related diseases. Does that mean brain age gap is suited as a marker for such newly identified age related, brain-based diseases?

The limits of BAG as a specific marker for novel age-related brain-based diseases

Again, the answer seems to be negative. The reason why brain age gap isn't a marker for novel brain-based diseases should by now be obvious, since it relies on reiteration of the argumentative steps above:

BAG will play a relevant role in the process of identifying individual mechanisms of aging as described in the previous section (*The limits of BAG as a (stable) marker for age*). It is possible to identify in the predictor which features it was sensitive to, and – by comparison with more specific markers – to identify which currently non-pathological features have been identified as indicating advanced age, as described in the section before (*The limits of BAG as a specific marker*). This will at first provide rough indicators only, which might ideally be improved upon. However, for each of the mechanisms so identified, brain age prediction will be in the same position in which it is relative to established brain-based disorders: it will not be specific to this mechanism. It will be the best we have while there are not yet any markers specific to the mechanism in question. But its role as a marker for such individual mechanisms will be transient, once a specific mechanism enters the focus of further research. Such research will likely not only generate more distinctive markers but might also turn up possible interventions.

BAG thus starts out as a not very specific marker for some mechanisms of aging – but often the best we have – and ceases to be the latter once more specific diagnostic tools are developed. What is more, when research into interventions becomes a valid research goal, what once was a mechanism of aging starts to become an age-related disease. Even if *BAG* happens to still be the best marker for the identification of the state of process in question, it ceases to be a marker for *age*. So, if *BAG* is neither a specific marker for established or novel brain-based disorder, nor a stable marker for age, what, if anything, is its role?

BAG as an investigative tool, and its repercussions

The answer to this question should have become apparent by now, given the discussion of actual use in two contexts. These are first and predominantly the research context and second, and still growing, the clinical context.

In the research context, *BAG* is in fact used as a tool for initial identification of possible causes of age-related change in the brain. As described above, the initial brain age prediction is used as a starting point for developing diverse forms of more specific markers for a plethora of different types of state, including established brain-based disorders on the one hand and other changes incurred during chronological aging on the other. Thus, while it initially sounds as if brain age prediction were one technique with a few variants, there already is a strong development towards differentiation of diverse techniques taking their starting point from current brain age prediction. Thus, in neuroscientific research, brain age prediction is an initial device for developing more precise measures of specific age-related phenotypes.

In the clinical context the role of *BAG* – and other organ-age markers – is the prediction of mortality. Depending on the result of such an initial prediction it is typically followed up by more precise investigations of the reasons of an evaluated risk of mortality and the identification of means to prevent the original prediction from coming true.

"Biomarkers of aging are suggested to predict future health and survival better than chronological age" [18].

Brain age prediction stands at the beginning of an investigation in the research context as an investigation of types of states or processes, in the medical case of tokens thereof. However, in both contexts brain age prediction does not initially make a difference between disease and mechanisms of aging. Because the investigative process does not make any difference between mechanisms it is sensitive to, it treats all of them as potential targets of investigation and intervention or development of new interventions. This is the common approach we take towards diseases. Simply by employing the same investigative process, mechanisms of aging are treated as potentially pathological states.

That markers such as *BAG* do not make a difference between mechanisms of aging and diseases does not prejudge the wider social decision of which specific processes to invest research and medical resources in. Neither does it decide the metaphysical question whether there is a difference between pathological processes and processes of aging. It does, however, presuppose a certain answer within scientific and clinical practice, namely a negative answer. At the same this is a purely pragmatic presupposition, which does not point to any properties of aging or of disease which we could identify before these investigations.

However, such pragmatic presuppositions can easily transform into conclusive social answers under two conditions, which would – even if otherwise optimal – lack societal justification. First, they must remain exclusively implicit in the respective practices, i.e. they must escape further interdisciplinary and social discussion. Second, they must become the basis of extra-disciplinary contexts of action, such as private science funding or health care. The latter, the implicit adoption of the disciplinary assumption that mechanisms of aging should be treated in the same way as pathologies, is already common in current private research funding. For this reason, it is even more important to intervene in the first condition and to ensure that interdisciplinary and societal discourse about a possible difference between mechanisms of aging and pathologies takes place.

A possible objection and outlook

The thesis up till now has been that the use of BAG and similar markers might tacitly transform aging into a pathology and this change in understanding might result in social decisions to intervene without there being sufficient social justification for these decisions. One might object that there is nothing tacit about the whole process. Rather it simply is the way in which scientific investigation and measurement of aging proceeds: Biomedical science firstly identifies a process commonly take to be a valid target of intervention – remember the long history of fictional and quack interventions – like the philosopher's stone or anti-wrinkle lotions – into the aging process. Subsequently it investigates its causes, and develops possible interventions as it does with all other pathological processes. These results are therefore not lacking in social justification at all, rather they are the outcome of the common process of scientific investigation applied to a consensually accepted target and, given a positive risk/benefit ratio, a good thing.

While I'd tentatively agree with the last claim, namely that interventions in the aging process with a positive risk/benefit ratio would be a good thing, I doubt the premises and the rest of the conclusion. First, aging has not been an uncontroversial valid target of intervention. There is a long fictional tradition discussing interventions into aging, but on a closer look, most of these stories are critical. Fictional Interventions into age usually do not turn out well, in most cases the beneficiaries forget to ask for eternal youth alongside the way, lose their enjoyment of life, suffer from the consecutive loss of their nearest and dearest or suffer some similar effects.

What is more: If interventions into aging were such an uncontroversial desire, shouldn't people say so when asked? But in fact public opinion is at best equally divided on the desirability of longevity interventions. A large international survey published in 2020 found a total of 50% of the participants opposed to developing technologies allowing people to live to 120 [37]. Previous surveys have shown similar results. Given these empirical results it is hard to argue that the aging process

is an uncontroversial valid target for developing interventions.

Second, previous – and some contemporary – theories of ageing do not take age to be pathological (see below). This alone shows that the process described above is only one possible (and admittedly possibly the best) way to "to measure and understand the ageing process".

That age is a pathological process is not just a controversial claim in the wider public as indicated by the survey results just mentioned. It is controversial even within the biomedical community. There still are research programs in the life sciences taking age to have an evolutionary role – [38] or to otherwise deny that it is a disease [39]. And beyond the life sciences, several authors quite explicitly insist on the role of biological aging in a human life and speak out loudly against any attempt to develop, much less use, interventions into the biological aging process (a critical overview is provided in [40]. Aging came in first when the British Medical Journal asked their readers to name non-diseases, i.e. "a human process or problem that some have defined as a medical condition but where people may have better outcomes if the problem or process was not defined in that way" [41].

To summarize: aging is neither consensually perceived to be a valid target for the development of interventions, nor is it uncontroversial to categorize it as pathological. Even if the process of investigation and pathologization as described above will turn out to be successful and generate interventions with great risk/benefit ratio, that alone neither makes it socially transparent nor justified.

Turned positively, it can be said that the differentiation between biological and chronological aging and the use of biomarkers of aging such as BAG confronts us with the question of which states and processes to target in research and development of medical technologies as an open social question. The answer should be given based on careful deliberation of the interests and needs involved, the cultural and medical means of coping with either, and the effects of trying to alleviate one or the other.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Competing Interests The Author declares that they have no competing interests.

References

1. Burke, S.N., and C.A. Barnes. 2006. Neural plasticity in the ageing brain. *Nature Reviews Neuroscience* 7 (1): 30–40. https://doi.org/10.1038/nrn1809.
2. Baecker, L., et al. 2021. Machine learning for brain age prediction: Introduction to methods and clinical applications. *eBioMedicine* 72: 103600. https://doi.org/10.1016/j.ebiom.2021.103600.
3. Franke, K. and C. Gaser. 2019. *Ten Years of BrainAGE as a Neuroimaging Biomarker of Brain Aging: What Insights Have We Gained?* Frontiers in Neurology, 10. https://doi.org/10.3389/fneur.2019.00789
4. Butler, E.R., et al. 2021. Pitfalls in brain age analyses. *Human Brain Mapping* 42 (13): 4092–4101. https://doi.org/10.1002/hbm.25533.
5. de Lange, A.-M.G., et al. 2022. Mind the gap: Performance metric evaluation in brain-age prediction. *Human Brain Mapping* 43 (10): 3113–3129. https://doi.org/10.1002/hbm.25837.
6. Cole, J.H., et al. 2018. Brain age predicts mortality. *Molecular Psychiatry* 23 (5): 1385–1392. https://doi.org/10.1038/mp.2017.62.
7. Steffener, J., et al. 2016. Differences between chronological and brain age are related to education and self-reported physical activity. *Neurobiology of Aging* 40: 138–144. https://doi.org/10.1016/j.neurobiolaging.2016.01.014.
8. Ronan, L., et al. 2016. Obesity associated with increased brain age from midlife. *Neurobiology of Aging* 47: 63–70. https://doi.org/10.1016/j.neurobiolaging.2016.07.010.
9. Wrigglesworth, J., et al. 2021. Factors associated with brain ageing - a systematic review. *BMC Neurology* 21 (1): 312. https://doi.org/10.1186/s12883-021-02331-4.
10. Anatürk, M., et al. 2021. Prediction of brain age and cognitive age: Quantifying brain and cognitive maintenance in aging. *Human Brain Mapping* 42 (6): 1626–1640. https://doi.org/10.1002/hbm.25316.
11. Cole, J.H., et al. 2019. Brain age and other bodily 'ages': Implications for neuropsychiatry. *Molecular Psychiatry* 24 (2): 266–281. https://doi.org/10.1038/s41380-018-0098-1.
12. Beheshti, I., N. Maikusa, and H. Matsuda. 2018. The association between "Brain-Age Score" (BAS) and traditional neuropsychological screening tools in Alzheimer's

disease. *Brain and Behavior* 8 (8): e01020. https://doi.org/10.1002/brb3.1020.
13. Niu, X., et al. 2020. Improved prediction of brain age using multimodal neuroimaging data. *Human brain mapping* 41 (6): 1626–1643. https://doi.org/10.1002/hbm.24899.
14. Cole, J.H. 2020. Multimodality neuroimaging brain-age in UK biobank: Relationship to biomedical, lifestyle, and cognitive factors. *Neurobiology of Aging* 92: 34–42. https://doi.org/10.1016/j.neurobiolaging.2020.03.014.
15. Rokicki, J., et al. 2021. Multimodal imaging improves brain age prediction and reveals distinct abnormalities in patients with psychiatric and neurological disorders. *Human Brain Mapping* 42 (6): 1714–1726. https://doi.org/10.1002/hbm.25323.
16 Leonardsen, E.H., et al. 2022. Deep neural networks learn general and clinically relevant representations of the ageing brain. *NeuroImage* 256: 119210. https://doi.org/10.1016/j.neuroimage.2022.119210.
17. Bashyam, V.M., et al. 2020. MRI signatures of brain age and disease over the lifespan based on a deep brain network and 14 468 individuals worldwide. *Brain* 143 (7): 2312–2324. https://doi.org/10.1093/brain/awaa160.
18. Hartmann, A., et al. 2021. *Ranking Biomarkers of Aging by Citation Profiling and Effort Scoring*. Frontiers in Genetics, 12. https://doi.org/10.3389/fgene.2021.686320
19. López-Otín, C., et al. 2013. The Hallmarks of Aging. *Cell* 153 (6): 1194–1217. https://doi.org/10.1016/j.cell.2013.05.039.
20. Gems, D., and J.P. de Magalhães. 2021. The hoverfly and the wasp: A critique of the hallmarks of aging as a paradigm. *Ageing Research Reviews* 70: 101407. https://doi.org/10.1016/j.arr.2021.101407.
21. Kudryashova, K.S., et al. 2020. Aging Biomarkers: From Functional Tests to Multi-Omics Approaches. *Proteomics* 20 (5–6): 1900408. https://doi.org/10.1002/pmic.201900408.
22. Lemoine, M. 2020. Defining aging. *Biology & Philosophy* 35 (5): 46. https://doi.org/10.1007/s10539-020-09765-z.
23. Luders, E., et al. 2018. Potential Brain Age Reversal after Pregnancy: Younger Brains at 4–6 Weeks Postpartum. *Neuroscience* 386: 309–314.
24. Luders, E., N. Cherbuin, and C. Gaser. 2016. Estimating brain age using high-resolution pattern recognition: Younger brains in long-term meditation practitioners. *NeuroImage* 134: 508–513. https://doi.org/10.1016/j.neuroimage.2016.04.007.
25. Boorse, C. 1977. Health as a Theoretical Concept. *Philosophy of Science* 4: 542–573.
26. Matthewson, J., and P.E. Griffiths. 2017. Biological Criteria of Disease: Four Ways of Going Wrong. *The Journal of Medicine and Philosophy: A Forum for Bioethics and Philosophy of Medicine* 42 (4): 447–466. https://doi.org/10.1093/jmp/jhx004.
27. Caplan, A. 2015. *How Can Aging Be Thought of as Anything Other Than a Disease?*, in *Handbook of the Philosophy of Medicine*, T. Schramme and S. Edwards, Editors, Springer Netherlands: Dordrecht. p. 1–8. https://doi.org/10.1007/978-94-017-8706-2_10-1
28. Cooper, R. 2002. Disease. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 33 (2): 263–282. https://doi.org/10.1016/s0039-3681(02)00018-3.
29. Hofmann, B. 2001. The technological invention of disease. *Medical Humanities* 27 (1): 10–19. https://doi.org/10.1136/mh.27.1.10.
30. Schramme, T. 2013. 'I hope that I get old before I die': Ageing and the concept of disease. *Theoretical Medicine and Bioethics* 34 (3): 171–187. https://doi.org/10.1007/s11017-013-9256-2.
31. Callahan, D. 1990. *What kind of life: the limits of medical progress 318*. New York: Simon and Schuster.
32. Weismann, A. 1891. *Essays upon heredity and kindred biological problems*. 2d ed. Vol. v.1, Oxford: Clarendon Press. https://www.biodiversitylibrary.org/item/23551. Accessed 18 Sept 2023.
33. Kirkwood, T.B.L., and T. Cremer. 1982. Cytogerontology since 1881: A reappraisal of August Weismann and a review of modern progress. *Human Genetics* 60 (2): 101–121. https://doi.org/10.1007/BF00569695.
34. Milholland, B., and J. Vijg. 2022. Why Gilgamesh failed: The mechanistic basis of the limits to human lifespan. *Nature Aging* 2 (10): 878–884. https://doi.org/10.1038/s43587-022-00291-z.
35. Buchanan, A. 2009. Human nature and enhancement. *Bioethics* 23 (3): 141–150. https://doi.org/10.1111/j.1467-8519.2008.00633.x.
36. Lewens, T. 2012. Human Nature: The Very Idea. *Philosophy & Technology* 25 (4): 459–474. https://doi.org/10.1007/s13347-012-0063-x.
37. Prudhomme, M. 2020. *SIENNA D3.5: Public views of human enhancement technologies in 11 EU and non-EU countries*. https://doi.org/10.5281/zenodo.4068194
38. Johnson, A.A., M.N. Shokhirev, and B. Shoshitaishvili. 2019. Revamping th e evolutionary theories of aging. *Ageing Research Reviews* 55: 100947.
39. Hayflick, L. 2007. Biological aging is no longer an unsolved problem. *Annals of the New York Academy of Sciences* 1100: 1–13.
40. Overall, C. 2003. *Aging, death, and human longevity : a philosophical inquiry xi, 264*. Berkeley: University of California Press.
41. Smith, R. 2002. In search of "non-disease." *BMJ* 324 (7342): 883–885. https://doi.org/10.1136/bmj.324.7342.883.