

RESEARCH ARTICLE

Interpretable brain age prediction using linear latent variable models of functional connectivity

Ricardo Pio Monti^{1,7*}, Alex Gibberd², Sandipan Roy³, Matthew Nunes³, Romy Lorenz^{4,5}, Robert Leech⁶, Takeshi Ogawa⁸, Motoaki Kawanabe^{7,8}, Aapo Hyvärinen^{9,10}

1 Gatsby Computational Neuroscience Unit, University College London, London, United Kingdom, **2** Department of Mathematics & Statistics, Lancaster University, Bailrigg, United Kingdom, **3** Department of Mathematical Sciences, University of Bath, Bath, United Kingdom, **4** MRC Cognition and Brain Sciences Unit, University of Cambridge, Cambridge, United Kingdom, **5** Department of Psychology, Stanford University, Stanford, CA, United States of America, **6** Centre for Neuroimaging Science, Kings College London, London, United Kingdom, **7** RIKEN Center for Advanced Intelligence Project (AIP), Kyoto, Japan, **8** Brain Information Communication Research Laboratory Group, Advanced Telecommunications Research Institute International (ATR), Kyoto, Japan, **9** Université Paris-Saclay, Inria, 91190 Palaiseau, France, **10** Department of Computer Science and HIIT, University of Helsinki, Helsinki, Finland

* ricardo.monti08@gmail.com



OPEN ACCESS

Citation: Monti RP, Gibberd A, Roy S, Nunes M, Lorenz R, Leech R, et al. (2020) Interpretable brain age prediction using linear latent variable models of functional connectivity. PLoS ONE 15(6): e0232296. <https://doi.org/10.1371/journal.pone.0232296>

Editor: Carlo Vittorio Cannistraci, Technische Universität Dresden, GERMANY

Received: December 4, 2019

Accepted: April 11, 2020

Published: June 10, 2020

Peer Review History: PLOS recognizes the benefits of transparency in the peer review process; therefore, we enable the publication of all of the content of peer review and author responses alongside final, published articles. The editorial history of this article is available here: <https://doi.org/10.1371/journal.pone.0232296>

Copyright: © 2020 Monti et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: With respect to the CamCAN data, the resting state fMRI data was employed. This can be accessed at: <https://camcan-archive.mrc-cbu.cam.ac.uk/dataaccess/>.

Abstract

Neuroimaging-driven prediction of brain age, defined as the predicted biological age of a subject using only brain imaging data, is an exciting avenue of research. In this work we seek to build models of brain age based on functional connectivity while prioritizing model interpretability and understanding. This way, the models serve to both provide accurate estimates of brain age as well as allow us to investigate changes in functional connectivity which occur during the ageing process. The methods proposed in this work consist of a two-step procedure: first, linear latent variable models, such as PCA and its extensions, are employed to learn reproducible functional connectivity networks present across a cohort of subjects. The activity within each network is subsequently employed as a feature in a linear regression model to predict brain age. The proposed framework is employed on the data from the CamCAN repository and the inferred brain age models are further demonstrated to generalize using data from two open-access repositories: the Human Connectome Project and the ATR Wide-Age-Range.

1 Introduction

The human brain changes during the lifespan of an adult, resulting in robust and reproducible changes in structure and function [1, 2]. Moreover, there is reason to hypothesize that deviations from the typical brain ageing trajectory may reflect latent neuropathological influences [3], serving to motivate further research into developing reliable biomarkers derived from brain imaging data. Such biomarkers could be fundamental in order to better understand and

With respect to the HCP data, we studied resting state fMRI data from HCP Young Adult dataset: <https://www.humanconnectome.org/study/hcp-young-adult/document/1200-subjects-data-release>. With respect to the ATR Wide-Age-Range data, the resting state fMRI data was studied: <https://bicr-resource.atr.jp/impact/>.

Funding: R.P.M. was supported by the Gatsby Charitable Foundation. A.H. was supported by a Fellowship from CIFAR, and from the DATAIA convergence institute as part of the "Programme d'Investissement d'Avenir", (ANR-17-CONV-0003) operated by Inria. M.K. was partially supported by MEXT Grant-in-Aid for Scientific Research (KAKENHI 18KK0284, 19H04924).

Competing interests: The authors have declared that no competing interests exist.

combat age-associated neurodegenerative diseases. To date, early studies have shown success in the context of traumatic brain injury [4] and schizophrenia [5].

Due to the significant potential benefits associated with brain-imaging driven biomarkers for age, there have been many statistical models proposed for healthy brain ageing. These models vary in complexity as well as in the class of neuroimaging data employed. One of the earliest demonstrations was that of [6], who employed voxel-based morphometry to demonstrate the structural changes which occur during healthy ageing. More recently, a wide range of sophisticated machine learning methods have been employed [7, 8, 9]. [4] employed Gaussian process regression to predict the biological age of subjects using structural neuroimaging data, demonstrating that such a model was able to accurately predict brain age. The resulting model was subsequently applied to subjects with traumatic brain injury (TBI), where the associated residuals (difference between predicted and true biological age) were shown to be significantly larger for subjects with TBI as compared with healthy subjects; the associated model consistently predicted subjects with TBI to be *older*, possibly a result of accelerated atrophy. This work was further extended by [10], who employed convolutional neural networks to obtain improved performance. In related work, [11] employ kernel regression with an application to the early identification of Alzheimer's disease.

While the vast majority of the literature has employed structural imaging modalities, there are also numerous examples of where functional imaging has been utilized. A pertinent example is [12], who employ resting-state fMRI together with support vector machines (SVMs) in order to accurately classify subjects as being either children (ages 7-11 years old) or adults (ages 24-30 years old). Furthermore, they observe an overall decrease in network connectivity as subjects mature. In related work, [13] identify ageing-driven changes in functional connectivity, highlighting decreased connectivity within the default mode network and the somato-motor network. Subsequently, [14] categorized the changes in functional connectivity that occur with healthy ageing in terms of various network measures.

More generally, the study of functional connectivity is itself an exciting avenue of modern neuroscientific research which has shown great potential for improving our understanding of the human brain function and architecture [15]. By way of example, changes in functional connectivity have been related to various neuropathologies such as Parkinson's disease [16] and Alzheimer's [17] as well as conditions such as Autism [18]. Recently, the changes in functional connectivity induced by ageing have begun to be studied. Initial studies have reported significant differences in the connectivity between younger and older subjects using resting-state fMRI [14]. Moreover, results appear to suggest there are important changes that occur in the connectivity not just between regions but also at the level of entire networks. However, despite recent advances, a holistic understanding of the relationship between healthy ageing and the associated changes in functional connectivity is still missing.

In this work we seek to build robust models of brain age based on the functional connectivity of individuals. This serves to combine the two prominent avenues of neuroscientific research: brain age prediction and analysis of functional connectivity. In particular, the methods presented in this work have two principal objectives:

1. To demonstrate that measures of functional connectivity can reliably be employed as features in machine learning models of brain age. To this end we build and validate models using three large open-source datasets: the Cambridge Center for Ageing and Neuroscience (CamCAN), the Human Connectome Project (HCP) and the ATR Wide-Age-Range datasets.
2. We further wish to interpret and inspect the proposed models in order to gain further insights into the changes in functional connectivity associated with ageing. This calls for the

use of parsimonious and simple predictive models together with features whose relationship with functional connectivity is clearly understood.

Throughout this paper, we put forward the thesis that for the potential impact of functional connectivity assessment to be met (i.e., in terms of developing powerful biomarkers) the research community needs to develop robust methods for data-analysis which can combine both supervised and unsupervised models of functional connectivity analysis. Instead of tweaking existing statistical methods, it is imperative to develop methods which are intuitive, interpretable, and insightful from a neurophysiological perspective. Such models must utilise as much experimental information as possible in order to investigate the factors which affect functional connectivity.

To further motivate our thesis, one should consider that most experiments to date operate on data from a single laboratory, or class of experiment which limits the generality of any obtained results. Such concerns have been recently recognised, particularly within the context of brain ageing [19, 20], and have given rise to multi-laboratory collaborations with data-sharing becoming more common. However, it is still highly unlikely that all subject features (and how these are measured) will be comparable across different experimental environments. Thus while data-sharing has seen much progress, it could be argued that the impact of these endeavours is still to come, and to achieve this, we need to develop methods which can combine information from across disparate, but informative experiments.

To this end we proceed in a two-step framework. First, we seek to learn robust features which summarize properties of functional connectivity across a cohort of subjects in an unsupervised manner. Due to our focus on interpretability, we focus on linear latent variable models, such as principal component analysis (PCA), independent component analysis (ICA) and their generalizations. The benefit of employing latent variable models such as PCA is that we may interpret the latent variables in terms of activity within functional connectivity networks, as proposed by [21] (see also Fig 2 below). Second, once features have been obtained in an unsupervised manner, they are subsequently used to predict brain age using standard linear regression models. We deliberately restrict ourselves to simple linear classifiers as they can be easily interrogated, allowing us to explicitly understand how each feature contributes to the predicted brain age. An overview of our two-stage approach is provided in Fig 1.

The remainder of this manuscript is organized as follows: in Section 2 we first review linear latent variable models and their implications for functional connectivity analysis. We then present our proposed two-step procedure. Experimental results, studying synthetic as well as real resting-state fMRI data, are presented in Section 3.

2 Materials and methods

We focus our analysis on resting-state fMRI time series data which is collected across a cohort of N subjects. For the i th subject, it is assumed we have access to fMRI measurements over p fixed regions of interest, denoted by $X^{(i)} \in \mathbb{R}^p$, as well as the subjects age, $a^{(i)} \in \mathbb{R}_+$. Throughout this work we approximately model the fMRI data for each subject with a stationary multivariate Gaussian distribution, $X^{(i)} \sim \mathcal{N}(0, \Sigma^{(i)})$, where $\Sigma^{(i)}$ denotes the covariance for subject i . Each entry in $\Sigma^{(i)}$ denotes the covariance between any pair of regions, which serves to define a measure of the functional connectivity [22]. As such, it follows that $\Sigma^{(i)}$ encodes a functional connectivity network over p regions where edges encode the marginal dependence structure.

The goal of the proposed methods is to learn interpretable and robust models to predict the biological age, $a^{(i)}$, of subjects given information relating only to their functional connectivity. To achieve this, we propose a two-step framework. Our approach first employs linear latent variable models in order to model high-dimensional connectivity matrices using a reduced

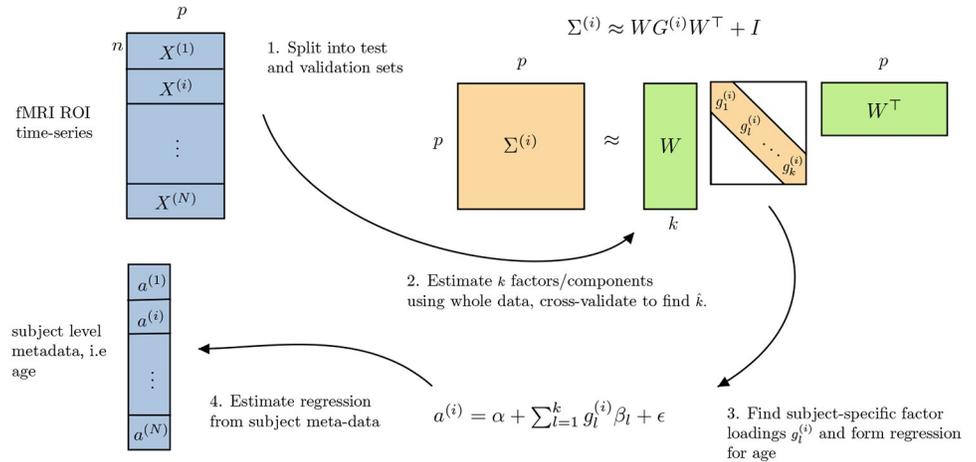


Fig 1. Pipeline for estimating networks, factor loadings, and predictive model for biological brain age. Inferred factors $W \in \mathbb{R}^{p \times k}$ describe networks which are reproducible across the entire population, the subject-specific factor loadings $g_l^{(i)}$ are then used to predict brain age. Once the factor loadings are estimated as above, using one experimental data-set (we use CamCAN data in our experiments), we can then assess how these factors perform for brain age prediction on completely held-out data-sets; we demonstrate how the model generalizes well using HCP and ATR Wide-Age-Range datasets.

<https://doi.org/10.1371/journal.pone.0232296.g001>

number of latent variables. We interpret such variables as corresponding to functional connectivity networks, allowing us to describe patterns in connectivity as being composed of various distinct networks. We note that such a two-step approach has previously been employed in the context of brain age prediction [11, 9]. However, as far as we are aware, this is the first work to directly interpret the role of linear latent variable models, such as PCA, as learning the relevant functional networks. This work thereby provides a clear motivation and interpretation for such a two-stage strategy.

Constraints on W introduced by PCA and related models	Pros	Cons
Plain PCA (Leonardi et al., 2013)	- Optimal low rank approximation to covariance	- Difficult to interpret w_k due to +/- entries - Not identifiable - Activations restricted to be uncorrelated
Non-negative ($w_{k,j} \geq 0$) (Non-neg. PCA; Sigg & Buhlmann, 2008)	- Spatial sparsity due to non-negativity	- No identifiability results - Regions not clustered into distinct networks
Non-negative + Orthonormal ($w_{k,j} \geq 0$, $w_k^T w_{k'} = 0$) (MHA; Monti & Hyvärinen, 2018)	- Regions clustered into networks - Model is identifiable - Component activations need not be uncorrelated	- Each region can only belong to one network
Spatial sparsity (Spatial ICA; Hyvärinen et al., 2001)	- Model is identifiable - Spatial patterns are maximally sparse	- Regions not clustered into distinct networks - Loadings not restricted to be non-negative - Activations restricted to be independent

Fig 2. Figure demonstrating the relationship between linear latent variable models, such as PCA and its extensions, to inferred networks. We highlight how introducing various structural constraints on the loading matrix, W , improves interpretability of such models.

<https://doi.org/10.1371/journal.pone.0232296.g002>

In Section 2.1 we discuss the various latent variable models employed, and highlight how introducing assumptions such as non-negativity can help further improve interpretability of results. We also discuss theoretical benefits associated with such assumptions. We then discuss the how the features (i.e., functional networks) inferred by the latent variable models may be used to build linear models for brain age.

2.1 Linear latent variable models for functional connectivity: PCA and its extensions

In this section we outline the linear latent variable models employed in the unsupervised learning stage of the proposed framework. We begin by discussing principal component analysis (PCA), a well-established technique for dimensionality reduction [23]. The common derivation for PCA poses it as an optimization problem seeking to learn the linear projection which maximizes explained variance within the projected space [24]. However, PCA can also be derived as inference under a simple linear latent variable model, which posits that observations $X^{(i)} \in \mathbb{R}^p$ are generated as a linear projection from low-dimensional latent variables, $Z^{(i)} \in \mathbb{R}^k$ [25]. When both observations and latent variables are taken to follow a multivariate Gaussian distributions we obtain the following generative model for observed data:

$$Z^{(i)} \sim \mathcal{N}(0, G^{(i)}) \tag{1}$$

$$X^{(i)} | Z^{(i)} = z^{(i)} \sim \mathcal{N}(Wz^{(i)}, \nu^{(i)}I) \tag{2}$$

where $G^{(i)} \in \mathbb{R}^{k \times k}$ is a diagonal matrix and $\nu^{(i)} \in \mathbb{R}_+$ denotes measurement noise. Eqs (1) and (2) serve to highlight how PCA can be seen as a low-rank model for the covariance matrix; by marginalizing over latent variables we obtain:

$$\Sigma^{(i)} = WG^{(i)}W^T + \nu^{(i)}I, \tag{3}$$

implying that the loading matrix, W , captures low-rank covariance structure. Learning the associated loading matrix, W , proceeds via maximizing the log-likelihood over observations across all N subjects:

$$\mathcal{L} = \sum_{i=1}^N p \log 2\pi + \log \det \Sigma^{(i)} + \text{tr}(\Sigma^{(i)-1}K^{(i)}), \tag{4}$$

where $\Sigma^{(i)}$ is as defined in Eq (3) and $K^{(i)}$ denotes the sample covariance matrix for the i th subject. In the context of PCA, the maximization is performed subject to the constraint that W be orthonormal,

$$\hat{W} = \arg \max_{W: W^T W = I} \{\mathcal{L}\}, \tag{5}$$

and a closed-form solution is obtained via eigendecomposition.

Following [21] it is possible to interpret each column of W as encoding functional networks or “eigenconnectivities”. While the loading matrix, W , is shared across all subjects, each diagonal entry of $G^{(i)}$ denotes the extent to which the associated network is expressed in subject i . This allows us to study connectivity as being composed of various distinct networks, resulting in significant benefits from the perspective of interpretability. We can further unpack Eq (3) as

follows (see also Fig 2 below):

$$\Sigma^{(i)} = \sum_{j=1}^k g_j^{(i)} W_j W_j^T + v^{(i)} I, \quad (6)$$

where W_j denotes the j th column of W and we write $g_j^{(i)}$ to denote the j th diagonal entry of the matrix $G^{(i)} \in \mathbb{R}^{k \times k}$. As such, we may interpret each W_j as encoding the j th network and $g_j^{(i)}$ as a measure of activity within the corresponding network in the i th subject.

There exist several extensions to the model described in Eqs (1) and (2), the prime example being factor analysis which allows the variances in Eq (2) to vary across dimensions. Recently, several extensions have been proposed where constraints such as non-negativity are introduced with the goal of improving the interpretability of results [26, 27, 28]. The motivation behind such methods stems from the fact that interpreting and visualizing PCA-based networks becomes very challenging, particularly in high-dimensions. Challenges arise from the fact that each principal component will correspond to a weighted sum of BOLD activities across all observed regions. As such, it is often difficult to identify which regions are the principal contributors to a certain principal component (and hence functional network) without applying ad-hoc post analysis. Furthermore, it is possible that some entries in the principal components may be negative, which further complicates the interpretation from the perspective of functional connectivity analysis.

The aforementioned issues can be mitigated via the introduction of non-negativity constraints on the loading matrix, W . This ensures that each principal component corresponds only to a weighted *positive* sum of activity over all brain regions. As such, the principal component can be directly interpreted as the contribution of each region to each functional network. Furthermore, the introduction of non-negativity will often yield sparsity in the sense that many of the entries of the principal components will be exactly zero [27]. It follows that such sparsity further facilitates the interpretation of the corresponding networks. From an optimization perspective, the loading matrix is inferred by maximizing the original log-likelihood objective, with the additional non-negativity constraint:

$$\hat{W} = \arg \max_{W: W \geq 0} \{\mathcal{L}\}. \quad (7)$$

It is important to note that the orthonormality constraint has been dropped in Eq (7), making the associated optimization problem less challenging. However, the combination of non-negativity and orthonormality, as enforced in [29], leads to several desirable properties. First, the loading matrix W has at most one non-zero entry per row. This implies that we may interpret the columns of W as encoding membership to k non-overlapping networks or clusters. Another very important benefit of introducing non-negativity and orthonormality constraints is that the matrix W is uniquely defined and identifiable. This is not the case in standard factor analytic models, where W is only identifiable up to an arbitrary rotation [30, 25]. Given that throughout this work we will directly interpret the columns of the loading matrix, W , as encoding functional connectivity networks, the lack of identifiability in PCA and factor analysis models is a significant limitation. We refer to the model presented in [29] as Modular Hierarchical Analysis (MHA). The associated optimization problem therefore becomes:

$$\hat{W} = \arg \max_{W: W^T W = I \text{ and } W \geq 0} \{\mathcal{L}\}. \quad (8)$$

MHA can therefore be seen to address the two important limitations of traditional models such as PCA and factor analysis; first that the presence of negative values in the loading

matrix complicates the interpretation of such matrices (addressed via the use of non-negativity constraints) and second is the fact that the latent variables are rotationally invariant (addressed via the further introduction of orthogonality). A further limitation of models such as PCA and factor analysis is that they implicitly assume latent variables must be uncorrelated. In many cases, especially when such models are applied on data relating to a cohort of subjects, such an assumption will not be valid, implying the associated generated models are misspecified. In contrast, MHA is able to identify and recover components even when they are uncorrelated. This is an important theoretical advantage, as MHA continues to enjoy the same identifiability properties even in the presence of correlated latent variables, and practical advantage, as we demonstrate in this work. Finally, we note that in the context of fMRI data, MHA corresponds to an intuitive generative model whereby latent variables capture the activity within each functional network. The optimization of Eqs (5), (7) and (8) is discussed in [S1 Appendix](#). Furthermore, we provide both Python and R code to implement MHA in [S1 Code](#).

Moreover, we note that model introduced by [26], termed Modular Connectivity Factorization (MCF), shares many similarities with MHA. In fact, both methods introduce non-negativity and orthonormality over the loading matrix, W . The fundamental difference, however, is that MCF is not associated with a linear latent variable model, and instead parameters are inferred as follows:

$$\hat{W} = \arg \max_{W: W^T W = I \text{ and } W \geq 0} \left\{ \sum_{i=1}^N \text{tr}(\Sigma^{(i)} K^{(i)})^2 \right\}, \quad (9)$$

where $\Sigma^{(i)}$ is defined as in [Eq \(6\)](#) and $K^{(i)}$ is the empirical covariance for the i th subject. A related approach was also proposed by [31].

Finally, it is important to note that whilst identifiability can be obtained via the combination of non-negativity and orthonormality, as is the case with the MHA model, it can also be obtained by relaxing the assumed distribution over latent variables, as is the case with independent component analysis (ICA) models. Formally, ICA is also a linear latent variable model, however, latent variables are no longer assumed to follow a Gaussian distribution [32]. While the relaxation of the Gaussianity assumption complicates the associated optimization, which must now be solved using gradient descent methods and accounting for the presence of multiple local optima due to the non-convex objective function [33], ICA has been widely employed in the study of functional connectivity [34, 35]. Moreover, we note that the “spatial” version of ICA used in fMRI reverses the roles of latent variables and loadings, which means that it is actually looking at the non-Gaussianity or sparsity of what we call here the loadings, corresponding to spatial patterns. [Fig 2](#) provides a visualization of the benefits obtained by introducing each of the aforementioned constraints. In particular, we note that it is the combination of non-negativity together with orthonormality which yields interpretable and identifiable networks. We empirically validate such claims by applying all of the aforementioned models to synthetic and real fMRI datasets below.

2.2 Predicting brain age using functional network activity

The previous section outlined the various flavours of latent variable models which can be employed in order to learn functional networks across a cohort of N subjects. The aforementioned models allow us to decompose observed functional connectivity patterns as a linear sum of networks encoded by the columns of the loading matrix, W . While the loading matrix is shared across all subjects (indicating the same networks are present across all subjects), the extent to which they contribute to the observed covariance of the i th subject is denoted by the diagonal entries of $G^{(i)}$, as stated in [Eq \(6\)](#).

We now consider the task of predicting the biological brain age, $a^{(i)}$, using inferred functional connectivity networks as features. In the interest of interpretability we limit ourselves to linear regression models of the form:

$$a^{(i)} = \sum_{j=1}^k \beta_j g_j^{(i)} + \epsilon^{(i)}. \tag{10}$$

Recall that $g_j^{(i)}$ corresponds to the j th diagonal entry of the matrix $G^{(i)}$. As such, the proposed models will essentially seek to predict the biological age of subjects by considering activity within each inferred functional network. In the case of the i th subject, the observed activity in network j is quantified by $g_j^{(i)} \in \mathbb{R}_+$. In practice, we will seek to quantify the activity of various functional networks on unseen subjects, defined to be subjects whose data was not employed to estimate loading matrix, W . We note that due to the orthonormality of W , together with Eq (6), we may estimate $g_j^{(i)}$ for data from unseen subjects, denoted by i^* , as follows:

$$\hat{g}_j^{(i^*)} = W_j^T \hat{\Sigma}^{(i^*)} W_j - \nu^{(i^*)}. \tag{11}$$

We note that Eq (11) requires the observation noise, $\nu^{(i^*)}$. This is not a concern for all subjects whose data is employed during the unsupervised learning of the latent variables, as parameters $\nu^{(i)}$ are inferred alongside loading matrix, W . However, the primary goal of this work is to build predictive models which can generalize to unseen subjects. In this context, an estimate of the observation noise, $\nu^{(i^*)}$, can be obtained as follows:

$$\hat{\nu}^{(i^*)} = \text{tr} \hat{\Sigma}^{(i^*)} - W^T \hat{\Sigma}^{(i^*)} W. \tag{12}$$

Although the class of models considered in Eq (10) may be considered amongst the simplest supervised regression models, they yield several important benefits when seeking to understand both the estimated parameters as well as the contribution of each of the features. In particular, each β_j corresponds to the regression coefficient summarizing the (linear) relationship between the activity of the j th network and biological age, conditional on all remaining networks. As such, if certain regression coefficients are deemed to be insignificant, we may conclude that the associated network is invariant during healthy ageing.

2.3 Hyper-parameter selection

The proposed two-stage estimation framework requires the input of only one hyper-parameter: the dimensionality of latent variables k . In the context of PCA and factor analysis, this hyper-parameter directly corresponds to the number of principal components or factors inferred, and a wide literature exists for tuning such a parameter [23]. One of the advantages of the latent variable models presented in Section 2.1 is that they each correspond to probabilistic models whose likelihood can be directly evaluated. As such, a logical choice to tuning hyper-parameter k is to directly maximize the log-likelihood over held out data.

In order to effectively perform hyper-parameter tuning as well as quantify the generalization performance of the proposed method, data was split into training, validation and test datasets as follows:

- First, a subset of subjects were held out as test data. As such, we obtain two datasets:

$$\left\{ X_{1:n}^{(i)}, a^{(i)} \right\}_{i \in S_{train}} \quad \text{and} \quad \left\{ X_{1:n}^{(i)}, a^{(i)} \right\}_{i \in S_{test}}$$

where $S_{train}, S_{test} \subset \{1, \dots, N\}$ denote the non-overlapping sets of training and test subjects

respectively. Recall N is the number of subjects present and we write $X_{1:n}^{(i)}$ to denote the n observations available for the i th subject.

- Training data is further split into training and validation datasets on a subject-by-subject basis.

Splitting the data in this manner allows for effective hyper-parameter tuning, using training and validation datasets, as well as for generalization performance to be measured using test dataset which corresponds to unseen subjects.

2.4 Experimental data

The data employed in this manuscript corresponds to resting-state fMRI data taken from three distinct open-access repositories. There were small variations in the resting-state functional MR image acquisition for each of the repositories considered: CamCAN [38], Human Connectome Project [37], and the ATR Wide Age Range [38]. The pre-processing employed on each dataset was as follows:

- CamCAN: This dataset was pre-processed by us. Data was motion corrected, spatially smoothed with a 5mm FWHM Gaussian kernel, registered into MNI152 standard space using FLIRT [39] via a skull-stripped high-resolution T1 image and resampled to 4x4x4mm voxel sizes. Each high resolution T1 image was segmented into grey and white matter and cerebrospinal fluid using SPM Dartel [40]. Mean timecourses for cerebrospinal fluid and white matter as well as 6 motion parameters were linearly filtered from each voxel to reduce non-neural noise.
- HCP: We used the pre-processed resting-state fMRI data from a random subset of healthy participants. Notably, the pipeline involved FIX ICA-based noise reduction process [40], to remove individual sources of physiological, non-physiological and motion related noise. Full details of the pre-processing pipeline can be found at <https://www.humanconnectome.org/study/hcp-young-adult/document/extensively-processed-fmri-data-documentation>.
- ATR: We used the preprocessed data. The pre-processing pipeline notably included regressing out the global grey matter signal as well as signals from cerebrospinal fluid and white matter, to remove sources of spurious variation

All three pre-processed fMRI datasets were subsequently processed as follows: a cortical parcellation based on resting-state functional connectivity analyses [42] was used to define 264 distinct 10mm diameter regions of interest (ROIs). The fMRI time course averaging across all voxels within each ROI was extracted. These 264 average time courses were then used in subsequent analyses. Full details are provided here <https://bcr-resource.atr.jp/var/www/webapp/bicresource/bicresource/staticfiles/pdf/Methods.pdf>.

3 Results

In this section we present a range of experimental results involving both synthetic and real resting-state fMRI datasets. Throughout this section, we contrast the performance of the various linear latent variable models presented in Section 2.1. In particular, we study the performance across the following methods: factor analysis (FA), PCA, non-negative PCA [27], MCF [26] and MHA [29] as well as ICA. In the case of ICA, we first employ PCA as a dimensionality reduction before employing the FastICA algorithm proposed by [43]. The implementations available in `Scikit Learn` were employed for Factor Analysis, PCA and ICA [44].

We first present results using synthetic data in Section 3.1. These simulation experiments serve as a numerical validation of the proposed two-stage procedure. Experiments relating to brain age prediction from resting-state fMRI data are subsequently presented in Section 3.2.

3.1 Synthetic data experiments

In this section we evaluate the performance of the proposed two-stage estimation framework using synthetic data. To this end, we generate artificial data whose properties approximately match those which are frequently reported in fMRI studies. The objective is then to quantify which of the linear latent variable models presented in Section 2.1 are able to both robustly recover the associated loading matrix, W , as well as learn the relevant factors which serve as accurate predictors of brain age on unseen subjects. Synthetic data was then generated in order to satisfy Eqs (1), (2) and (10). This is achieved as follows:

- First, we randomly generated a factor loading matrix, $W \in \mathbb{R}^{p \times k}$, which satisfied the constraints of both non-negativity and orthonormality. The reason for introducing both constraints is that we will seek to quantify how reliably each latent variable model can recover W , and it is therefore imperative to ensure we generate W from an identifiable model (see discussion in Section 2.1). In order to achieve this a dense matrix, W , was sampled with each entry following a uniform distribution over the interval $[0, 1]$. Subsequently, for each row only the entry with the largest value was retained with all other entries set to zero. Finally, the norm of each column was set to one.

- Second, the factor loadings for the i th subject, $g^{(i)} \in \mathbb{R}^k$, were randomly generated as follows:

$$g_j^{(i)} \sim \mathcal{N}(2.5, 1.0), \quad \text{for } j = 1, \dots, k$$

with all negative samples being discarded.

- The regression coefficients, $\beta \in \mathbb{R}^k$, were drawn uniformly at random from the interval $[0, 10]$.
- Finally, we are able to randomly generate observations and ages for each subject as follows:

$$X^{(i)} \sim \mathcal{N}(0, WG^{(i)}W^T + \nu^{(i)}), \quad (13)$$

$$a^{(i)} \sim \mathcal{N}(\beta^T g^{(i)}, \epsilon). \quad (14)$$

Recall that $G^{(i)} \in \mathbb{R}^{k \times k}$ is a diagonal matrix consisting of entries $g_j^{(i)}$.

We note that the choices for sampling distributions of both the factor loadings, $g^{(i)}$, as well as the regression coefficients, β , are necessarily somewhat heuristic. However, care was taken to ensure the implied distributions over subject ages approximately matched the empirical distributions observed within the CamCAN repository.

We note that throughout experiments we consider the performance of each method whilst varying two distinct factors: the number of observations per subject, n , and the number of training subjects, N . Furthermore, throughout simulations we fix the dimensionality of observations to be $p = 50$ and the number latent factors to be $k = 5$.

Given artificial data generated as described above, we look to quantify the performance of each of the linear latent variable models using the following two metrics:

1. Accurate recovery of the loading matrix, W . This is quantified in terms of the squared error between the true loading matrix and the estimated loading matrix.

2. Accurate brain age prediction over unseen subjects. In line with other literature, this is quantified in terms of the mean absolute error between true and predicted brain ages [11, 8].

3.1.1 Synthetic data results. We begin by considering the performance of each linear latent variable model as the number of observations per subject, n , increases for a fixed number of training subjects, $N = 25$. The results are presented in Fig 3. We note that both in terms of recovery of the loading matrix, W , as well as in terms predicting the ages over unseen subjects, the introduction of regularity constraints, be they in the form of non-negativity, orthonormality or non-Gaussianity or sparsity (as in ICA), leads to improvements.

We also study the performance of the various latent variable models when the number of training subjects, N , increases and the number of observations is fixed at $n = 100$ per subject. These results are presented in Fig 4. In terms of recovery of the loading matrix, W , we again observe that introducing regularity constraints leads to significant improvements. In terms of

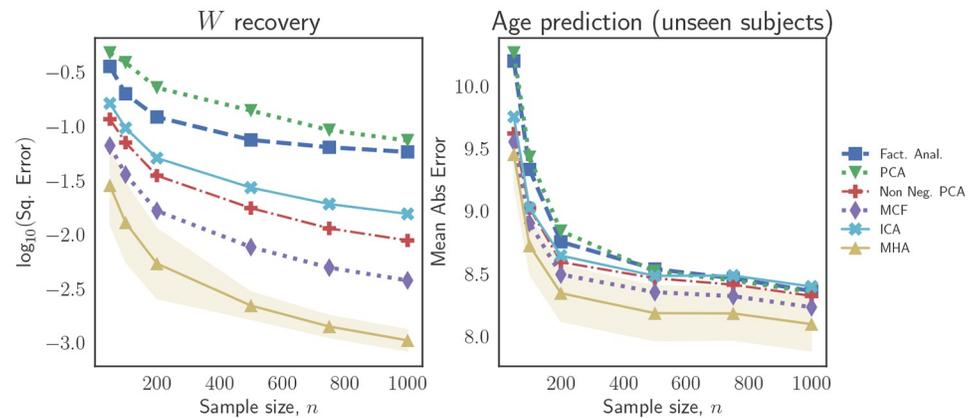


Fig 3. Simulation results for recovery of the true loading matrix (left panel) and prediction of brain age for unseen subjects (right panel) as the number of observations per subject, n , increases. We note that the introduction of regularity constraints (e.g., non-negativity or orthonormality) on the loading matrix leads to improvement in performance.

<https://doi.org/10.1371/journal.pone.0232296.g003>

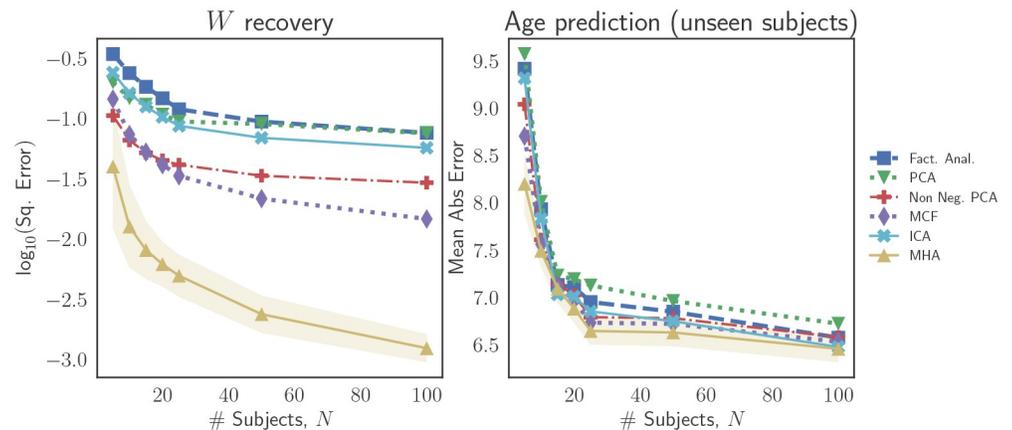


Fig 4. Simulation results for recovery of the true loading matrix (left panel) and prediction of brain age for unseen subjects (right panel) as the number of training subjects, N , increases. We note that the introduction of regularity constraints (e.g., non-negativity or orthonormality) on the loading matrix leads to improvement in performance.

<https://doi.org/10.1371/journal.pone.0232296.g004>

predictions over unseen subjects (as shown in the right panel of Fig 4), the improvements due to the introduction of regularity conditions begin to fade as the number of training subjects increases. In particular, beyond a certain number of training subjects (approximately 25 in the case of these experiments), the improvement in out-of-sample predictions begins to plateau.

3.2 Resting-state fMRI data experiments

While the previous section presented results relating to synthetic data, here we present experimental results where the proposed two-step procedure is applied to three open-source resting-state fMRI datasets. The datasets considered correspond to the Cambridge Center for Ageing and Neuroscience (CamCAN) repository, the Human Connectome Project (HCP) repository, and the ATR Wide-Age-Range repository. The purpose of employing three distinct datasets is to effectively measure the generalization performance of the proposed approach on unseen data. As such, data from the HCP and Wide-Age-Range repositories was not employed during any of the model training and instead used exclusively as unseen test data. It is important to note that in addition to significant inter-subject variability [45], fMRI data also suffers from the presence of several other well-documented issues such as variable scanner performance or noise [46, 47, 48]. As such, validating the performance of the proposed brain age prediction models in this way will provide a more realistic measure of their generalization performance.

3.2.1 CamCAN repository results. Resting-state fMRI data was collected from a total of 647 subjects from the CamCAN repository. Subject ages ranged from 18 to 88 years of age (average age of 54.31 ± 18.56 , 318 males and 329 females). The CamCAN dataset was employed as the principal dataset in the proposed two-step procedure, implying that it was employed to learn both the functional network structure in the unsupervised learning stage and the linear regression models in the supervised learning stage. As such, the data was split into training, validation and test subsets as described in Section 2.3.

Step 1: Unsupervised functional network inference. The first stage of the proposed framework involves the estimation of reproducible functional connectivity networks via the use of the various linear latent variable models discussed in Section 2.1. The number of functional networks inferred corresponds directly to the dimensionality of latent variables, which is determined by hyper-parameter k . As each linear latent variable model can be interpreted as a probabilistic model, we select hyper-parameter k by maximizing the log-likelihood over the validation dataset. This resulted in the choice of $k = 5$ when the loading matrix was restricted to be both non-negative and orthonormal, as proposed by [26] and [29]. While it is possible that the choice of hyper-parameter may vary across distinct latent variable models (e.g., for PCA or factor analysis), we choose to keep the choice of k fixed across all models as this facilitates model comparison and interpretation of results.

The left panel of Fig 5 visualizes the results when the MHA linear latent variable model was employed (Figures produced using the `plot_glass_brain` function from the `nilearn` python module [49]). We note that, as discussed in Section 2.1, the MHA linear latent variable model effectively clusters regions into sub-networks via the introduction of non-negativity and orthonormality constraints. As such, each plot in the left panel of Fig 5 visualizes spatially remote brain regions which have been clustered together, indicating that these regions share strong positive correlations. We note that these correlations (i.e., edges in a network) are omitted for clarity in Fig 5. The results demonstrate that the inferred networks are spatially homogeneous and symmetric across both hemispheres. Furthermore, many of the inferred networks correspond to widely reported networks and regions: network 1 captures the default model network (DMN) and network 2 overlaps with the salience network, while networks 3 and 4 correspond to a higher-level visual network and the somatomotor network respectively. For

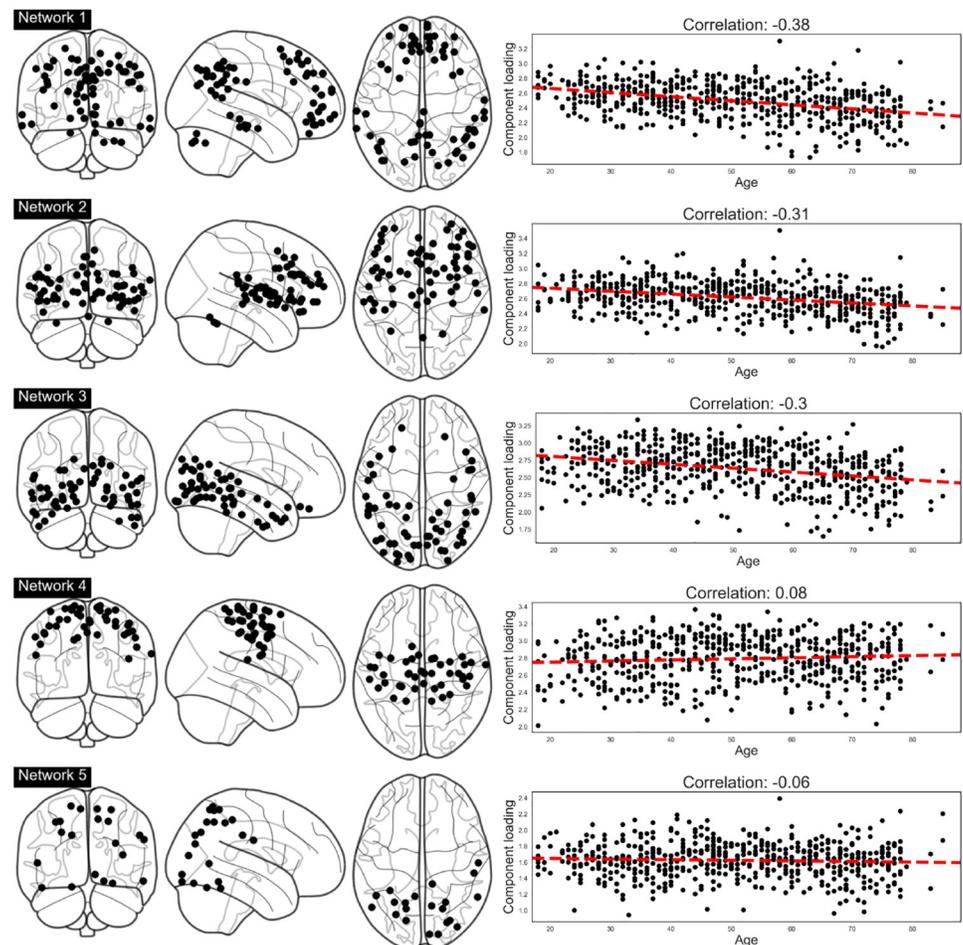


Fig 5. Left panel: Inferred networks as recovered when non-negativity and orthonormality constraints are introduced over the loading matrix, W . Networks are spatially consistent and symmetric. Right panel: visualization of network activities against subject age demonstrating (mostly negative) linear trends with healthy aging.

<https://doi.org/10.1371/journal.pone.0232296.g005>

comparison, we include equivalent plots for all other latent variable models considered in visualized in Fig 6, presented in the Supplementary Material. We note that alternative methods, such as PCA, which did not enforce the combination of both non-negativity and orthonormality, yielded results which were visibly less clustered and more difficult to interpret.

The right panel of Fig 5 visualizes the correlation between the activity of each network (as defined in Eq (11)) with the age of each subject. For networks 1-3 we observe a significant negative correlation between the activity and age, suggesting that ageing induces a drop in activity of such networks. These results are in line with related research on ageing induced differences in functional connectivity. In particular, the decrease in activity of the DMN (network 1), has been widely reported [19, 50, 51].

Step 2: Supervised training of brain age prediction models. Recall that the overall objective of the proposed framework was build interpretable models of biological brain age. To this end, the features recovered from linear latent variable models were employed as features in a linear regression framework to predict the brain age of each subject. In particular, the five distinct the linear latent variable models detailed in Section 2.1 were employed to learn reproducible sub-networks parameterized by a loading matrix, $W \in \mathbb{R}^{p \times k}$. The activity within each

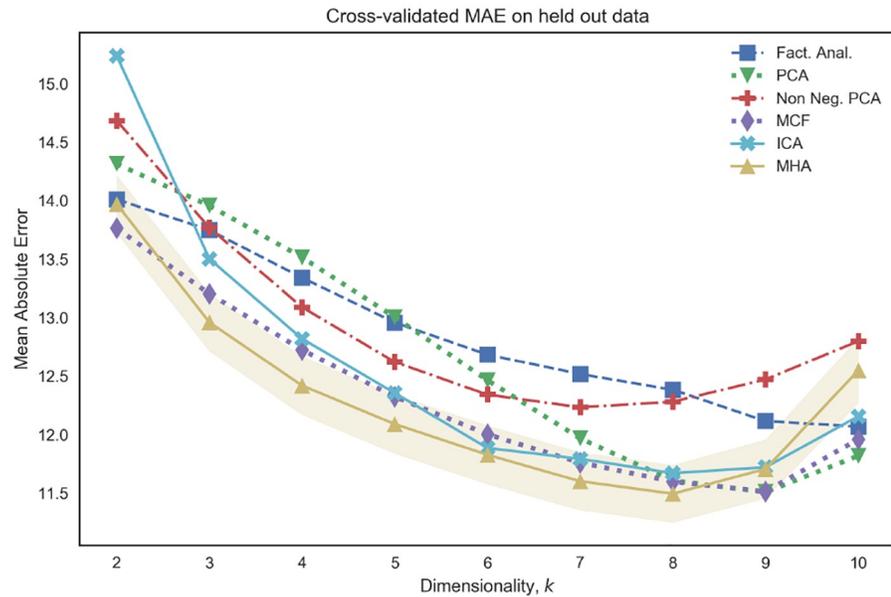


Fig 6. Inferred networks using alternative linear latent variable models. In the case of models such as PCA and factor analysis, networks were obtained by thresholding entries of W so only non-negative entries considered.

<https://doi.org/10.1371/journal.pone.0232296.g006>

functional network, defined as in Eq (11), was subsequently employed as features to predict biological age using linear regression.

We note that the CamCAN repository, as well as HCP and ATR repositories, each contained over a hundred subjects each. This is in contrast to typical fMRI studies, where the sample size is often in the range of 20 to 30 subjects [52, 48]. Furthermore, recall that the goal of experiments presented are to quantify performance on unseen resting-state fMRI data with a view to providing an indication of how each of the linear latent variable models employed would perform in a typical fMRI study. As such, throughout the remainder of this section we report the performance, in terms of mean absolute error, over random subsets of 30 subjects from each repository. This corresponds to a form of bootstrapping, where we average results over a random sample of possible cohorts. In practice, we report results over 1000 random subsets of 30 subjects for each of the three repositories considered.

Fig 7 visualizes the mean absolute error on unseen test data for various choices of $k \in \{2, \dots, 10\}$. We note that the combination of linear regression with the use of non-negativity and orthonormality constraints, as advocated by both the MCF and MHA models,

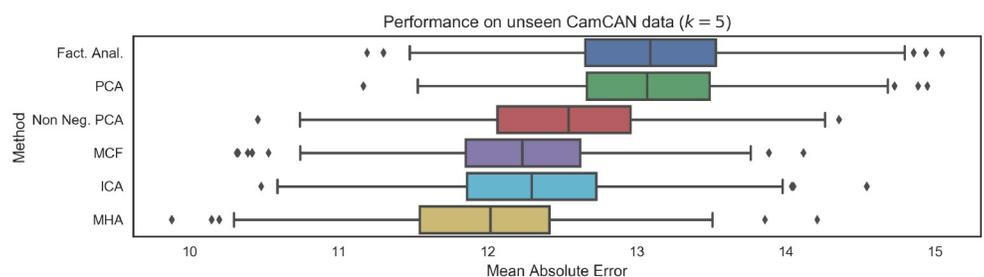


Fig 7. Mean Absolute Error (MAE) performance for a varying number of networks, as determined by k (x-axis), on unseen test data from CamCAN. We note that the combination of non-negativity and orthonormality (MHA and MCF) yields competitive results across a wide range of k .

<https://doi.org/10.1371/journal.pone.0232296.g007>

leads to competitive performance over a range of choices of k . In particular, such algorithms out-perform both non-negative PCA and PCA, suggesting that the introduction of such constraints serves to improve the predictive properties of the model. Moreover, we note that Fig 7 indicates the presence of a bias-variance trade-off that is often encountered in supervised learning whereby performance on unseen test data begins to deteriorate as the number of parameters (in our case k) increases beyond a certain value.

As mentioned previously, the choice of $k = 5$ was selected in by maximizing log-likelihood over a validation dataset (i.e., in an entirely unsupervised manner—data regarding subject ages was not considered). Fig 8 visualizes the performance on the unseen test dataset for the specific choice of $k = 5$, for all possible choices of linear latent variable models. The results indicate that as additional constraints are introduced to the loading matrix, the generalization capabilities of the models also improve. As such, MCF and MHA, which introduce the most stringent constraints corresponding to *both* non-negativity and orthonormality, obtain the best generalization performance. We also note that ICA is also competitive. Moreover, non-negative PCA, which relaxes the requirement for orthonormality, is the next most competitive latent variable model. Finally, PCA and factor analysis, which relax all the aforementioned constraints, obtain the worst generalization performance.

3.2.2 Transfer onto HCP and ATR Wide-Age-Range repositories. The results of Section 3.2.1 provide a measure of performance, in terms mean absolute error in predicted brain age, within a large-scale resting-state fMRI dataset. However, it is widely accepted that in addition subject-specific noise, there are several other significant contributors to noise in fMRI data: these include issues related to scanner noise and frequency of acquisition of images [46, 47, 48]. As a result, in order to thoroughly verify the generalization performance of the proposed methods, we employ resting-state fMRI data from the HCP and ATR Wide-Age-Range repositories. We note that data from the aforementioned repositories was employed only for testing purposes, as such it was not employed to learn the network structure across subjects, nor to tune the parameters of the linear regression models. For a summary of the characteristics of HCP and ATR Wide-Age-Range datasets see Fig 9 and S1 Table in the Supplementary Material.

Prediction of biological age on both the HCP and ATR Wide-Age-Range repositories was performed as follows: First, the loading matrix, \hat{W} was employed to obtain estimated activity within each network, as detailed in Eqs (11) and (12). Subsequently, predictions of biological age were obtained using Eq (10). At each stage both \hat{W} and $\hat{\beta}$ are the parameters inferred using the CamCAN dataset (i.e., there was no fine-tuning of parameters). As a result, performance on both HCP and ATR Wide-Age-Range datasets provide a robust measure of generalization performance to entirely unseen data.

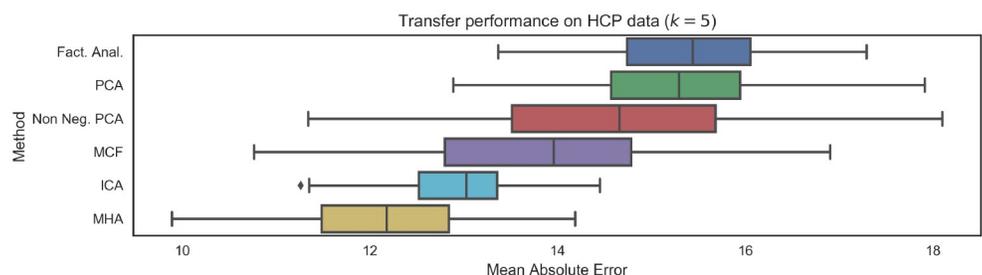


Fig 8. Mean Absolute Error (MAE) performance on unseen testing data from CamCAN repository when the dimensionality of latent variables is fixed to $k = 5$ (implying we infer 5 networks). We note that as regularity constraints are introduced, in particular non-negativity and orthonormality, predictive performance improves.

<https://doi.org/10.1371/journal.pone.0232296.g008>

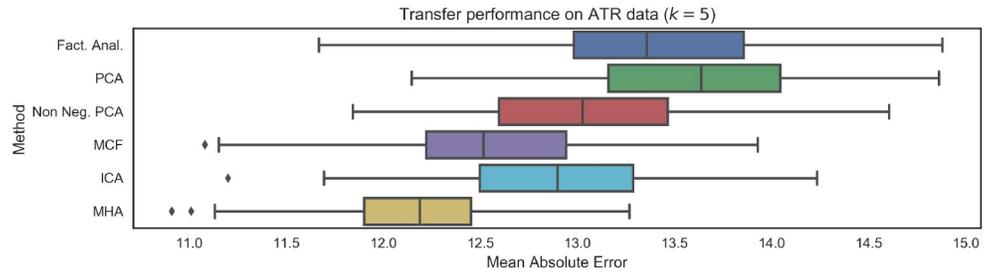


Fig 9. Histogram visualizing age distribution for each of the repositories employed. We note that the CamCAN dataset has the widest range of all repositories considered, validating its use as a the primary dataset in our study.

<https://doi.org/10.1371/journal.pone.0232296.g009>

Results on the HCP data are provided in Fig 10. As expected, the mean absolute errors are larger for each of the distinct latent variable models when compared to the results of on the CamCAN dataset (Fig 8), which will be partially the result of varying scanner noise and image acquisition properties. Importantly we note that, as with the CamCAN dataset, there once again a relationship between the introduction of additional constraints (in the form of non-negativity, orthonormality or non-Gaussianity) and generalization performance. As before, methods such as PCA and factor analysis which do not introduce any constraints had the weakest performance as well as the largest drop in performance.

The HCP results presented above serve to partially validate the predictive models trained using the CamCAN dataset. However, one significant limitation of the HCP dataset is that subject ages only range from 22 to 37 years of age. This is particularly relevant in the context of brain age biomarkers, as many neurodegenerative diseases of interest will be associated with advanced ages. As a result, we further validated the generalization capabilities of the proposed

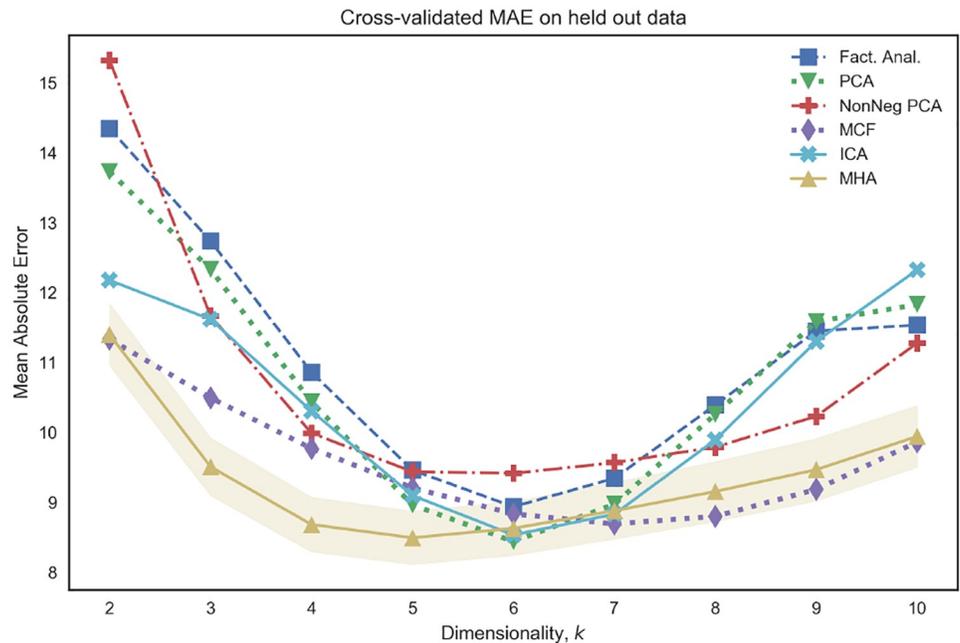


Fig 10. Mean Absolute Error (MAE) performance on unseen data from HCP repository. Results are broadly consistent with performance on the CamCAN data, indicating good generalization. We note that the introduction of non-negativity or orthogonality constraints leads to improved generalization. The number of functional networks was $k = 5$.

<https://doi.org/10.1371/journal.pone.0232296.g010>

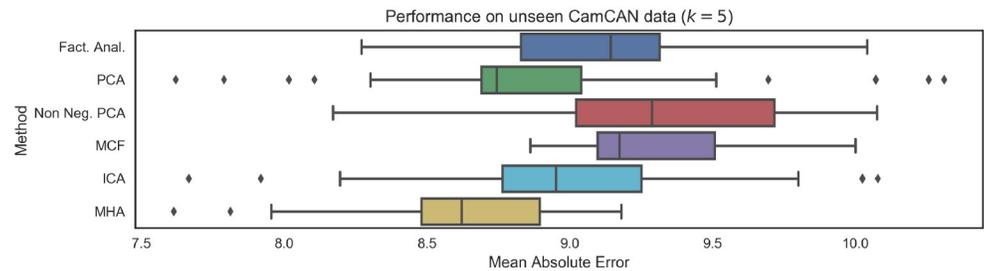


Fig 11. Mean Absolute Error (MAE) performance on unseen data from ATR Wide-Age-Range repository. Results are broadly consistent with performance on the CamCAN data, indicating good generalization. Further, as with the HCP data, we note that the introduction of non-negativity or orthogonality constraints leads to improved generalization. The number of functional networks considered was $k = 5$.

<https://doi.org/10.1371/journal.pone.0232296.g011>

brain age prediction models on the ATR Wide-Age-Range dataset, which had subjects ranging from 20 to 70 years of age. Results, presented in Fig 11 are consistent with results on the CamCAN and HCP datasets, again indicating that the introduction of constraints non-negativity and orthonormality constraints improves generalization performance.

3.3 Extension to non-independent latent variable models

The results presented above employ linear latent variable models where the inferred latents are assumed to be independent. This is clearly stated in the generative model considered in Eq (1) where the covariance of latent variables, $G^{(i)}$, is assumed to be diagonal. Note that in the case of PCA, factor analysis and MHA, since latent variables are assumed to be multivariate Gaussian, the fact the covariance is diagonal implies the latent variables are independent. However, such an assumption will often fail in practice, implying that the empirical covariance structure over latent variables will not be diagonal. In this section we seek to exploit this by directly introducing the off-diagonal entries of the latent variable covariances, $G^{(i)}$, as features in our linear regression models for biological age. As such, whilst Eq (10) considered a linear model where only the diagonal entries of each $G^{(i)}$ were employed to predict biological ages of each subject, we now consider linear regression models of the following form:

$$a^{(i)} = \sum_{j=1}^k \sum_{l \geq j} \beta_{jl} g_{jl}^{(i)} + \epsilon^{(i)}. \quad (15)$$

Note that in Eq (15) we employ the full upper triangular entries of the covariance matrix as features. This is equivalent to vectorizing the covariance matrix and removing duplicate entries due to symmetry. As such, whilst k features were employed in Eq (10), we now consider a linear models with $\binom{k}{2}$ features; many of which will seek to predict the biological age of individuals based on the off diagonal entries of each $G^{(i)}$. It is important to note that the model presented in Eq (10) is a special case of Eq (15).

As in Section 3.2, we proceed in a two-stage approach whereby we first estimate the loading matrices for the various linear latent variable models employed and subsequently train linear regression models using the full vectorized covariance matrix as features.

Fig 12 visualizes the MAE error on unseen test data as a function of the dimensionality of latent variables, k . We note that for all choices of k the reported errors are smaller than those reported in Fig 7. This provides empirical evidence that the off-diagonal entries of the latent variable covariances are discriminative features for brain age prediction, and therefore can be seen as evidence that models which assume diagonal covariance structure over latents are

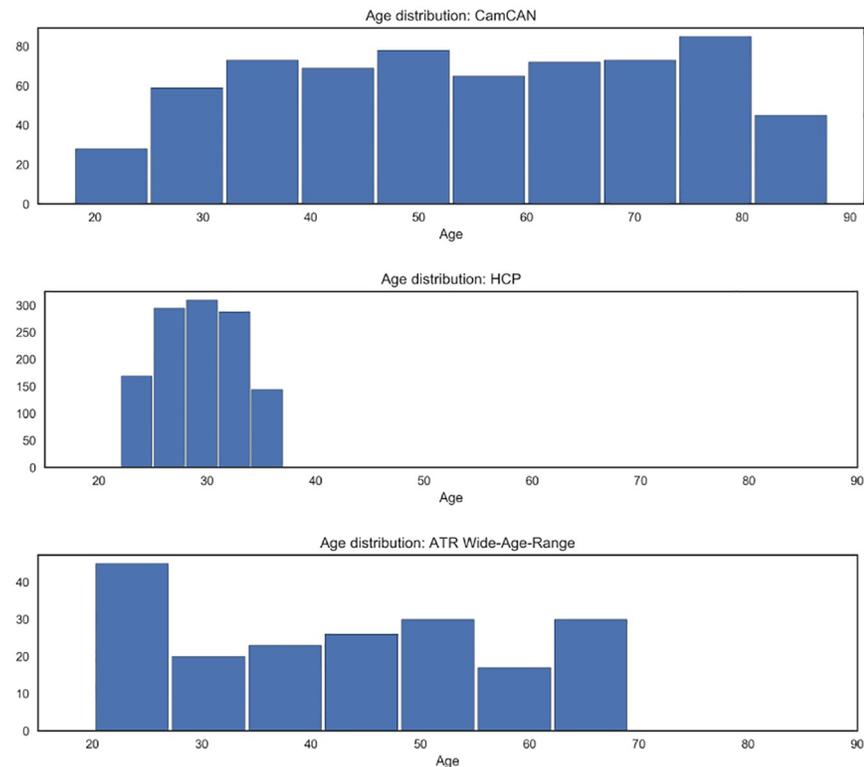


Fig 12. Mean Absolute Error (MAE) performance for a varying number of networks, as determined by k (x-axis), on unseen test data from CamCAN when latent variables are no longer assumed to have an isotropic covariance structure and the full vectorized covariance is employed as features in the linear regression models. We note that MHA is able to directly accommodate such a scenario and hence is competitive for all choices of latent variable dimension, k .

<https://doi.org/10.1371/journal.pone.0232296.g012>

misspecified. Fig 13 provides further visualizations in the case where $k = 5$. We note that the MHA model performs competitively, this is to be expected as this model directly accommodates the possibility of non-independent latent variables [29]. Moreover, we note that MHA performs particularly well when the number of networks is small (when dimension of latent variables, k , is less than or equal to 5), which is useful when we wish to prioritize the

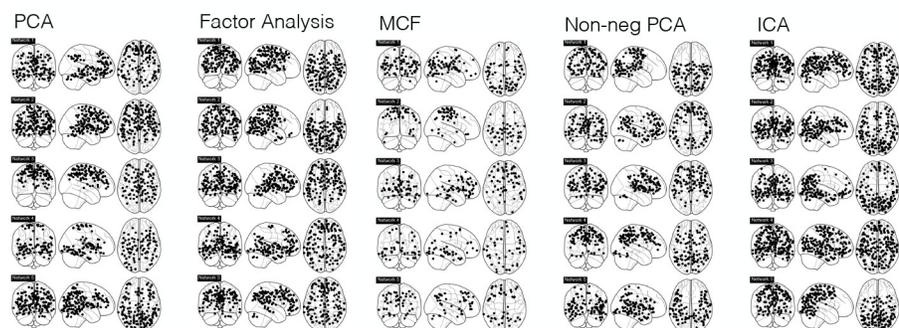


Fig 13. Mean absolute error (MAE) performance on unseen testing data from CamCAN repository when the dimensionality of latent variables is fixed to $k = 5$ (implying we infer 5 networks). Note that latent variables are no longer assumed to have an isotropic covariance structure and the full vectorized covariance is employed as features in the linear regression models.

<https://doi.org/10.1371/journal.pone.0232296.g013>

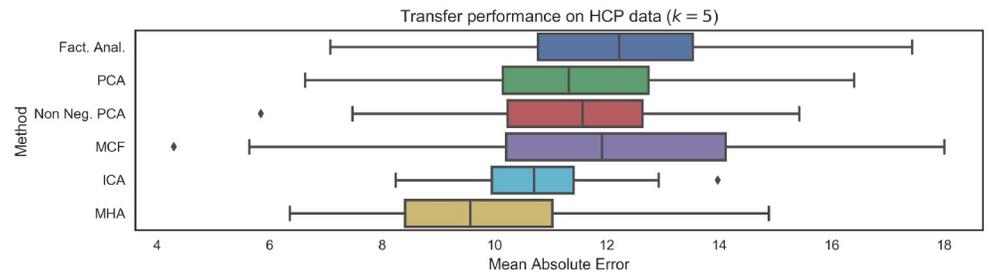


Fig 14. Mean Absolute Error (MAE) performance on unseen data from HCP repository. Results are broadly consistent with performance on the CamCAN data, indicating good generalization. We note that the introduction of non-negativity or orthogonality constraints leads to improved generalization. The number of functional networks was $k = 5$.

<https://doi.org/10.1371/journal.pone.0232296.g014>

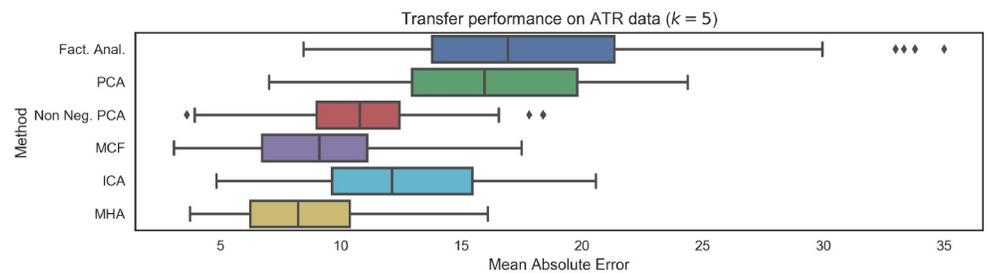


Fig 15. Mean absolute error (MAE) performance on unseen data from ATR Wide-Age-Range repository. Results are broadly consistent with performance on the CamCAN data, indicating good generalization. Further, as with the HCP data, we note that the introduction of non-negativity or orthogonality constraints leads to improved generalization. The number of functional networks considered was $k = 5$.

<https://doi.org/10.1371/journal.pone.0232296.g015>

interpretability of results. Finally, the performance of various methods, as depicted in Fig 12, shows similar trends as in Fig 7; there is once again a bias-variance trade-off associated with the choice of k and the introduction of non-negativity or non-Gaussianity constraints (as in MCF or ICA) leads to improved generalization performance. Finally, whilst Fig 12 only shows generalization performance to unseen subjects from the CamCAN cohort, we also present results for generalization performance to brain age prediction on the HCP and ATR Wide-Age-Range datasets in Figs 14 and 15 of the Supplementary Material.

4 Conclusion

It is widely accepted that ageing has pronounced effects on the functional architecture of the human brain [14, 9]. In the current study we have presented and validated a two-stage framework through which to train interpretable and robust models of biological brain age based on functional connectivity. In particular, the proposed framework first employs linear latent variable models to uncover reproducible networks which are present throughout a cohort of subjects. A variety of such latent variable models are considered many of which extend PCA by introducing constraints such as non-negativity over the loading matrix. Our experiments suggest that whilst PCA is a natural candidate for dimensionality reduction, and can be interpreted as recovering latent *eigenconnectivities*, the introduction of constraints such as non-negativity can serve to greatly improve both interpretability and predictive performance. While ICA improves on PCA by introducing spatial sparsity, we found that MHA as well as MCF lead to better results, especially in the case of a small number of networks. Reasons for

this improvement include using a combination of non-negativity and orthogonality that leads to disjoint networks, as well as explicit modelling of connectivity between the components.

Given inferred functional networks and their activations we train linear predictive models of biological brain age where in the interest of interpretability we deliberately restrict ourselves to linear models. This allows us to directly interrogate the effects of each functional network on the predicted brain age (as shown in Fig 5). In line with other results in the literature, we find a decrease in activation in the default mode network, salience network and higher-level visual network as biological age increases.

The proposed two-stage framework is first validated on the data from the CamCAN repository and subsequently further applied to two further open-access repositories: the HCP and ATR Wide-Age-Range repositories. The use of data from two additional repositories serves to provide a clear empirical indication of the generalization capabilities of the proposed approach. This is especially relevant in the context of fMRI data, where artefacts such as scanner noise can often cause significant challenges [48].

We note that the brain age prediction errors presented in this work are not competitive with alternative methods which are based on alternative imaging modalities, such as structural imaging data [53, 10]. This is to be expected for two reasons. First, the imaging modality employed in this work, resting-state fMRI data, is both noisier and likely to be less age-indicative than structural measures. Second, in this work we deliberately restrict ourselves to building simple yet interpretable models of brain age. As such, we restrict ourselves to consider only linear classifiers as these allow for clear model interpretation and interrogation, while noting that the use of more expressive models (e.g., nonlinear models) in the second stage should naturally lead to improved performance.

Furthermore, it is important to note that whilst this work demonstrates the feasibility of functional connectivity driven models of biological brain age, all subjects included in these studies were healthy. As such, whilst such models could eventually be employed to develop biomarkers, further experimentation and validation will be required in future. Moreover, an avenue for further research would be to consider performing classification instead of regression in the second stage of the proposed method. Whilst a natural task would be to discriminate between healthy controls and subjects with some neuropathology, such an approach could also be employed in the context of task-based fMRI as well as to study changes in functional connectivity induced by various distinct tasks [54] or neuropathologies [55, 56]. In particular, task-based fMRI has been widely reported as displaying non-stationary functional connectivity structure [57, 58, 59, 60]. As such, seeking to discriminate between various cognitive tasks, for example as considered by [61], [62], [63, 64], could be an exciting future application. Moreover, while in this work we have considered linear latent variable models such as PCA, future work could consider alternative latent variable modes such as latent position graphs [65] and causal models [66, 67, 68].

Supporting information

S1 Appendix. Technical details of the MHA algorithm.

(PDF)

S1 Code. Python and R implementations of the MHA algorithm.

(PDF)

S1 Table. Table detailing number of subjects studied in each of the three datasets considered. In the case of the HCP datasets, 80 subjects were randomly selected out of all possible

subjects.
(PDF)

S1 Fig. Age distributions of subjects across repositories.
(PNG)

S2 Fig. Functional connectivity networks inferred by PCA and alternative models.
(PARTIAL)

S3 Fig. Generalization performance of brain age prediction on HCP and ATR Wide-Age-Range datasets.
(PNG)

Acknowledgments

The authors wish to thank Steve Smith for valuable feedback and discussions.

Author Contributions

Conceptualization: Ricardo Pio Monti, Alex Gibberd, Sandipan Roy, Matthew Nunes, Romy Lorenz, Robert Leech, Motoaki Kawanabe, Aapo Hyvärinen.

Data curation: Ricardo Pio Monti, Romy Lorenz, Robert Leech, Takeshi Ogawa, Motoaki Kawanabe, Aapo Hyvärinen.

Formal analysis: Ricardo Pio Monti, Sandipan Roy, Matthew Nunes, Romy Lorenz, Motoaki Kawanabe, Aapo Hyvärinen.

Funding acquisition: Robert Leech, Motoaki Kawanabe, Aapo Hyvärinen.

Investigation: Ricardo Pio Monti, Alex Gibberd, Sandipan Roy, Matthew Nunes, Romy Lorenz, Robert Leech, Motoaki Kawanabe, Aapo Hyvärinen.

Methodology: Ricardo Pio Monti, Alex Gibberd, Sandipan Roy, Matthew Nunes, Romy Lorenz, Robert Leech, Motoaki Kawanabe, Aapo Hyvärinen.

Project administration: Ricardo Pio Monti, Romy Lorenz, Motoaki Kawanabe, Aapo Hyvärinen.

Resources: Ricardo Pio Monti, Motoaki Kawanabe, Aapo Hyvärinen.

Software: Ricardo Pio Monti, Motoaki Kawanabe, Aapo Hyvärinen.

Supervision: Motoaki Kawanabe, Aapo Hyvärinen.

Validation: Ricardo Pio Monti, Alex Gibberd, Sandipan Roy, Matthew Nunes, Romy Lorenz, Robert Leech, Takeshi Ogawa, Motoaki Kawanabe, Aapo Hyvärinen.

Visualization: Ricardo Pio Monti, Alex Gibberd, Sandipan Roy, Matthew Nunes, Romy Lorenz, Robert Leech, Takeshi Ogawa, Motoaki Kawanabe, Aapo Hyvärinen.

Writing – original draft: Ricardo Pio Monti, Alex Gibberd, Sandipan Roy, Matthew Nunes, Romy Lorenz, Robert Leech, Takeshi Ogawa, Motoaki Kawanabe, Aapo Hyvärinen.

Writing – review & editing: Ricardo Pio Monti, Alex Gibberd, Sandipan Roy, Matthew Nunes, Romy Lorenz, Robert Leech, Takeshi Ogawa, Motoaki Kawanabe, Aapo Hyvärinen.

References

1. Lim S., Han C. E., Uhlhaas P. J., and Kaiser M. Preferential detachment during human brain development: age- and sex-specific structural connectivity in diffusion tensor imaging (dti) data. *Cerebral Cortex*, 25(6):1477–1489, 2013. <https://doi.org/10.1093/cercor/bht333> PMID: 24343892
2. Raz N. and Rodrigue K. M. Differential aging of the brain: patterns, cognitive correlates and modifiers. *Neuroscience & Biobehavioral Reviews*, 30(6):730–748, 2006. <https://doi.org/10.1016/j.neubiorev.2006.07.001>
3. Cole J. H., Ritchie S. J., Bastin M. E., Hernández M. V., Maniega S. M., Royle N., et al. Brain age predicts mortality. *Molecular Psychiatry*, 23(5):1385, 2018. <https://doi.org/10.1038/mp.2017.62> PMID: 28439103
4. Cole J. H., Leech R., Sharp D. J., and Initiative A. D. N. Prediction of brain age suggests accelerated atrophy after traumatic brain injury. *Annals of Neurology*, 77(4): 571–581, 2015. <https://doi.org/10.1002/ana.24367> PMID: 25623048
5. Koutsouleris N., Davatzikos C., Borgwardt S., Gaser C., Bottlender R., Frodl T., et al. Accelerated brain aging in Schizophrenia and beyond: a neuroanatomical marker of psychiatric disorders. *Schizophrenia bulletin*, 40(5):1140–1153, 2013. <https://doi.org/10.1093/schbul/sbt142> PMID: 24126515
6. Good C. D., Johnsrude I. S., Ashburner J., Henson R. N., Friston K. J., and Frackowiak R. S. A voxel-based morphometric study of ageing in 465 normal adult human brains. *Neuroimage*, 14(1):21–36, 2001. <https://doi.org/10.1006/nimg.2001.0786> PMID: 11525331
7. Franke K., Gaser C., Manor B., and Novak V. Advanced brainage in older adults with type 2 diabetes mellitus. *Frontiers in Aging Neuroscience*, 5:90, 2013. <https://doi.org/10.3389/fnagi.2013.00090> PMID: 24381557
8. Lancaster J., Lorenz R., Leech R., and Cole J. H. Bayesian optimization for neuroimaging pre-processing in brain age classification and prediction. *Frontiers in Aging Neuroscience*, 10:28, 2018. <https://doi.org/10.3389/fnagi.2018.00028> PMID: 29483870
9. Smith S. M., Vidaurre D., Alfaro-Almagro F., Nichols T. E., and Miller K. L. Estimation of brain age delta from brain imaging. *NeuroImage*, 2019. <https://doi.org/10.1016/j.neuroimage.2019.06.017>
10. Cole J. H., Poudel R. P., Tsagkrasoulis D., Caan M. W., Steves C., Spector T. D., et al. Predicting brain age with deep learning from raw imaging data results in a reliable and heritable biomarker. *NeuroImage*, 163:115–124, 2017. <https://doi.org/10.1016/j.neuroimage.2017.07.059> PMID: 28765056
11. Franke K., Ziegler G., Klöppel S., Gaser C., Initiative A. D. N., et al. Estimating the age of healthy subjects from t1-weighted MRI scans using kernel methods: exploring the influence of various parameters. *Neuroimage*, 50(3):883–892, 2010. <https://doi.org/10.1016/j.neuroimage.2010.01.005> PMID: 20070949
12. Dosenbach N. U., Nardos B., Cohen A. L., Fair D. A., Power J. D., Church J. A., et al. Prediction of individual brain maturity using fMRI. *Science*, 329(5997):1358–1361, 2010. <https://doi.org/10.1126/science.1194144> PMID: 20829489
13. Geerligs L. et al. Reduced specificity of functional connectivity in the aging brain during task performance. *Human Brain Mapping*, 35:319–330, 2012. <https://doi.org/10.1002/hbm.22175> PMID: 22915491
14. Geerligs L., Renken R. J., Saliassi E., Maurits N. M., and Lorist M. M. A brain-wide study of age-related changes in functional connectivity. *Cerebral Cortex*, 25(7): 1987–1999, 2014. <https://doi.org/10.1093/cercor/bhu012> PMID: 24532319
15. Sporns O. *Discovering the Human Connectome*. MIT press, 2012.
16. Wu T., Wang L., Chen Y., Zhao C., Li K., and Chan P. Changes of functional connectivity of the motor network in the resting state in Parkinson's disease. *Neurosci. Lett.*, 460(1):6–10, 2009.
17. Damoiseaux J. S., Prater K. E., Miller B. L., and Greicius M. D. Functional connectivity tracks clinical deterioration in Alzheimer's disease. *Neurobiology of Aging*, 33(4), 2012. <https://doi.org/10.1016/j.neurobiolaging.2011.06.024>
18. Cherkassky V. L., Kana R. K., Keller T. A., and Just M. A. Functional connectivity in a baseline resting-state network in Autism. *Neuroreport*, 17(16):1687–1690, 2006. <https://doi.org/10.1097/01.wnr.0000239956.45448.4c> PMID: 17047454
19. Geerligs L., Rubinov M., Henson R. N., et al. State and trait components of functional connectivity: individual differences vary with mental state. *Journal of Neuroscience*, 35(41):13949–13961, 2015. <https://doi.org/10.1523/JNEUROSCI.1324-15.2015> PMID: 26468196
20. Geerligs L., Tsvetanov K. A., and Henson R. N. Challenges in measuring individual differences in functional connectivity using fMRI: the case of healthy aging. *Human Brain Mapping*, 38(8):4125–4156, 2017. <https://doi.org/10.1002/hbm.23653> PMID: 28544076

21. Leonardi N., Richiardi J., Gschwind M., Simioni S., Annoni J. M., Schluep M., et al. Principal components of functional connectivity: A new approach to study dynamic brain connectivity during rest. *Neuroimage*, 83: 937–950, 2013. <https://doi.org/10.1016/j.neuroimage.2013.07.019> PMID: 23872496
22. Smith S. M. The future of fMRI connectivity. *Neuroimage*, 62(2):1257–1266, 2012. <https://doi.org/10.1016/j.neuroimage.2012.01.022> PMID: 22248579
23. Jolliffe I. *Principal component analysis*. Springer, 2011.
24. Hotelling H. Analysis of a complex of statistical variables into principal components. *Journal of educational psychology*, 24(6):417, 1933.
25. Harman H. H. *Modern Factor Analysis*. Univ. of Chicago Press, 1960.
26. Hirayama J., Hyvärinen A., Kiviniemi V., Kawanabe M., and Yamashita O. Characterizing variability of modular brain connectivity with constrained principal component analysis. *PLoS One*, 11(12): e0168180, 2016. <https://doi.org/10.1371/journal.pone.0168180> PMID: 28002474
27. Sigg C. D. and Buhmann J. M. Expectation-maximization for sparse and non-negative PCA. In *Proceedings of the 25th international conference on Machine learning*, pages 960–967. ACM, 2008.
28. Zass R. and Shashua A. Non-negative sparse PCA. In *Advances in Neural Information Processing Systems*, pages 1561–1568, 2007.
29. Monti R. P. and Hyvärinen A. A Unified Probabilistic Model for Learning Latent Factors and Their Connectivities from High-Dimensional Data. In *34th Conference on Uncertainty in Artificial Intelligence*, 2018.
30. Bishop C. M. *Pattern Recognition and Machine Learning*. Springer, 2006.
31. Hyvärinen A., Hirayama J. I., Kiviniemi V., and Kawanabe M. Orthogonal Connectivity Factorization: Interpretable Decomposition of Variability in Correlation Matrices. *Neural Computation*, 28(3):445–484, 2016. https://doi.org/10.1162/NECO_a_00810 PMID: 26735746
32. Hyvärinen A., Karhunen J., and Oja E. *Independent Component Analysis*. Wiley, 2001.
33. Himberg J., Hyvärinen A., and Esposito F. Validating the independent components of neuroimaging time series via clustering and visualization. *Neuroimage*, 22(3): 1214–1222, 2004. <https://doi.org/10.1016/j.neuroimage.2004.03.027> PMID: 15219593
34. Esposito F., Scarabino T., Hyvärinen A., Himberg J., Formisano E., Comani S., et al. Independent component analysis of fMRI group studies by self-organizing clustering. *Neuroimage*, 25(1):193–205, 2005. <https://doi.org/10.1016/j.neuroimage.2004.10.042> PMID: 15734355
35. de Ven van V. G., Formisano E., Prvulovic D., Roeder C. H., and Linden D. E. Functional connectivity as revealed by spatial independent component analysis of fMRI measurements during rest. *Human Brain Mapping*, 22(3):165–178, 2004. <https://doi.org/10.1002/hbm.20022>
36. Taylor J., Williams N., Cusack R., Auer T., Shafto M., Dixon M., et al. The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample. *NeuroImage*, 18, 2015.
37. Van Essen D. C., Smith S. M., Barch D. M., Behrens T. E., Yacoub E., Ugurbil K., et al. The WU-Minn Human Connectome Project: an overview. *Neuroimage*, 80:62–79, 2013. <https://doi.org/10.1016/j.neuroimage.2013.05.041> PMID: 23684880
38. Ogawa T., Aihara T., Shimokawa T., and Yamashita O. Large-scale brain network associated with creative insight: combined voxel-based morphometry and resting-state functional connectivity analyses. *Scientific reports*, 8(1):6477, 2018. <https://doi.org/10.1038/s41598-018-24981-0> PMID: 29691468
39. Smith S. M., Jenkinson M., Woolrich M. W., Beckmann C. F., Behrens T. E., Johansen-Berg H., et al. Advances in functional and structural MR image analysis and implementation as FSL. *Neuroimage*, 23: S208–S219, 2004. <https://doi.org/10.1016/j.neuroimage.2004.07.051> PMID: 15501092
40. Ashburner J. Computational anatomy with the SPM software. *Magnetic Resonance Imaging*, 27(8):1163–1174, 2009. <https://doi.org/10.1016/j.mri.2009.01.006> PMID: 19249168
41. Salimi-Khorshidi G., Douaud G., Beckmann C. F., Glasser M. F., Griffanti L., and Smith S. M. Automatic denoising of functional MRI data: combining independent component analysis and hierarchical fusion of classifiers. *Neuroimage*, 90:449–468, 2014. <https://doi.org/10.1016/j.neuroimage.2013.11.046> PMID: 24389422
42. Power J. D., Cohen A. L., Nelson S. M., Wig G. S., Barnes K. A., Church J. A., et al. Functional network organization of the human brain. *Neuron*, 72(4):665–678, 2011. <https://doi.org/10.1016/j.neuron.2011.09.006> PMID: 22099467
43. Hyvärinen A. Fast and robust fixed-point algorithms for independent component analysis. *IEEE transactions on Neural Networks*, 10(3):626–634, 1999. <https://doi.org/10.1109/72.761722> PMID: 18252563
44. Pedregosa F., Varoquaux G., Gramfort A., Michel V., Thirion B., Grisel O., et al. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

45. Kelly C., Biswal B. B., Craddock R. C., Castellanos F. X., and Milham M. P. Characterizing variation in the functional connectome: promise and pitfalls. *Trends in Cognitive Sciences*, 16(3):181–188, 2012. <https://doi.org/10.1016/j.tics.2012.02.001> PMID: 22341211
46. Bennett C. M. and Miller M. B. How reliable are the results from functional magnetic resonance imaging? *Annals of the New York Academy of Sciences*, 1191(1):133–155, 2010. <https://doi.org/10.1111/j.1749-6632.2010.05446.x> PMID: 20392279
47. Friedman L., Glover G. H., Consortium F., et al. Reducing interscanner variability of activation in a multi-center fMRI study: controlling for signal-to-fluctuation-noise-ratio (SFNR) differences. *Neuroimage*, 33(2):471–481, 2006. <https://doi.org/10.1016/j.neuroimage.2006.07.012> PMID: 16952468
48. Poldrack R. A., Mumford J. A., and Nichols T. E. *Handbook of functional MRI data analysis*. Cambridge University Press, 2011.
49. Abraham A., Pedregosa F., Eickenberg M., Gervais P., Mueller A., et al. Machine learning for neuroimaging with scikit-learn. *Frontiers in neuroinformatics*, 8:14, 2014. <https://doi.org/10.3389/fninf.2014.00014> PMID: 24600388
50. Grady C., Sarraf S., Saverino C., and Campbell K. Age differences in the functional interactions among the default, frontoparietal control, and dorsal attention networks. *Neurobiology of Aging*, 41:159–172, 2016. <https://doi.org/10.1016/j.neurobiolaging.2016.02.020> PMID: 27103529
51. Liem F., Geerligs L., Damoiseaux J. S., and Margulies D. S. Functional Connectivity in Aging, 2019.
52. Cremers H. R., Wager T. D., and Yarkoni T. The relation between statistical power and inference in fMRI. *PLoS one*, 12(11):e0184923, 2017. <https://doi.org/10.1371/journal.pone.0184923> PMID: 29155843
53. Cole J. H. and Franke K. Predicting age using neuroimaging: innovative brain ageing biomarkers. *Trends in Neurosciences*, 40(12):681–690, 2017. <https://doi.org/10.1016/j.tins.2017.10.001> PMID: 29074032
54. Zippo A. G., Castiglioni I., Lin J., Borsa V. M., Valente M., and Biella G. E. Short-term classification learning promotes rapid global improvements of information processing in human brain functional connectome. *Frontiers in Human Neuroscience*, 13:462, 2019a. <https://doi.org/10.3389/fnhum.2019.00462>
55. Lorenz R., Violante I. R., Monti R. P., Montana G., Hampshire A., and Leech R. Dissociating frontoparietal brain networks with neuroadaptive bayesian optimization. *Nature communications*, 9(1):1–14, 2018. <https://doi.org/10.1038/s41467-018-03657-3>
56. Zippo A. G., Del Grosso V., Patera A., Riccardi M. P., Tredici I. G., Bertoli G., et al. Chronic pain alters microvascular architectural organization of somatosensory cortex. *bioRxiv*, page 755132, 2019b.
57. Calhoun V. D., Miller R., Pearlson G., and Adali T. The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron*, 84(2): 262–274, 2014. <https://doi.org/10.1016/j.neuron.2014.10.015> PMID: 25374354
58. Monti R. P., Hellyer P., Sharp D., Leech R., Anagnostopoulos C., and Montana G. Estimating time-varying brain connectivity networks from functional MRI time series. *NeuroImage*, 103:427–443, 2014. <https://doi.org/10.1016/j.neuroimage.2014.07.033> PMID: 25107854
59. Monti R. P., Anagnostopoulos C., Montana G., et al. Learning population and subject-specific brain connectivity networks via mixed neighborhood selection. *The Annals of Applied Statistics*, 11(4):2142–2164, 2017a. <https://doi.org/10.1214/17-AOAS1067>
60. Monti R. P., Lorenz R., Braga R. M., Anagnostopoulos C., Leech R., and Montana G. Real-time estimation of dynamic functional connectivity networks. *Human Brain Mapping*, 38(1):202–220, 2017b. <https://doi.org/10.1002/hbm.23355>
61. Chung A. W., Pesce E., Monti R. P., and G. Montana. Classifying hcp task-fMRI networks using heat kernels. In *2016 International Workshop on Pattern Recognition in Neuroimaging (PRNI)*, pages 1–4. IEEE, 2016.
62. Lorenz R., Simmons L. E., Monti R. P., Arthur J. L., Limal S., Laakso I., et al. Efficiently searching through large tacs parameter spaces using closed-loop bayesian optimization. *Brain stimulation*, 12(6):1484–1489, 2019. <https://doi.org/10.1016/j.brs.2019.07.003> PMID: 31289013
63. Monti R., Lorenz R., Hellyer P., Leech R., Anagnostopoulos C., and G. Montana. Graph embeddings of dynamic functional connectivity reveal discriminative patterns of task engagement in hcp data. In *2015 International Workshop on Pattern Recognition in Neuroimaging*, pages 1–4. IEEE, 2015.
64. Monti R. P., Lorenz R., Hellyer P., Leech R., Anagnostopoulos C., and Montana G. Decoding time-varying functional connectivity networks via linear graph embedding methods. *Frontiers in Computational Neuroscience*, 11:14, 2017c. <https://doi.org/10.3389/fncom.2017.00014>
65. Athreya A., Fishkind D. E., Tang M., Priebe C. E., Park Y., Vogelstein J. T., et al. Statistical inference on random dot product graphs: a survey. *The Journal of Machine Learning Research*, 18(1):8393–8484, 2017.

66. Khemakhem I., Kingma D. P., Monti R. P., and Hyvärinen A. Variational autoencoders and nonlinear ica: A unifying framework. *arXiv preprint arXiv:1907.04809*, 2019.
67. Monti R. P., Zhang K., and Hyvärinen A. Causal discovery with general non-linear relationships using non-linear ica. *arXiv preprint arXiv:1904.09096*, 2019.
68. Sasaki H., Takenouchi T., Monti R., and Hyvärinen A. Robust contrastive learning and nonlinear ica in the presence of outliers. *arXiv preprint arXiv:1911.00265*, 2019.