Phenotype discovery from population brain imaging

Weikang Gong [a,*], Christian F. Beckmann [a,b,c], Stephen M. Smith [a]

[a] Centre for Functional MRI of the Brain (FMRIB), Nuffield Department of Clinical Neurosciences, Wellcome Centre for Integrative Neuroimaging, University of Oxford, Oxford, UK
[b] Radboud University Medical Centre, Department of Cognitive Neuroscience, Nijmegen, the Netherlands
[c] Donders Institute for Brain, Cognition and Behaviour, Radboud University Nijmegen, Nijmegen, the Netherlands

ARTICLE INFO

ABSTRACT

Neuroimaging allows for the non-invasive study of the brain in rich detail. Data-driven discovery of patterns of population variability in the brain has the potential to be extremely valuable for early disease diagnosis and understanding the brain. The resulting patterns can be used as imaging-derived phenotypes (IDPs), and may complement existing expert-curated IDPs. However, population datasets, comprising many different structural and functional imaging modalities from thousands of subjects, provide a computational challenge not previously addressed. Here, for the first time, a multimodal independent component analysis approach is presented that is scalable for data fusion of voxel-level neuroimaging data in the full UK Biobank (UKB) dataset, that will soon reach 100,000 imaged subjects. This new computational approach can estimate modes of population variability that enhance the ability to predict thousands of phenotypic and behavioural variables using data from UKB and the Human Connectome Project. A high-dimensional decomposition achieved improved predictive power compared with widely-used analysis strategies, single-modality decompositions and existing IDPs. In UKB data (14,503 subjects with 47 different data modalities), many interpretable associations with non-imaging phenotypes were identified, including multimodal spatial maps related to fluid intelligence, handedness and disease, in some cases where IDP-based approaches failed.

1. Introduction

Large-scale multimodal brain imaging has enormous potential for boosting epidemiological and neuroscientific studies, generating markers for early disease diagnosis and prediction of disease progression, and the understanding of human cognition, by means of linking to clinical or behavioural variables. Recent major studies have been acquiring brain magnetic resonance imaging (MRI), genetics and demographic/behavioural data from large cohorts. Examples are the UK Biobank (UKB) (Miller et al., 2016), the Human Connectome Project (HCP) (Van Essen et al., 2013) and the Adolescent Brain Cognitive Development (ABCD) study (Jernigan et al., 2018). These studies involve multimodal data, meaning that several distinct types of MRI data are acquired, mapping activity, functional networks, structural connectivity, white matter microstructure, and organisation and volumes of different brain tissues and sub-structures (Miller et al., 2016). However, the multimodal, high-dimensional and noisy nature of such big datasets makes many existing analytical approaches for extracting interpretable information impractical (Smith and Nichols, 2018).

Traditionally, large-scale neuroimaging studies first summarize the imaging data into interpretable image-derived phenotypes (IDPs) (Miller et al., 2016; Elliott et al., 2018), which are scalar quantities derived from raw imaging data (e.g., regional volumes from structural MRI, mean task activations from task MRI, resting-state functional connectivities between brain parcels). This knowledge-based approach is simple and efficient, and effectively reduces the high-dimensional data into interpretable, compact, convenient features. However, there may well be a large loss of information, due to such "expert-hand-designed" features not capturing important sources of subject variability (or even just losing sensitivity by the pre-defined spatial sub-areas being suboptimal), as well as ignoring cross-modality relationships. Further, such uni-modal compartmentalised analyses do not utilise the fact that for many biological effects of interest we expect there to be biological convergence across different data modalities, i.e., changes in the underlying biological phenotype likely manifest themselves

* Corresponding author.
E-mail address: weikang.gong@ndcn.ox.ac.uk (W. Gong).

across multiple quantitative phenotypes, so that a joint analysis effectively increases both the power of detecting such effects and the interpretability of the findings.

In contrast to such uni-modal analyses, data-driven multivariate approaches (i.e., unsupervised machine learning) have been proposed, which perform simultaneous decomposition of voxel-level data directly, generally representing data as the summation of a number of "components" or "modes". Each mode is formed as the outer product of two vectors: one is a vector of subject weights (describing the relative strength of expression of that mode in each subject), and a vector of voxel weights (in effect a spatial map for each data modality, describing the spatial localisation of the mode). The subject weight vectors (one per mode) can be considered "features" (similar to IDPs, but being data-driven) for use in further modelling, such as for the prediction of non-imaging variables. They are often either based on eigendecomposition, such as multi-set canonical correlation analysis (mCCA) (Kettenring, 1971; Klami et al., 2015), or based on variations of independent component analysis (ICA) (Calhoun et al., 2006; Liu et al., 2009; Beckmann and Smith, 2005; Groves et al., 2011). Among them, FMRIB's Linked ICA (FLICA) (Groves et al., 2011) is an efficient approach which has been successfully applied to identify brain systems that are involved in lifespan development and diseases (Groves et al., 2012; Douaud et al., 2014), attention deficit hyperactivity disorder (Ball et al., 2018), preterm brain development (Ball et al., 2017) and cognition and psychopathology (Alnæs et al., 2018). FLICA has advantages compared with uni-modal analysis on IDPs, including: (1) It leverages the cross-modality information of multimodal data, so has the potential to detect patterns that are not discoverable in any single modality; (2) It is a data-driven objective approach which automatically discovers meaningful patterns in voxel-level multimodal data by searching for spatial non-Gaussian sources that have been shown to likely reflect real structured features in neuroimaging data (Griffanti et al., 2014). While this approach has been applied successfully to medium-sized cohort data (Groves et al., 2012; Douaud et al., 2014; Ball et al., 2018; 2017; Alnæs et al., 2018), the original algorithms for carrying out FLICA do not scale well with increasing data size, and are unable to analyze large datasets such as UKB, where dozens of different modalities over tens of thousands of subjects are available. Importantly, because the core FLICA algorithms are multivariate, acting in a complex way simultaneously across all subjects, modalities and voxels using Variational Bayesian updates of parameters, this problem cannot be solved through simple parallelisation or other algorithmically simple methods for distributing computations across a large cluster, and so cannot be addressed simply by increasing the number of processors or memory available.

To tackle this problem, we propose an approach that embeds advanced data compression techniques across the different data dimensions into the FLICA approach. We use a multimodal extension of MELODIC's Incremental Group Principal component analysis (Smith et al., 2014) (mMIGP, applied across modalities) and online dictionary learning (Mairal et al., 2010) (DicL, applied within-modalities) to efficiently reduce the size of multimodal neuroimaging data. The reduced data are then characterised through FLICA in terms of underlying modality-specifc maps and subject loading vectors. Here we refer to this combination of techniques as Bigdata FLICA, or BigFLICA for short. Two important advantages of the proposed approach are: (1) Preserving key information in original data but also reducing the effects of stochastic domain-specific noise; (2) Increasing the computational efficiency of the FLICA algorithm for extremely large population datasets. BigFLICA is scalable for simultaneously analyzing all the multimodal data of the full 100,000-subjects UKB dataset using only a modest computing cluster (Fig. 1).

We first demonstrate the effectiveness of our approach through extensive simulations. Then, in real data, we quantify performance in terms of the prediction accuracy of non-imaging-derived phenotypes (nIDPs) (Liégeois et al., 2019; Kong et al., 2018), such as health outcome measures. Using voxel-level imaging data of 81 modalities from 1003 subjects in the HCP and 47 modalities from 14,053 subjects in the UKB, we show that BigFLICA can perform comparably with original FLICA (Groves et al., 2011) in terms of the prediction accuracy for nIDPs (158 in HCP and 8787 in UKB). Most importantly, we systematically investigated whether there are benefits to jointly fusing multimodal data together, instead of analysing them separately. We show that significant improvements in the prediction accuracy of nIDPs are found when comparing a high-dimensional BigFLICA with other widely-used data analysis strategies: (1) doing single-modality ICA and concatenating the results across modalities and (2) using existing IDPs (5812 in HCP and 3913 in UKB). In particular, the improvements in prediction of many health outcome and cognitive variables are large, more than doubling prediction accuracy for some variables. Furthermore, we investigate the relationship between modes derived by BigFLICA and IDPs. We find that although the modes were estimated from the same set of voxel-level data, they have complementary information which can be combined together to further increase the prediction accuracy of nIDPs. Finally, we applied BigFLICA to analyze the UKB data and extracted 750 components. Existing multimodal ICA cannot estimate this many modes from this many subjects. We found several interpretable associations between modes of BigFLICA and nIDPs, including modes that relate to *fluid intelligence, handedness, age started wearing glasses or contact lenses* and *hypertension*. In many cases BigFLICA can find associations with nIDPs with greater statistical sensitivity than was possible with IDPs. Overall, BigFLICA demonstrated the advantages of data-driven joint multimodal modelling in the analysis of biobank-scale multimodal datasets.

2. Methods

2.1. Brief overview of the proposed approach: BigFLICA

FLICA (Groves et al., 2011) is a Bayesian ICA approach for multimodal data fusion. The input of FLICA is K modalities' data matrices $Y^{(k)}$ with dimensions $N \times P_k$, $k = 1, \ldots, K$, where P_k is the number of features (e.g., voxels) and N is the number of subjects. FLICA aims to find a joint L-dimensional decomposition of all $Y^{(k)}$: $Y^{(k)} = HW^{(k)}X^{(k)} + E^{(k)}$, where $H_{(N \times L)}$ is the shared subject mode (mixing matrix) across modalities (a vector of subject weights for each mode), so is a 'link' across different modalities, $W^{(k)}_{(L \times L)}$ is a positive diagonal mode-weights matrix (one overall weight per modality per mode), $X^{(k)}_{(L \times P_k)}$ is the independent (spatial) feature maps for the L components of a modality (one map per modality per mode), and $E^{(k)}_{(N \times P_k)}$ is the modality-specific Gaussian noise term (Fig. 1). We propose two efficient approaches that can either be used separately or combined together to reduce the size of the original data matrices, and therefore reduce the computational load of the original FLICA. An overview of BigFLICA is shown in Fig. 1.

The first approach, termed multimodal extension of MELODIC's Incremental Group Principal component analysis (Smith et al., 2014) (mMIGP), aims to reduce the subject dimension to a linear combination of the original subjects. mMIGP is a time- and memory-efficient approximation of principal component analysis (PCA) on feature-concatenated multimodal data. To this end, if we aim to get a $L^\star$ decomposition, we first apply MIGP (Smith et al., 2014) separately within each modality to estimate $\widetilde{U}^{(k)}_{(N \times L^\star)}$, which is

Fig. 1. Overview of the proposed approach for jointly analyzing a biobank-scale multimodal neuroimaging dataset. Currently for the UKB dataset (voxel-level data, 14,503 subjects, 47 modalities), the total data size is approximately 800 GB, and if we directly feed these data into FLICA and extract 750 components, we will need approximately 1066 GB CPU memory and 1680 h computation time. Our new approach, BigFLICA, used multimodal MIGP and dictionary learning to preprocess the multimodal data; this is efficient and memory friendly, and much of this preprocessing can be easily parallelized. BigFLICA only used 50 GB memory and 73 h to analyze the same dataset using a 24-core compute server.

an approximation of an $L^\star$-dimensional PCA decomposition of one modality $Y^{(k)}$. This step can be done in parallel across modalities. Then, we concatenate all $\widetilde{U}^{(k)}$ in the component dimension and apply another MIGP to get $U_{(N \times L^\star)}$, which is an $L^\star$-dimensional approximate PCA decomposition of all modalities together. Finally, we project the original data of each modality $Y^{(k)}$ to the PCA-reduced space using U. If no further reduction (e.g., dictionary learning as below) is to be applied, the data that could then be fed into the core FLICA would be the K component-by-feature matrices $V^{(k)}$ of size $L^\star \times P_k$, and FLICA would then extract L ($L < L^\star$) components from these (Methods). This step almost adds little computational cost compared with the original FLICA, because a similar PCA step is needed to initialize the parameters of the original FLICA, but this approach is feasible for large numbers of subjects and modalities. Although different modalities usually have different overall signal-to-noise ratios (SNR), which is largely ignored by this mMIGP step, the subsequent FLICA can take this into account by the modality-specific noise terms, and a high-dimensional mMIGP is used to capture modes with even small variations in each modality.

It is known that voxels are correlated in both a local fashion (local spatial autocorrelation) and across brain networks (long range correlation); hence, effective feature subsampling could hope to capture all important information in the data but also reduce the cost of spatial modelling in FLICA (Hoyos-Idrobo et al., 2019). Therefore, we incorporate an approach, termed sparse online Dictionary Learning (Mairal et al., 2010) (DicL), to reduce the dimension of feature (e.g., voxel) space that can capture both local and distant spatial correlation structure. Specifically, for each modality, we use DicL to model the $V^{(k)}$ as a sparse linear combination of $L^{\star\star}$ basis elements: $V^{(k)} = A^{(k)} D^{(k)}$, where $D^{(k)}$ is the *sparse spatial dictionary basis*, and $A^{(k)}$ is the *feature loadings*. By minimizing the reconstruction error, and enforcing sparsity in the dictionary basis $D^{(k)}$, we aim to achieve an optimal subsampling of feature space. The inputs of FLICA are then K smaller matrices $A^{(k)}$, which are only of dimension $L^\star \times L^{\star\star}$, and FLICA then extracts L ($L < \min(L^\star, L^{\star\star})$) components from these (**Methods**). Compared with doing FLICA with the original K large $N \times P_k$ matrices, using the DicL preprocessed data can greatly reduce the computation load of FLICA. DicL can easily be parallelized across modalities and is memory friendly, which further increases efficiency (Fig. 1).

2.2. FLICA model

The input to FLICA is K modalities' data matrices $Y^{(k)}$ with each modality's dimensions being $N \times P_k, k = 1, \dots, K$, where P_k is the number of features (e.g., voxels) in modality k and N is the number of subjects. FLICA aims to find a joint L-dimensional decomposition of all $Y^{(k)}$:

$$Y^{(k)} = HW^{(k)}X^{(k)} + E^{(k)}, \ \ k = 1, \dots, K \tag{1}$$

where $H_{(N \times L)}$ is the shared subject mode (mixing matrix) across modalities, so is a 'link' across different modalities, $W^{(k)}_{(L \times L)}$ is a positive diagonal mode-weights matrix, $X^{(k)}_{(L \times P_k)}$ is the independent (spatial) feature maps for the L components of modality k, and $E^{(k)}_{(N \times P_k)}$ is the Gaussian noise term.

2.3. Multimodal extension of MELODIC's incremental group principal component analysis for subject-space dimension reduction

We propose a multimodal extension of our previous MIGP approach (Smith et al., 2014), termed mMIGP, to reduce the subject dimension of multimodal data. MIGP has been extensively validated in simulations and real neuroimaging data for finding an approximate PCA decomposition in a time- and memory-efficient way (Smith et al., 2014). Suppose that our multimodal data are K matrices $Y^{(k)}, k = 1, \dots, K$ with dimensions $N \times P_k$, where N is the number of subjects and P_k is the number of features (e.g. voxels) in a modality. In mMIGP, each feature is z-score normalized first. Then, an MIGP is applied to each modality separately to find an $L^\star$-dimensional approximate PCA decomposition. Specifically, we want to find an approximation of a singular value decomposition (SVD) of each $Y^{(k)}$:

$$Y^{(k)} = \widetilde{U}^{(k)} \widetilde{S}^{(k)} (\widetilde{V}^{(k)})^\tau, \ \ k = 1, \dots, K \tag{2}$$

where $\widetilde{U}^{(k)}_{(N \times L^\star)}$ and $\widetilde{V}^{(k)}_{(P_k \times L^\star)}$ are the left and right singular vectors, while $\widetilde{S}^{(k)}_{(L^\star \times L^\star)}$ are the singular values. A naive SVD on $Y^{(k)}$ scales quadratically with N, which is not efficient when N is large. To find the approximation, MIGP sequentially feeds a subset of (columns of) $Y^{(k)}$ into an SVD, so that these subsets are reduced to a low-dimensional representation. The low-dimensional representation is

then concatenated with another subset of (columns of) $Y^{(k)}$, and is fed into another SVD to find the low-dimensional representation of them. The final SVD approximation is found after one pass of all data. The computational complexity of MIGP scales linearly with N. For a detailed description, please see Appendix A of the MIGP paper (Smith et al., 2014).

The third step is to concatenate all $\widetilde{U}^{(k)}$ in the component dimension and apply another MIGP for finding a $L^\star$-dimensional approximate PCA decompositions U of size $N \times L^\star$, which is a low-dimensional representation of multimodal data in the analysis. Finally, the z-score normalized data $Y^{(k)}$ of each modality is projected onto the U by:

$$V^{(k)} = U^\tau Y^{(k)} \, k = 1, 2, \ldots, K \tag{3}$$

the $V^{(k)}, k = 1, 2, \ldots, K$ are the inputs of the subsequent FLICA algorithm. Therefore, the total size of data output by this stage is $L^\star \sum_{k=1}^{K} P_k$, which is smaller than the original input size $N \sum_{k=1}^{K} P_k$. The fractional reduced data size is $L^\star/N$, and the $L^\star$ can be fixed when more subjects are introduced, so it is scalable in the big-data analysis. In practice, we usually choose $L^\star$ based on the percentage of explained variance of SVD in the third step.

If we feed $V^{(k)}, k = 1, 2, \ldots, K$ into FLICA to estimate $L^\star$ FLICA modes, the output subject mode matrix $H^\star$ is of the size $L^\star \times L$, so we then simply multiply this by U to get the final subject-mode matrix:

$$H = U H^\star \tag{4}$$

The mMIGP approach is equivalent to performing an approximate PCA on feature-concatenated data. The advantage is that it does not need to fit all data into the memory, and even can be parallelized across modalities (Smith et al., 2014). This approach is also equivalent to applying mCCA across all modalities (Parra, 2018).

2.4. Sparse dictionary learning for voxel-space dimension reduction

If the resolution of the data is high and the number of modalities is large, applying just the mMIGP reduction still leaves FLICA as being memory and computationally expensive. Therefore, we propose a method that can effectively reduce the voxel dimension, and preserve the important spatial information for subsequent FLICA spatial modelling. Although the most obvious ways of voxel subsampling are either to apply regular spatial downsampling (similarly, local voxel clustering) or apply PCA *within each modality*, the former only focuses on the local patterns (Hoyos-Idrobo et al., 2019) (and does not adapt downsampling to local variations in redundant information across voxels) while the later empirically finds more global and noise patterns in neuroimaging data, and does not work at all well empirically in this context (see also Allen et al., 2014 and references therein).

The method we used here is sparse Dictionary Learning (DicL) (Mairal et al., 2010), which effectively performs 'voxel grouping' in both local and global fashion. It can be used directly on each of the original *z*-score normalized modalities, i.e., $Y^{(k)}, k = 1, 2, \ldots, K$, or on the mMIGP reduced data, i.e., $V^{(k)}, k = 1, 2, \ldots, K$. Taking the former as an example, the sparse DicL is adopted here:

$$Y^{(k)} = A^{(k)} D^{(k)} \, k = 1, 2, \ldots, K \tag{5}$$

where $D^{(k)}$ is *sparse spatial dictionary basis*, and $A^{(k)}$ is the *feature loadings* with each column representing a linear combination of information from a group of voxels which might either be a local cluster or spatially distributed network. By minimizing an l_1-regularized sparse-coding objective function, a local optimal solution can be obtained:

$$l(A^{(k)}, D^{(k)}) = \sum_{i=1}^{P_k} \|Y_i^{(k)} - A^{(k)} D_i^{(k)}\|_F^2 + \lambda \|D_i^{(k)}\|_1$$
$$s.t. \ \ \mathcal{C} = \{A^{(k)} \in \mathbb{R}^{N \times L^{\star\star}} s.t. \ \forall j = 1, \ldots, L^{\star\star}, (A_j^{(k)})^\tau A_j^{(k)} \leqslant 1\} \tag{6}$$

where subscript i represents the ith column of the corresponding matrix, and λ is a regularization parameter. The l_1-regularization term enforces that the learned spatial loadings $D^{(k)}$ are sparse. The objective function can be efficiently optimized by a block-coordinate descent optimizer with warm restarts. It has been implemented in the SPAMS package (http://spams-devel.gforge.inria.fr/). Compared with simply using PCA in this step, sparse DicL has three advantages: (1) the spatial loading matrix $D^{(k)}$ can be sparse, so a smaller number of voxels are involved in each column of the dictionary; (2) the columns of the dictionary do not need to be orthogonal to each other, which is more flexible; (3) an "overcomplete" dictionary is allowed, i.e., the number of dictionary basis vectors can exceed the minimum of N and P_k, which further increases the flexibility.

After the above modality-wise DicL, the final inputs to FLICA are the matrices $A^{(k)}, k = 1, 2, \ldots, K$, of size $N \times L^{\star\star}$ if we use $Y^{(k)}, k = 1, 2, \ldots, K$, or $L^\star \times L^{\star\star}$ if we use $V^{(k)}, k = 1, 2, \ldots, K$. Note that (unlike the typical approach of feeding spatial PCA eigenvectors into ICA) we are not feeding the spatial dictionary basis ($D^{(k)}$) into the FLICA core modelling, but the feature loadings ($A^{(k)}$). To get the spatial loading matrices from FLICA, we do voxel-wise multiple regression where the target variable is a voxel and the design matrix is the FLICA subject mode. We could change the order by applying DicL first and then mMIGP, but this empirically has a lower computation efficiency.

2.5. Evaluation of BigFLICA in simulations

We simulated 500 subjects, and each had two modalities, which were both $30 \times 30 \times 30$ images. We first simulated K ground-truth (independent) spatial maps X; each of these was a $30 \times 30 \times 30$ image. The spatial maps were a weighted sum of two Gaussian white noise images, where the first one was $30 \times 30 \times 30$ with weight 0.05, and the second was a $5 \times 5 \times 5$ cube randomly located in the full image with weight 0.95. Then, random positive component weights W, Gaussian random subject loadings H and Gaussian white noise terms E were simulated. Finally, after vectorizing each spatial map and noise term, the data for a single modality Y was generated as $Y = HWX + \sigma E$, where σ was a parameter to control the signal-to-noise ratio (SNR). A small amount of spatial smoothing using a Gaussian kernel was applied to spatial maps X and noise terms E to mimic real image data. Each of the two modalities also had 5 unique spatial maps that were not shared by each other. The voxels were z-score normalized before feeding into the subsequent FLICA analysis. The SNR was defined as: $var(vec(HWX))/var(vec(\sigma E))$.

Performance evaluation: When FLICA was applied to the simulated data, the number of independent components was always set to the ground truth K. The performance was measured by the similarity between estimated subject-mode matrix $H^\star$ and the ground truth H. The similarity was measured by the greedy matching of the components based on maximum correlation and then estimating the mean correlations across components.

Evaluation of mMIGP for subject-space dimension reduction: After generating simulated data, we reduced the data to varying dimensions ($L^\star = 50, 100, 200, 300, 400$) using mMIGP, and then fed the reduced data into FLICA. This was compared with the original FLICA. The number of ground-truth components was set to 25,35,45 and the SNR was set to 4,1,0.25,0.04. All simulations were repeated 50 times.

Evaluation of DicL for voxel-space dimension reduction: To evaluate the influence of the DicL parameters on the subsequent FLICA results, we performed the DicL on simulated data using varying parameter combinations ($\lambda = 0.1$ to 16 and $L^{\star\star} = 100$ to 3000) followed by FLICA (nIC = 25, 50, 100). This was compared with the original FLICA. The SNR was set to 4,1,0.25,0.04, and the number of

iterations for the DicL was set to 50, because we empirically find that this number of iterations is sufficient for DicL to converge to a stable result in simulation and real data. All simulations were repeated 50 times.

2.6. HCP and UK Biobank data

The voxel/vertex-wise neuroimaging data of 81 different modalities of 1003 subjects from the HCP S1200 data release were used in this paper (Van Essen et al., 2013). The preprocessing was conducted by the HCP team using an optimized pipeline (Glasser et al., 2013). The 81 modalities included (1) 25 resting-state ICA dual-regression spatial maps (z-score normalized); (2) 47 unique task contrast maps as z-statistics from 7 different fMRI tasks; (3) 3 T1-image derived modalities (grey matter volume, surface area, surface thickness); (4) 6 Tract-Based Spatial Statistics (TBSS) features from diffusion MRI (FA, L1, L2, L3, MD, MO) (Smith et al., 2006). In addition, 158 nIDPs were used here, which was the same as our previous study (Smith et al., 2015). Names of nIDPs are in **Supplementary File 1**.

The UK Biobank imaging data were mainly preprocessed by FSL (Smith et al., 2004; Jenkinson et al., 2012) and FreeSurfer (Fischl, 2012) following an optimized pipeline (Alfaro-Almagro et al., 2018) (https://www.fmrib.ox.ac.uk/ukbiobank/). The voxel-wise neuroimaging data of 47 modalities of 14,053 subjects were used in this paper, including: (1) 25 "modalities" from the resting-state fMRI ICA dual-regression spatial maps (z-score normalized); (2) 6 modalities from the emotion task fMRI: 3 contrasts (shapes, faces, faces>shapes) of z-statistics and 3 *contrasts of parameter estimate* maps; (3) 10 diffusion MRI derived modalities (9 TBSS features, including FA, MD, MO, L1, L2, L3, OD, ICVF, ISOVF (Smith et al., 2006; Zhang et al., 2012) and a summed tractography map of 27 tracts from AutoPtx in FSL); (4) 4 T1-MRI derived modalities (grey matter volume and Jacobian map (which shows expansion/contraction generated by the nonlinear warp to standard space, and hence reflects local volume) in the volumetric space, and cortical area and thickness in the Freesurfer's fsaverage surface space); (5) 1 susceptibility-weighted MRI map (T2-star); (6) 1 T2-FLAIR MRI derived modality (white matter hyperintensity map estimated by BIANCA Griffanti et al., 2016). A detailed description is in **Table A.6**. In addition, the 8787 nIDPs were included, but we retained the 7245 of those that have at least 1000 non-missing values (subjects). Names of nIDPs are in **Supplementary Files**. Group-level resting-state independent component spatial maps and task activation z-statistic maps are in the **Supplementary Files**.

When carrying out nIDP prediction, a total of 13 and 54 confounding variables were regressed out from nIDPs using linear regression in the HCP and the UKB datasets respectively (**Supplementary Materials**). Subjects with a missing modality were imputed by the mean value of all other subjects. We did not impute the missing nIDPs.

2.7. Comparing BigFLICA with the original FLICA on real data

On real data, we do not know the ground truth components, and the data may not follow the assumptions of ICA. Therefore, we rely on the performance of predicting nIDPs as a surrogate criterion to evaluate different methods. We applied the proposed mMIGP approach to HCP data and a subset of 1036 UKB subjects (so that the original FLICA is computationally tractable). Elastic-net regression, from the glmnet package (Zou and Hastie, 2005), was used to predict the nIDPs using FLICA's subject modes as model regressors (features). This approach is widely-used and has been shown to achieve a robust and state-of-the-art performance in many neuroimaging studies (Cui and Gong, 2018; Jollans et al., 2019). To evaluate the model performance, for each nIDP, we used 5-fold cross validation, and computed Pearson correlation between the predicted and true values of each nIDP across the 5 test sets. As there are tuning parameters within the Elastic-net regression, in each training set, we performed a nested 5-fold cross validation to tune the model parameters, and used the best model selected in the nested 5-fold cross validation to do the prediction in the test set. When comparing any two approaches, the same training-validation-testing split was used. The prediction accuracy was quantified as the Pearson correlation between predicted and the true values of each nIDP in the test sets.

To evaluate MIGP preprocessing, we reduced the dimension to varying L^* (from 100 to 500) using MIGP first and then used FLICA to extract $L = 50$ components. The original FLICA was also applied to extract 50 components. To evaluate DicL preprocessing, we used the DicL (dictionary dimension $= 2000$ and sparsity parameter $\lambda = 1$) to reduce the data dimension of each modality followed by the FLICA to extract varying numbers of components (nIC$= 25, 50, 100, 200, 300$). The original FLICA was also applied to extract the same numbers of components. The prediction accuracy of BigFLICA was compared with the original FLICA applied on non-reduced data.

2.8. Statistical significance of difference of prediction accuracy between two approaches

To compare the overall prediction accuracy of two approaches (e.g., BigFLICA with mMIGP preprocessing vs. the original FLICA), we estimate the statistical significance of the difference between the prediction correlations across nIDPs. Starting with a total of p nIDPs, we first exclude nIDPs where both methods have low prediction accuracy ($r < 0.1$, as these would likely just add noise to the comparison), resulting in p_1 nIDPs. Then we want to test whether the overall prediction accuracy of p_1 nIDPs with one method is significantly higher than another method. A naive approach would be a simple paired t-test, but the correlation structures among nIDPs makes the samples dependent with each other, so that the p-value of a t-test is not valid. Note that the paired t-test can also be formulated as a linear regression model, and in the framework of linear regression, sample correlation can be taken into account by the generalized least squares approach. Specifically, we assume: $y = x\beta + e$, where y is the difference of the prediction accuracy between two methods and x is a column of ones. The t statistic for coefficient β can be calculated as (Kariya and Kurata, 2004):

$$\beta = (x'V^{-1}x)^{-1}x'V^{-1}y$$
$$s = (x'V^{-1}x)^{-1}y'(V^{-1} - V^{-1}x(x'V^{-1}x)^{-1}x'V^{-1})y/(p_1 - 1)$$
$$t = \beta/\sqrt{\text{diag}(s)}$$

where V is the sample covariance matrix, i.e., the covariance among nIDPs. Note that when $V = I$, the model is equivalent to a paired t-test. In general we do not know V, but can obtain a good estimate of V by calculating the covariance of nIDPs using the nIDPs-by-subject matrix. We used the lscov function in Matlab to perform the above estimation.

2.9. Parameter settings of running BigFLICA in the full HCP and UKB datasets

We applied BigFLICA approach to extract a varying number of target components in two datasets. In HCP, we used FLICA with DicL preprocessing only (dictionary dimension 2000 and $\lambda = 1$). In UKB, we used FLICA with both mMIGP and DicL preprocessing (dictionary dimension 5000, $\lambda = 1$ and mMIGP dimension 1000 ($>95\%$ explained variance)). The number of FLICA VB iterations is 1000.

2.10. Comparing BigFLICA with multiple independent single-modality ICA decomposition

ICA is a widely-used approach for decomposing single-modality neuroimaging data, including functional MRI (Smith et al., 2015) but also in structural MRI (Zeighami et al., 2015) and diffusion MRI (Li et al., 2012). A natural question arises whether BigFLICA is able to combine multimodal information more effectively than the single-modality approaches such as ICA (we used the fastICA algorithm (Hyvarinen, 1999)), which ignores inter-modality relationships.

We first performed ICA on each modality of HCP and UKB data separately to extract 25,100, 250, 500 and 750 components. For a given component number, we built a prediction model using the concatenated ICA subject modes (across modalities) to predict each of the nIDPs. To be fair, for BigFLICA, we extract the same number of ICs to build the prediction model. For example, in the UKB data and a 25-dimensional decomposition, the predictor is a Subject $\times (25 \times 47)$ matrix for single-modality ICA, where 25 is the number of components in each modality and 47 is the total number of modalities. For BigFLICA, the predictor is a Subject $\times 25$ matrix. This is arguably a fair comparison because each of the BigFLICA modes potentially contains information from all modalities. The method to build a predictive model and evaluate this is the same as above, except that when we used the concatenated ICA subject modes, we added a univariate screening step in the training set to select the top 300 most informative features according to their correlation with an nIDP in the training set. This step, in general, boosts the predictive accuracy because the dimensionality of concatenated ICA modes is usually very high, so that many of the modes are pure noise with respect to any given nIDP. Therefore, the univariate screening can help the elastic-net regression to filter out noisy features effectively. We did not perform univariate screening when using the BigFLICA subject modes to predict nIDPs.

2.11. Comparing BigFLICA with hand-curated imaging-derived phenotypes

A popular choice of data analysis strategy is to extract imaging features based on expert knowledge (e.g., regional volumes and thickness, and resting-state functional connectivities between brain regions), often referred to as IDPs (Miller et al., 2016). Brain IDPs have been shown to genetically correlate with many SNPs in our previous genome-wide association study (GWAS) in UK Biobank (Elliott et al., 2018), and they have been shown to change in many psychiatric diseases (Kelly et al., 2017; Van Rooij et al., 2017; Hibar et al., 2018).

We extracted 5812 IDPs from the HCP, including (1) 199 structural MRI features from Freesurfer as provided by the HCP; (2) 4700 regional mean task activations from 47 independent task contrasts using a 100-dimensional parcellation atlas (Schaefer et al., 2017); (3) 625 functional connectivities (FCs) based on a 25-dimensional ICA parcellation with partial correlation to estimate FCs; (4) 288 regional mean TBSS features (FA, L1, L2, L3, MD, MO) using the Johns Hopkins University tract atlas. The names of these IDPs are given in the **Supplementary File 3**.

We used 3913 IDPs from UKB, including global and local features from the 6 imaging modalities (T1, T2-FLAIR, swMRI, tfMRI, rfMRI, and dMRI) (Smith et al., 2020). The names of these IDPs are given in the **Supplementary File 4**.

We built prediction models that use IDPs or BigFLICA modes to predict each of the nIDPs using the same strategy as above. The FLICA dimension is set to 25, 100, 250, 500, 750. In addition, we also concatenated IDPs and each of the BigFLICA subject modes together to predict the nIDPs, and the performance is compared with using IDPs alone. We used a univariate screening step to select the top 300/500 most informative IDPs according to their correlation with an nIDP in the inner-fold (i.e., training set) of HCP/UKB. Finally, we also built models that use IDPs to predict each of the FLICA subject modes and vice versa, aiming to evaluate the shared variances between features extracted by these two different approaches in the same data.

3. Results

3.1. Evaluation of BigFLICA in simulations

We first applied BigFLICA on simulated data to evaluate the performance of mMIGP and DicL as data preprocessing approaches under different parameter settings and data signal-to-noise ratios. The mean correlation of extracted components with simulated ground truth was compared with the corresponding result from the original FLICA (Methods Section 2.5).

For mMIGP, Fig. 2a shows that, in most of the situations, the BigFLICA with mMIGP preprocessing gave similar results to the original FLICA, and both FLICA and BigFLICA accurately find the underlying ground truth in most cases. This is in agreement with results of simulations in the MIGP paper (Smith et al., 2014) that it can accurately approximate a full-data PCA in different situations. The optimal dimension of mMIGP is different among simulations; sometimes a relative low dimension can achieve an accurate estimation of components (e.g. Fig. 2a first three columns), while in other cases a high dimension is needed (e.g. Fig. 2a the fourth column).

For DicL, Fig. 2b shows that in almost all circumstances: (1) increasing the dictionary dimensions will boost the performance of subsequent FLICA analysis; (2) the optimal sparsity parameters are usually between $\lambda = 0.5$ to 2, and they have similar performance; (3) In most cases the optimal performance given by DicL matches that of non-reduced analysis (noted in figure legends). Therefore, in the real data analysis, when using the DicL approach, we always use a very high dimensional DicL decomposition and fix the sparsity parameter to $\lambda = 1$.

3.2. Computation time comparison

Table 1 shows the comparison of the computation time and memory requirement of BigFLICA with the original FLICA in the UKB dataset. All code was implemented in Python 2.7, and both BigFLICA and FLICA were run using 24 cores on a single compute node with Intel Xeon CPU E7-8857 v2 @ 3.00 GHz CPU and 2048 GB RAM. The computation time includes: (1) Preprocessing of data using mMIGP and DicL (BigFLICA only); (2) Initialization of FLICA parameters; (3) FLICA VB parameter updates. For the 100,000-subjects data, BigFLICA greatly decreases the computation time and memory usage from an unrealistic amount to a modest configuration for a modern HPC cluster, which therefore allows for the possibility of data-driven population phenotype discovery.

3.3. Real data: comparing BigFLICA with the original FLICA based on the prediction accuracy of nIDPs

As there is no ground truth available, we tested modes of BigFLICA have a similar prediction accuracy of nIDPs compared with the original FLICA, using data from the HCP, and a subset of 1036 subjects from the UKB (Methods Section 2.7). Elastic-net regression with nested 5-fold cross-validation was used to predict each of the nIDPs. This approach is widely-used and has been shown to achieve a robust and state-of-the-art performance in many neuroimaging studies (Cui and Gong, 2018; Jollans et al., 2019). Pearson correlation between each of the predicted and the true nIDPs in the outer test fold is used to quantify accuracy.
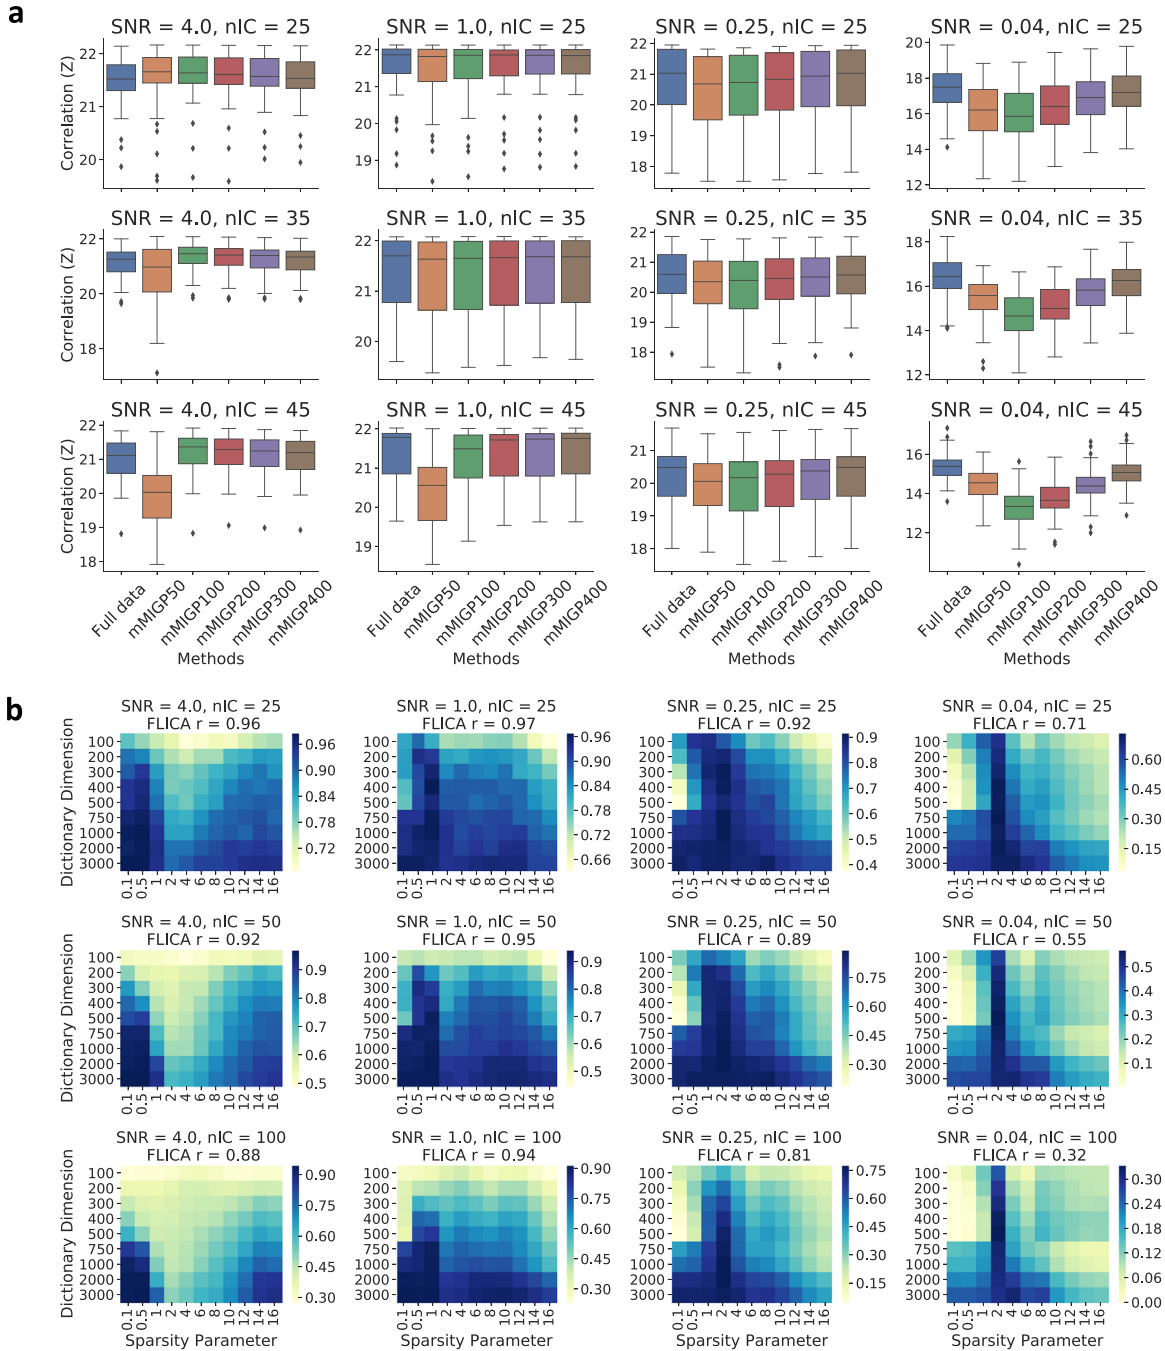
Fig. 2. Evaluation of multimodal extension of MIGP (mMIGP) and dictionary learning (DicL) as the data preprocessing steps for the FLICA using simulations. BigFLICA achieves similar performance as compared with original FLICA that uses the full data. a, Evaluation of mMIGP preprocessing. We compared the correlations (Z-transformed) of extracted components with ground truth across 50 simulations using the original FLICA (the left column of each figure) and the mMIGP preprocessed FLICA (other columns). The mMIGP dimensions vary between 50 and 400; the SNRs are between 4 and 0.04 (left to right), and the number of FLICA and ground truth components are 25, 35, 45 (top to bottom). As there are 500 subjects, the reduction factor is from 10 to 1.25. b, Evaluation of DicL preprocessing. We compared the correlations of extracted components with ground truth using the original FLICA (FLICA results given in the titles of each figure) and the DicL preprocessed FLICA with different sparsity parameters and dictionary dimensions (cells of the heatmaps). The SNRs are between 4 and 0.04 (left to right), and the number of FLICA and ground truth components are 25, 50, 100 (top to bottom). As there are 27,000 original features per modality, the reduction factor is from 270 to 9.

The statistical significance of differences of prediction accuracy between two approaches are estimated by a weighted paired *t*-test approach (Methods Section 2.8).

Fig. 3 shows the Bland–Altman plots comparing the prediction accuracy of nIDPs between original FLICA and BigFLICA with mMIGP preprocessing only (Fig. 3a), and with DicL preprocessing only (Fig. 3b), and with both data reduction approaches (Fig. 3c), in

the UKB and HCP datasets. In these comparisons, mMIGP reduced the data to approximately 1/10 to 1/2 of the original data size, and DicL reduced data to approximately 1/75 of the original data size. Overall, BigFLICA can estimate similar sets of modes with comparable prediction accuracy in real multimodal neuroimaging data, i.e., the difference of the correlation between two methods is centered around zero across a wide range of mean correlation values (which

Table 1

Comparison of computation time and amount of RAM usage of BigFLICA with the original FLICA in the UKB dataset (14,503 subjects, 47 different modalities). BigFLICA greatly increases computational efficiency in different settings. Both BigFLICA and FLICA were run on the same computer using all 24 cores in all computation stages with Intel Xeon CPU E7-8857 v2 @ 3.00 GHz and 2 TB RAM.

| Approaches		Number of components					100 K subjects
		nIC = 25	nIC = 100	nIC = 250	nIC = 500	nIC = 750	750 components (estimated)
Computation time (h)	The original FLICA	160 h	300 h	580 h	1020 h	1680 h	12,000 h
	BigFLICA (mMIGP preprocessing)	23 h	54 h	135 h	315 h	565 h	630 h
	BigFLICA (mMIGP + DicL preprocessing)	52 h	53 h	58 h	65 h	73 h	120 h
Peak RAM (GB)	The original FLICA	801 GB	821	879	963	1066	6000
	BigFLICA (mMIGP preprocessing)	66	88	136	215	297	297
	BigFLICA (mMIGP + DicL preprocessing)	50	50	50	50	50	50
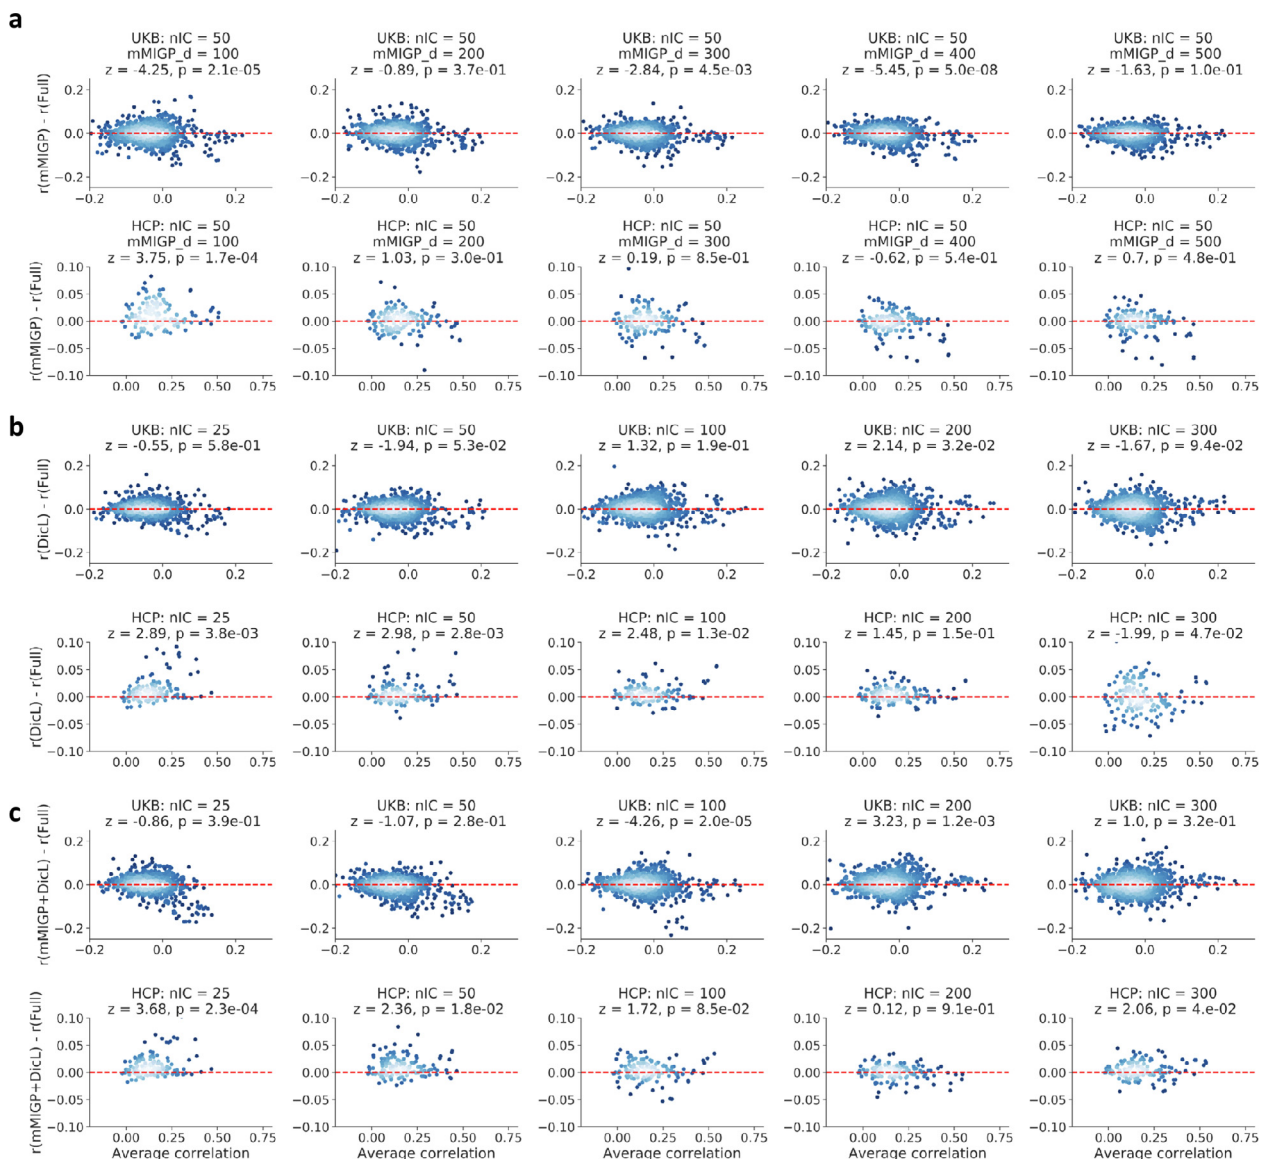
Fig. 3. Comparison of prediction accuracy of nIDPs between BigFLICA and the original FLICA. Overall, for most of the comparisons, the differences of prediction accuracy are not significant. In each of the BlandAltman plots, each point represents the prediction of one nIDP, where the x-axis is the average prediction correlation of the two approaches while the y-axis is the difference, i.e., BigFLICA - FLICA. The z- and p-values in the titles reflected the statistical significance of the differences. The Bonferroni correction 0.05 threshold corresponds to a raw p-value of 1.7e3. a, Comparing FLICA with mMIGP preprocessing with the original FLICA. We used a subset of 1036 subjects in the UKB dataset (top) and the HCP (bottom). The number of estimated FLICA components is set to 50, and mMIGP dimensions are set from 100 to 500. b, Comparing FLICA with DicL preprocessing with the original FLICA. We used a subset of 1036 subjects in the UKB dataset (top) and the HCP (bottom). The dictionary dimension is set to a high value of 2000, and the sparsity parameter is set to $\lambda = 1$ for all modalities. The number of estimated FLICA components are set from 25 to 300. c, Comparing FLICA with both mMIGP and DicL preprocessing combined, with the original FLICA. The mMIGP dimension is set to 500, and other settings are the same as in b. We use only a subset of UKB here so that running the original FLICA is computationally feasible. The lighter the blue, the higher the density of points.

are also reflected in the insignificant *p*-values of weighted paired *t*-test), which demonstrates that the mMIGP and DicL approaches are effective to reduce data and preserve key information in the data.

Finally, we tested whether the set of nIDPs which were predicted better with BigFLICA than the original FLICA are relatively consistent across different nIC. We calculated the differences of the prediction accuracy between BigFLICA and original FLICA, and then tested if they are correlated across different nICs (**Fig. A.6**). However, we did not observe a significant correlation. This might mean that the nIDP-related information in extracted imaging features is different at different scales (ICA dimensionalities), but as the space spanned by the imaging features is increased with increasing dimensionality, it seems more likely that the reduction in prediction of some nIDPs with increasing ICA dimensionality is a result of overfitting.

3.4. Comparing BigFLICA with multiple independent single-modality ICA decomposition

We also compared BigFLICA outputs against features pooled across those from separate ICA processing of each modality (Methods Section 2.10). Fig. 4a shows that BigFLICA has a worse prediction performance than via running ICA separately on each modality when the dimensionality *L* is low. This is because at low dimensional decomposition, single-modality ICA is most efficient because the constraints imposed on the degrees-of-freedom implied in the FLICA model is insufficient to capture the important data variation into joint components. However, when *L* becomes large, the prediction accuracy becomes better than the single-modality ICA (e.g., $\geqslant 250$ in UKB). This is because, at high dimensional decomposition, BigFLICA effectively combines multimodal information by considering cross-modal correlation in the data decomposition stage. Although the cross-modal correlation is considered in the final prediction stage when using single-modality ICA, the fact that BigFLICA identifies and takes advantage of correlated information between modalities at an earlier stage in feature generation helps improve the prediction performance.

In **Fig. A.7**, we also compared, in the UKB data, the 750-dimensional BigFLICA decomposition with the 25-dimensional ICA decomposition concatenated across modalities, i.e., we have 25×47 features in the single-modality ICA. In this comparison, the number of features for the two methods are almost the same, but we can see that BigFLICA clearly outperforms the single-modality ICA.

3.5. Comparing BigFLICA with hand-curated imaging-derive phenotypes

We compared the predictive performance of BigFLICA with IDPs in both HCP and UKB datasets (Methods Section 2.11). Fig. 4b shows that, in the UKB data, when the number of modes is low, BigFLICA has a worse predictive power than the joint performance of 3913 IDPs, due to the same insufficient degree-of-freedom reason as above. However, when the dimensionality becomes higher, BigFLICA is clearly outperforming the IDPs, owing to jointly fusing multimodal voxelwise data by considering cross-modality correlation. In the HCP data, the performance is overall similar. These results indicate that BigFLICA can potentially explain more phenotypic and behavioural variances than IDPs.

In more detail, **Table A.2** shows that, in the UKB dataset, the high-dimensional BigFLICA (nIC = 750) has improved prediction accuracy for many nIDPs that relate to *cognition phenotypes* and *health outcomes* compared with IDPs. These tables do not include nIDPs where both methods have low predictive power ($r < 0.1$). In the HCP dataset (**Table A.3**), BigFLICA (nIC = 100) also shows im-

proved prediction accuracy in many cognitive and health outcomes variables compared with using IDPs.

Further, when we concatenated the modes of BigFLICA and IDPs together to predict nIDPs, as shown in Fig. 4c, the combined feature sets have a significant improvement of prediction accuracy than the IDPs alone in the UKB data. There are almost no differences for the same comparison in the HCP data. This suggests that BigFLICA and IDPs may contain some complementary information of nIDPs.

To investigate the relationships between BigFLICA and IDPs further, we built prediction models that used modes of BigFLICA to predict each of the IDPs, to further characterise information overlap and complementarity between the two approaches. As shown in **Fig. A.8a** and **b**, different types of IDPs can be predicted differently, and the resting-state functional connectivities always had the worst accuracy in both the HCP and the UKB datasets, because they are (relatively) noisy. However, when using BigFLICA modes to predict 6 new summary features of the connectivity matrices (derived by applying ICA to the matrix of subjects by network matrix edges) (Elliott et al., 2018), the accuracy is very high (r range from 0.85 to 0.89 for a 100 dimensional BigFLICA decomposition). In addition, when we used IDPs to predict modes of BigFLICA, as shown in **Fig. A.8c** and **d**, the prediction correlation almost showed a bimodal distribution, which means that some of the FLICA modes can be predicted by the IDPs (mean $r \approx 0.8$) while others cannot (mean $r \approx 0.2$). These results further demonstrate that BigFLICA and IDPs span significant complementary variance.

3.6. Examples of BigFLICA modes in the 14k UKB dataset

We now give four examples of significant associations between BigFLICA modes and nIDPs, namely, *Fluid intelligence, Age started wearing glasses or contact lenses, Handedness* and *hypertension*. In Fig. 5, we show the top four most strongly associated modalities in FLICA modes that correlate with a given nIDP. **Fig. A.16** shows the population cross-subject mean maps for several task and rest fMRI modalities fed into FLICA. This helps give interpretive context for the FLICA mode maps, which depict subject variability in the activity/connectivity relative to these group mean maps.

For *Fluid intelligence*, using all modes (ICs) from the 750 dimensional BigFLICA decomposition as features (predictors) in multivariate elastic-net prediction, a cross-validated prediction correlation of $r = 0.26$ is achieved. When we correlated each of the BigFLICA modes and IDPs with the fluid intelligence score in the UKB, we found that several task-fMRI-related BigFLICA modes have the strongest associations (Fig. 5a). The first (IC 25) involves task contrast "faces" and "faces > shapes" and the second (IC 57) involves contrast "shapes" and "face" (see **Table A.6** for the full list of these modalities). As the correlation of the mode IC 25 (i.e., its subject weights vector) with fluid intelligence is negative ($r = -0.14$), this means that the negative-weights voxels (such as in the anterior insula) are positively correlated with intelligence. The fMRI task (Hariri faces-shapes matching Hariri et al., 2002) has, as expected, the greatest population average activation in sensory-motor areas (plus some amygdala involvement due to the emotionally negative nature of the faces), as seen in **Fig. A.16**. However, the main brain areas involved in these modes are distinct, including anterior cingulate cortex, frontal pole, inferior frontal gyrus, and anterior insula; it is therefore interesting that the areas found by BigFLICA to be modulated in these components (and found to associate with intelligence) are more "frontal, cognitive" areas than the sensory-motor areas primarily activated on average. The top associations between fluid intelligence and IDPs also involve task-fMRI IDPs (**Table A.4**), but these were a factor of two weaker than associations with BigFLICA modes.
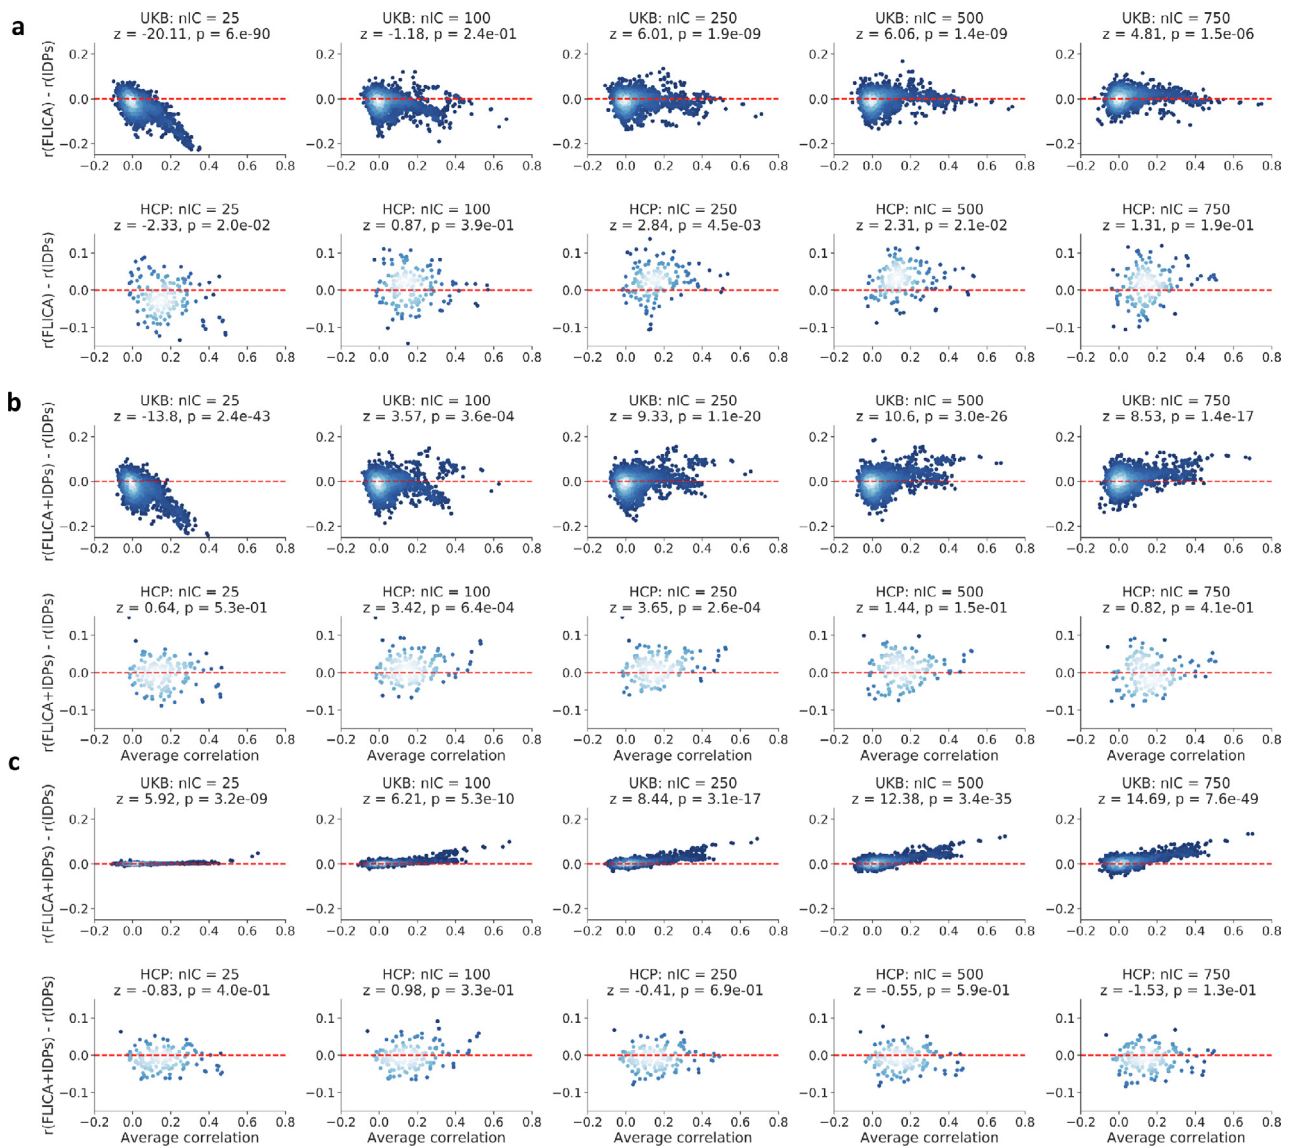
Fig. 4. Comparison of prediction accuracy of nIDPs between BigFLICA against single-modality ICA and the IDPs. Overall, for high-dimensional BigFLICA decompositions in the UKB dataset, BigFLICA achieved statistical significant increases of prediction accuracy of nIDPs compared with single-modality ICA and IDPs. Combining BigFLICA and IDPs together future improves compared with IDPs alone. In each of the Bland–Altman plots, each point represents the prediction of an nIDP, where the *x*-axis is the average prediction correlation of the two approaches, while the *y*-axis is the difference. The *z*- and *p*-values in the titles reflected the statistical significance of the differences. The Bonferroni correction 0.05 threshold corresponds to a raw *p*-value of 1.7e3. a, Comparing BigFLICA with the concatenation of single-modality ICA outputs. Top: UKB; Bottom: HCP. The number of FLICA components is set from 25 to 750. b, Comparing BigFLICA with IDPs. Top: UKB; Bottom: HCP. The number of IDPs is 3913 in UKB and 5812 in the HCP. c, Comparing the concatenation of BigFLICA and IDPs against IDPs only. Top: UKB; Bottom: HCP. The lighter the blue, the higher the density of points.

For *Age started wearing glasses or contact lenses*, BigFLICA achieved a prediction correlation of $r = 0.16$. Several resting-state connectivity and task modalities showed associations in primary visual areas (Fig. 5b), which is consistent with the fact that this is a vision-related health variable. Lower age of first wearing glasses is correlated with stronger activity in primary visual areas, and also with strength of resting-fMRI connectivity (or functional coherence) within the relevant areas of group-average connectivity; interestingly, in nearby distinct (but still primary visual) areas, there is reduction of correlation (blue voxels), suggesting greater differentiation of primary visual areas.

For *Handedness*, BigFLICA achieved a prediction correlation of $r = 0.23$. BigFLICA identified several multimodal, lateralized (or laterally asymmetric) modes, including resting-state mode 14 (left-lateralized language network), task, surface area and white matter tracts (Fig. 5c). There are several resting-state connectivity-related IDPs correlated with handedness (**Table A.4**), consistent with a re-

cent study (Wiberg et al., 2019) that also used UKB IDPs, while no IDPs related to other modalities are found significant; in both cases the maximum IDP correlation only reached $r = 0.12$, whereas the strongest association with BigFLICA modes was almost double this.

For a health variable *hypertension* (Fig. 5d), BigFLICA achieved a prediction correlation of $r = 0.22$. Several TBSS-related modalities showed consistent associations in the External Capsule tracts. Meanwhile, white matter hyperintensity (T2-Lesion volume) in the corresponding areas is also higher in people with hypertension. Several consistent findings have been reported in the literature (Moon et al., 2005; Allen et al., 2016; Hannawi et al., 2018).

3.7. BigFLICA comparison with mCCA and reproducibility

We tested whether BigFLICA (independent components-based spatial modelling) was better than mCCA (eigendecomposition based modelling, which could be considered to be similar to the output of BigFLICA without running the final core FLICA unmixing
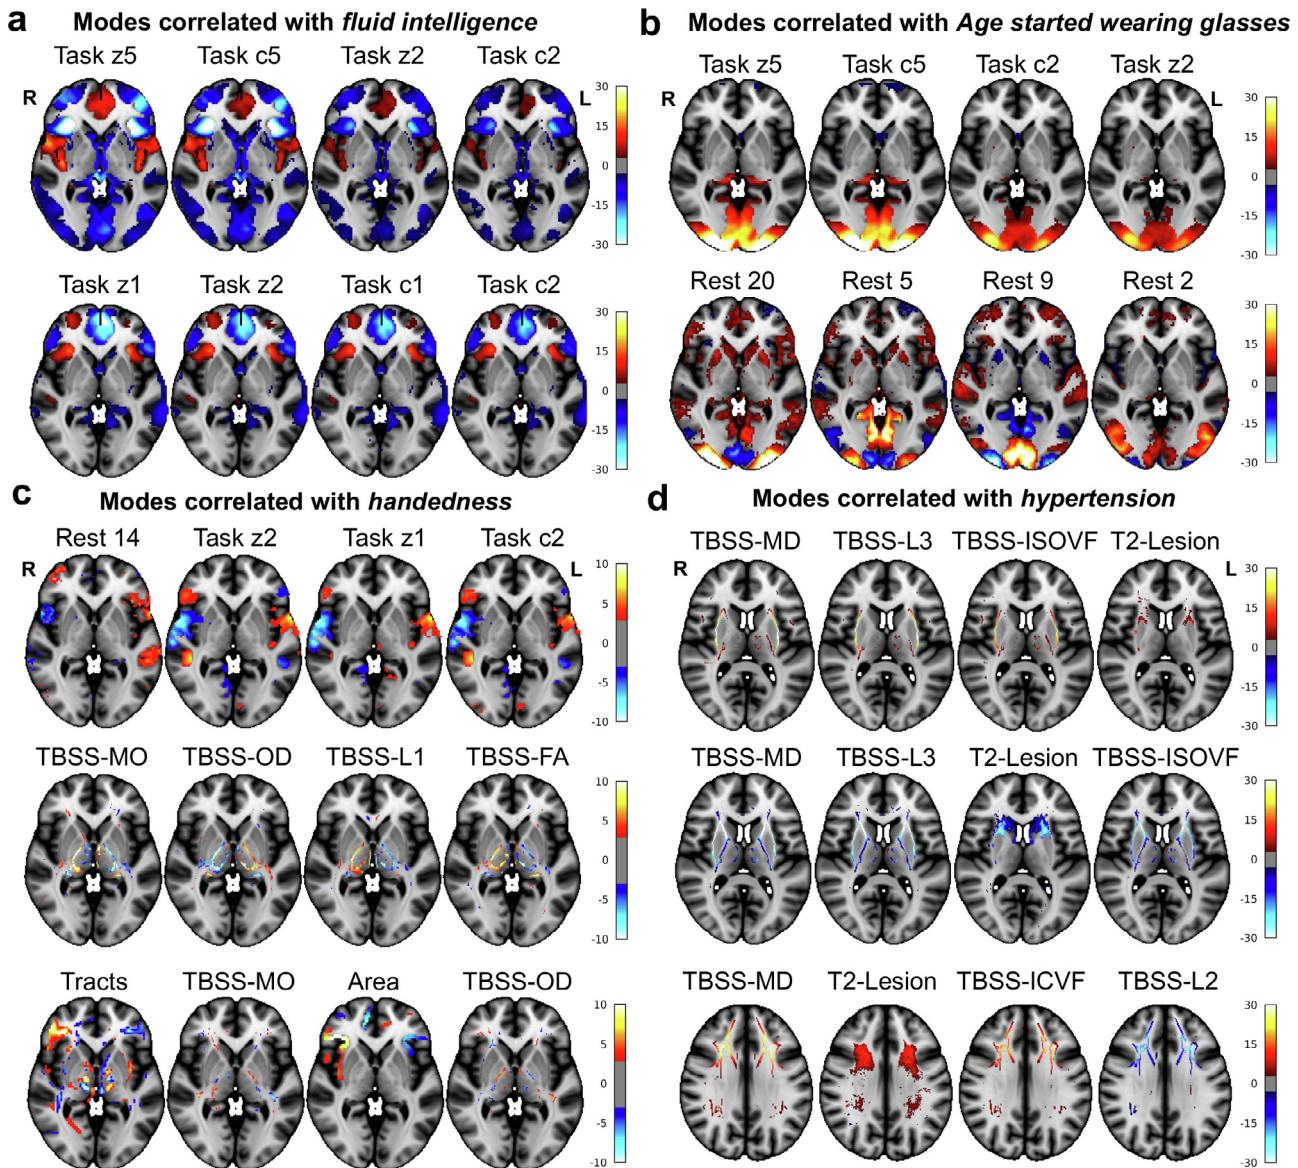
Fig. 5. Examples of BigFLICA modes in the 14k UKB dataset. For each subfigure, each row shows one IC (BigFLICA mode or independent component) with top 4 most strongly associated modalities. a, Two BigFLICA modes that significantly correlate with *fluid intelligence* (IC25: $r = -0.14$; IC57: $r = -0.12$). b, Two BigFLICA modes that significantly correlate with *Age started wearing glasses or contact lenses* (IC164: $r = -0.10$; IC13: $r = -0.05$). c, Three BigFLICA modes that significantly correlate with *handedness* (IC235: $r = -0.23$; IC569: $r = 0.07$; IC232: $r = -0.04$). d, Three BigFLICA modes that significantly correlate with *hypertension* (IC259: $r = 0.12$; IC13: $r = -0.11$; IC319: $r = -0.09$). The Bonferroni corrected 0.05 threshold corresponds to an uncorrected p-value of 9.2×10^{-9} (corrected for number of components (750) and number of nIDPs (7245)). All of the above correlations passed the Bonferroni threshold except for IC232 with uncorrected $p = 2.1 \times 10^{-7}$.

- note that to enable mCCA to run requires the same mMIGP initial processing that we have added in this work) in three ways. The number of extracted components was the same when performing this comparison. First, for the prediction accuracy of nIDPs, **Fig. A.9** shows that, in the UKB data, BigFLICA has a (very slightly) improved prediction accuracy compared with mCCA. Then, we proposed a hypothesis that modes of BigFLICA are more parsimonious features of nIDPs compared with mCCA, or in other word, a smaller number of modes of BigFLICA can predict the nIDPs. Results shown in **Fig. A.10** validate this hypothesis: for a given number of components and a given nIDP, BigFLICA modes have a (on average) higher proportion of zero weights in the elastic-net predictions, when compared with mCCA modes. The advantage is that a more parsimonious representation usually has a better biological interpretability. Finally, we estimated and compared the split-half reproducibility of BigFLICA and mCCA. As shown in **Fig. A.11** (right), BigFLICA has a much higher between-subject reproducibility than

mCCA (median BigFLICA correlation greater than 0.9 in all cases, while many mCCA dimensionalities have median correlation less than 0.5).

3.8. Reproducibility of BigFLICA

To test whether BigFLICA's spatial independent components are estimated reliably, the whole UKB dataset was divided into two parts: the first part contained 7000 subjects and the second part contained the remaining 7503 subjects. We applied BigFLICA to the two parts separately. After estimating the subject modes, we reconstructed the z-score normalized (voxel-wise) spatial maps of each modality by regressing the subject mode against the mMIGP-reduced data. The spatial independent components of each modality were concatenated spatially and greedily paired, based on the absolute correlation between two runs. When we computed the correlations, only voxels whose absolute z-scores that are both

larger than 3 in two runs were preserved (to reduce noise, given that there are huge numbers of empty voxels across all modalities for a given FLICA component in general; this does not bias the metric of reproducibility towards finding common similar patterns). **Fig. A.11** (left) shows that the FLICA components have very high reproducibility in the split-half test across a varying number of components.

We further tested the reproducibility of prediction accuracy of BigFLICA for predicting nIDPs. We ran BigFLICA separately on two halves of the dataset as above, and then predicted each of the nIDPs using elastic net regression in each half. **Fig. A.12** shows the prediction accuracy of all nIDPs. We can see that the prediction accuracy of nIDPs are highly correlated, especially for nIDPs that have a high prediction accuracy ($r = 0.57$ for a 25 dimensional decomposition, and $r > 0.7$ for higher dimensions).

Finally, we investigated the influence of sample size on the reproducibility. We performed the same split-half reproducibility test on two random subsets of 1500 and 3500 subjects. **Fig. A.13** shows that across different numbers of components, BigFLICA has a reproducibility of from 0.6 to 0.7, which is expected to be lower than the 7000 subject case. We can also see that there is a slight increase of reproducibility when the number of components increases. This different behaviour compared with the 7000-subject cases may be due to the fact that many of the components are empty when the number of components is larger than 250 (i.e., there is not enough data to support the ICA dimensionality), so that we can only compute the reproducibility index based on the non-empty components.

3.9. Stability of BigFLICA prediction

We evaluated the stability of prediction accuracy of BigFLICA against different train-test subject splits. We estimated the prediction correlation of each nIDP using 5 different train-test subject splits, and calculated the mean accuracy of prediction correlations for each nIDPs. We then computed the difference of the mean prediction correlation and the prediction correlation of one of the five random predictions. As shown in **Fig. A.14**, the differences are centered around zero with extremely small spread, which demonstrated the stability of BigFLICA against different random train-test splits.

Among the 25 resting-state dual-regression spatial maps included in the analysis, four of them had been originally identified as non-neural components (Miller et al., 2016). Non-neural components likely reflect non-neuronal physiology and therefore may help prediction, particularly for the nIDPs that relate to basic physiology (e.g., blood pressure). Therefore, we tested the prediction accuracy of nIDPs when we exclude the four non-neural modalities (rfMRI components). **Fig. A.15** shows that an increased prediction accuracy was observed for some nIDPs in the 25-dimensional decomposition when compared with using all 47 modalities, a similar prediction accuracy was observed for 100- and 250-dimensional decompositions, while a decreased prediction accuracy was observed for 500- and 750-dimensional decompositions. The set of nIDPs that are better predicted by inclusion of 4 non-neural components at high dimensions are those that relate to physical measures. Researchers may choose to include artifactual rfMRI components (e.g., where this helps maximise nIDP associations), or may exclude them (e.g., to maximise interpretability of associations).

3.10. Contribution of different modalities in a BigFLICA decomposition

Besides using BigFLICA for exploring the relationships between imaging and non-imaging phenotypic and behavioural data, we can also use it to investigate the relationship between different modalities. For each mode, BigFLICA estimates a vector of positive num-

bers reflecting the contributions of different modalities (i.e., the diagonal of each $W_{(k)}$, where the higher the number, the more important is one modality to a mode). We concatenated all such vectors across all modes so that it is a mode-by-modality matrix W, and normalized each column to sum to one. Six examples of such matrices are shown in **Fig. A.17**, with different numbers of estimated modes in the UKB dataset.

We then calculate each row's sum (across columns) in W, thereby reflecting the overall contribution of each modality in the BigFLICA decomposition. As shown in **Fig. A.18**, across all FLICA dimensionalities (numbers of estimated modes), each of the 25 resting-state fMRI dual-regression spatial maps usually has a low overall contribution, followed by task fMRI maps, while modalities reflecting more about structure of the brain (e.g., structural MRI and diffusion MRI) generally have high overall contributions. The relative differences of modality contribution between functional MRI-related modalities and structural/diffusion MRI-related modalities become larger with increasing number of estimated modes. We further estimated the total shared variances between a lower dimensional BigFLICA decomposition and a higher dimensional decomposition. **Table A.5** shows that a higher dimensional decomposition explains almost all variances of a lower dimensional decomposition (upper triangle of the table), while a lower dimensional decomposition can explain a large proportion of the variances of a higher dimensional decomposition.

3.11. Relationship between different modalities in a BigFLICA decomposition

We calculated the cosine similarity between different columns of W (using the 750-dimensional BigFLICA decomposition), to measure the similarity of different modalities in terms of their contribution to the BigFLICA decomposition, i.e., the more similar information two modalities carry, the more likely they will have similar contribution to a mode. **Fig. A.19a** shows that the modality relationship matrix is clearly grouped into three large clusters. The first is all resting-state modalities, while the second is the task fMRI maps, and the third is the diffusion MRI, structural MRI-related modalities and swMRI. The white matter hyperintensity map (T2 lesions) forms a single cluster. As a comparison, we also performed a 50-dimensional ICA decomposition within each modality, and calculated the shared variances between every pair of 50 ICs in two modalities using a simple multivariate regression model. As shown in **Fig. A.19b**, we also observed a similar pattern as **Fig. A.19a**. The main difference is that in **Fig. A.19a**, there are relatively stronger correlations within resting-state modalities and between resting-state and other modalities, but weaker correlations between task modalities and structural related modalities. These results reflect the fact that the multimodal modelling effects of BigFLICA learn different inter-modality relationships compared with single-modality ICA.

4. Discussion

In this paper, we presented BigFLICA, a multimodal data fusion approach which is scalable and tuneable to analyze the full UK-Biobank neuroimaging dataset, and other large-scale multimodal imaging studies. To the best of our knowledge, this is the first approach for data-driven (unsupervised) multimodal analysis in a brain imaging dataset of this size and complexity. Building on the top of the powerful FLICA model, we proposed a two-stage dimension-reduction approach that combines an incremental group-PCA (mMIGP) and dictionary learning (DicL) to effectively preprocess the multimodal dataset and reduce the computational load of the final FLICA, while maintaining or even improving performance, with as much as a 150-fold "intelligent" re-

duction in data size. We provide effective ways of choosing the hyper-parameters of BigFLICA, so that it is free of tuning except for choosing the final number of estimated components. Although this approach is motivated by the need for analyzing extremely big neuroimaging data, it is also applicable to other kinds of data such as genetics and behavioural measures. An easy-to-use version of this software will be integrated into an upcoming version of the FSL software package (Smith et al., 2004; Jenkinson et al., 2012). BigFLICA results on UKB will also be released via the UKB database as new data-driven IDPs (image features), for further epidemiological and neuroscientific research.

A strength of our work is that, unlike previous work that was limited to more moderate datasets and a few phenotypic and behavioural variables (Calhoun et al., 2006; Liu et al., 2009; Beckmann and Smith, 2005; Groves et al., 2011; Sui et al., 2012), we used two of the largest, high-quality multimodal neuroimaging datasets, and thousands of phenotypic and behavioural variables to validate the proposed approach. We demonstrated that BigFLICA is not only much faster than the original FLICA (and can be run on very large data that is simply not analysable with FLICA or other existing methods), but also estimates similar modes with a comparable performance for predicting the non-imaging-derived phenotypes in real data (when tested on a large data subset that is just small enough to allow for comparison against FLICA). We provide insights into the advantages of data-driven multimodal fusion in big datasets by quantitative analysis (Calhoun and Sui, 2016; Uludağ and Roebroeck, 2014). First, when comparing BigFLICA with simpler IDP-based approaches (and also single-modality ICA approaches), we demonstrated that a high-dimensional BigFLICA has improved predictive power overall. We demonstrated the value of multimodal fusion instead of analyzing each modality separately. Second, when combining high-dimensional BigFLICA-derived features with IDPs together, the predictive power increased further compared with using either method alone. In addition, when we used BigFLICA-derived features and manually created (with expert knowledge) IDPs to predict each other, they cannot predict each other perfectly (although they are derived from the same imaging data). This indicates that BigFLICA-derived features and IDPs can be complementary to each other, both therefore providing potentially important imaging biomarkers that capture different signal in the imaging data. An interesting finding is that although a high-dimensional BigFLICA has a much higher predictive power than a low dimensional decomposition, a low dimensional decomposition can still explain more than 80% of the total variance of the high dimensional decomposition. This suggests that some the phenotypic and behavioural variables are explained by only small proportions of variance of imaging data. Third, in addition to the value of using BigFLICA-derived features for relating imaging to non-imaging data, BigFLICA components (particularly at lower dimensionalities) may allow us to learn more about how the different brain imaging modalities (and hence different spatial and biological aspects of the brain's structure and function) relate to each other. Finally, when new primary data becomes available from new subjects, this data would need to have new IDPs calculated (at the subject level) and then combined with existing IDPs from previous subjects, for a complete between-subject analysis. The approach presented here, while not avoiding any of the necessary new computations, will make these efficient. Further, note that it is alternatively possible to use an existing decomposition and apply this to new subjects by projecting them onto the existing spatial bases to generate new subject weights (loadings) against the ICA features.

We see opportunities to improve the current approach. First, BigFLICA is limited to linear feature estimation, while the "ideal, true" information in imaging data may be highly nonlinear. Therefore, a nonlinear extension of BigFLICA, which might be achieved with kernel methods or deep neural networks, is an important

area of further research. Second, BigFLICA is an unsupervised dimension reduction and feature generation approach. However, integrating some supervision, i.e., the target variable (such as disease outcomes), into the dimension reduction may boost the performance of the algorithm. Additionally, because BigFLICA generates data-driven features, as opposed to expert-created IDPs, the biological or anatomical interpretation of features is often likely not to be immediately obvious, requiring potentially intensive expert study. Future work could attempt to automate this interpretation process, for example by relating features to existing anatomical templates and atlases, and even by mining imaging literature. Finally, BigFLICA, or extensions, may be an effective way of discovering imaging confound factors (Li et al., 2020) that cannot be found by traditional approaches.

Data availability

BigFLICA-derived features will be available from the UK Biobank database. For UK Biobank, all source data (including raw and processed brain imaging data, derived IDPs, and non-imaging measures) is available from UK Biobank via their standard data access procedure (see http://www.ukbiobank.ac.uk/register-apply). For HCP, data can be downloaded via website (http://humanconnectome.org/data) and ConnectomeDB.

Code availability

BigFLICA code is available at https://github.com/weikanggong/BigFLICA, and will also be released as part of an upcoming version of FSL. Matlab software for performing prediction using elastic-net regression is available at https://github.com/vidaurre/NetsPredict.

Declaration of Competing Interest

Authors declare that they have no conflict of interest.

CRediT authorship contribution statement

Weikang Gong: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Writing - original draft, Writing - review & editing. **Christian F. Beckmann:** Conceptualization, Methodology, Writing - original draft, Writing - review & editing, Supervision, Project administration, Funding acquisition. **Stephen M. Smith:** Conceptualization, Methodology, Writing - original draft, Writing - review & editing, Supervision, Project administration, Funding acquisition.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at 10.1016/j.media.2021.102050

References

Alfaro-Almagro, F., Jenkinson, M., Bangerter, N.K., Andersson, J.L., Griffanti, L., Douaud, G., Sotiropoulos, S.N., Jbabdi, S., Hernandez-Fernandez, M., Vallee, E., et al., 2018. Image processing and quality control for the first 10,000 brain imaging datasets from UK Biobank. Neuroimage 166, 400–424.

Allen, B., Muldoon, M.F., Gianaros, P.J., Jennings, J.R., 2016. Higher blood pressure partially links greater adiposity to reduced brain white matter integrity. Am. J. Hypertens. 29 (9), 1029–1037.

Allen, G.I., Grosenick, L., Taylor, J., 2014. A generalized least-square matrix decomposition. J. Am. Stat. Assoc. 109 (505), 145–159.

Alnæs, D., Kaufmann, T., Doan, N.T., Córdova-Palomera, A., Wang, Y., Bettella, F., Moberget, T., Andreassen, O.A., Westlye, L.T., 2018. Association of heritable cognitive ability and psychopathology with white matter properties in children and adolescents. JAMA Psychiatry 75 (3), 287–295.

Ball, G., Aljabar, P., Nongena, P., Kennea, N., Gonzalez-Cinca, N., Falconer, S., Chew, A.T., Harper, N., Wurie, J., Rutherford, M.A., et al., 2017. Multimodal image analysis of clinical influences on preterm brain development. Ann. Neurol. 82 (2), 233–246.

Ball, G., Malpas, C.B., Genc, S., Efron, D., Sciberras, E., Anderson, V., Nicholson, J.M., Silk, T.J., 2018. Multimodal structural neuroimaging markers of brain development and ADHD symptoms. Am. J. Psychiatry 176 (1), 57–66.

Beckmann, C.F., Smith, S.M., 2005. Tensorial extensions of independent component analysis for multisubject fMRI analysis. Neuroimage 25 (1), 294–311.

Calhoun, V.D., Adali, T., Giuliani, N., Pekar, J., Kiehl, K., Pearlson, G., 2006. Method for multimodal analysis of independent source differences in schizophrenia: combining gray matter structural and auditory oddball functional data. Hum. Brain Mapp. 27 (1), 47–62.

Calhoun, V.D., Sui, J., 2016. Multimodal fusion of brain imaging data: a key to finding the missing link (s) in complex mental illness. Biol. Psychiatry Cogn. Neurosci. Neuroimaging 1 (3), 230–244.

Cui, Z., Gong, G., 2018. The effect of machine learning regression algorithms and sample size on individualized behavioral prediction with functional connectivity features. Neuroimage 178, 622–637.

Douaud, G., Groves, A.R., Tamnes, C.K., Westlye, L.T., Duff, E.P., Engvig, A., Walhovd, K.B., James, A., Gass, A., Monsch, A.U., et al., 2014. A common brain network links development, aging, and vulnerability to disease. Proc. Natl. Acad. Sci. 111 (49), 17648–17653.

Elliott, L.T., Sharp, K., Alfaro-Almagro, F., Shi, S., Miller, K.L., Douaud, G., Marchini, J., Smith, S.M., 2018. Genome-wide association studies of brain imaging phenotypes in UK Biobank. Nature 562 (7726), 210.

Fischl, B., 2012. Freesurfer. Neuroimage 62 (2), 774–781.

Glasser, M.F., Sotiropoulos, S.N., Wilson, J.A., Coalson, T.S., Fischl, B., Andersson, J.L., Xu, J., Jbabdi, S., Webster, M., Polimeni, J.R., et al., 2013. The minimal preprocessing pipelines for the human connectome project. Neuroimage 80, 105–124.

Griffanti, L., Salimi-Khorshidi, G., Beckmann, C.F., Auerbach, E.J., Douaud, G., Sexton, C.E., Zsoldos, E., Ebmeier, K.P., Filippini, N., Mackay, C.E., et al., 2014. ICA-based artefact removal and accelerated fMRI acquisition for improved resting state network imaging. Neuroimage 95, 232–247.

Griffanti, L., Zamboni, G., Khan, A., Li, L., Bonifacio, G., Sundaresan, V., Schulz, U.G., Kuker, W., Battaglini, M., Rothwell, P.M., et al., 2016. BIANCA (brain intensity abnormality classification algorithm): a new tool for automated segmentation of white matter hyperintensities. NeuroImage 141, 191–205.

Groves, A.R., Beckmann, C.F., Smith, S.M., Woolrich, M.W., 2011. Linked independent component analysis for multimodal data fusion. Neuroimage 54 (3), 2198–2217.

Groves, A.R., Smith, S.M., Fjell, A.M., Tamnes, C.K., Walhovd, K.B., Douaud, G., Woolrich, M.W., Westlye, L.T., 2012. Benefits of multi-modal fusion analysis on a large-scale dataset: life-span patterns of inter-subject variability in cortical morphometry and white matter microstructure. Neuroimage 63 (1), 365–380.

Hannawi, Y., Yanek, L., Kral, B.G., Vaidya, D., Becker, L., Becker, D., Nyquist, P.A., 2018. Hypertension is associated with white matter disruption in apparently healthy middle-aged individuals. Am. J. Neuroradiol. 39 (12), 2243–2248.

Hariri, A.R., Tessitore, A., Mattay, V.S., Fera, F., Weinberger, D.R., 2002. The amygdala response to emotional stimuli: a comparison of faces and scenes. Neuroimage 17 (1), 317–323.

Hibar, D., Westlye, L.T., Doan, N.T., Jahanshad, N., Cheung, J., Ching, C.R., Versace, A., Bilderbeck, A., Uhlmann, A., Mwangi, B., et al., 2018. Cortical abnormalities in bipolar disorder: an MRI analysis of 6503 individuals from the ENIGMA bipolar disorder working group. Mol. Psychiatry 23 (4), 932.

Hoyos-Idrobo, A., Varoquaux, G., Kahn, J., Thirion, B., 2019. Recursive nearest agglomeration (ReNA): fast clustering for approximation of structured signals. IEEE Trans. Pattern Anal. Mach. Intell. 41 (3), 669–681.

Hyvarinen, A., 1999. Fast and robust fixed-point algorithms for independent component analysis. IEEE Trans. Neural Netw. 10 (3), 626–634.

Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M., 2012. FSL. Neuroimage 62 (2), 782–790.

Jernigan, T. L., Brown, S. A., Dowling, G. J., 2018. The adolescent brain cognitive development study. Journal of research on adolescence: the official journal of the Society for Research on Adolescence 28 (1), 154.

Jollans, L., Boyle, R., Artiges, E., Banaschewski, T., Desrivières, S., Grigis, A., Martinot, J.-L., Paus, T., Smolka, M.N., Walter, H., et al., 2019. Quantifying performance of machine learning methods for neuroimaging data. NeuroImage 199, 351–365.

Kariya, T., Kurata, H., 2004. Generalized Least Squares. John Wiley & Sons.

Kelly, S., Jahanshad, N., Zalesky, A., Kochunov, P., Agartz, I., Alloza, C., Andreassen, O., Arango, C., Banaj, N., Bouix, S., et al., 2017. Widespread white matter microstructural differences in schizophrenia across 4322 individuals: results from the ENIGMA Schizophrenia DTI working group. Mol. Psychiatry.

Kettenring, J.R., 1971. Canonical analysis of several sets of variables. Biometrika 58 (3), 433–451.

Klami, A., Virtanen, S., Leppäaho, E., Kaski, S., 2015. Group factor analysis. IEEE Trans. Neural Netw. Learn. Syst. 26 (9), 2136–2147.

Kong, R., Li, J., Orban, C., Sabuncu, M.R., Liu, H., Schaefer, A., Sun, N., Zuo, X.-N., Holmes, A.J., Eickhoff, S.B., et al., 2018. Spatial topography of individual-specific cortical networks predicts human cognition, personality, and emotion. Cereb. Cortex 29 (6), 2533–2551.

Li, H., Smith, S.M., Gruber, S., Lukas, S.E., Silveri, M.M., Hill, K.P., Killgore, W.D., Nickerson, L.D., 2020. Denoising scanner effects from multimodal MRI data using linked independent component analysis. NeuroImage 208, 116388.

Li, Y.-O., Yang, F.G., Nguyen, C.T., Cooper, S.R., LaHue, S.C., Venugopal, S., Mukherjee, P., 2012. Independent component analysis of DTI reveals multivariate microstructural correlations of white matter in the human brain. Hum. Brain Mapp. 33 (6), 1431–1451.

Liégeois, R., Li, J., Kong, R., Orban, C., Van De Ville, D., Ge, T., Sabuncu, M.R., Yeo, B.T., 2019. Resting brain dynamics at different timescales capture distinct aspects of human behavior. Nat. Commun. 10 (1), 2317.

Liu, J., Pearlson, G., Windemuth, A., Ruano, G., Perrone-Bizzozero, N.I., Calhoun, V., 2009. Combining fMRI and SNP data to investigate connections between brain function and genetics using parallel ICA. Hum. Brain Mapp. 30 (1), 241–255.

Mairal, J., Bach, F., Ponce, J., Sapiro, G., 2010. Online learning for matrix factorization and sparse coding. J. Mach. Learn. Res. 11 (Jan), 19–60.

Miller, K.L., Alfaro-Almagro, F., Bangerter, N.K., Thomas, D.L., Yacoub, E., Xu, J., Bartsch, A.J., Jbabdi, S., Sotiropoulos, S.N., Andersson, J.L., et al., 2016. Multimodal population brain imaging in the UK Biobank prospective epidemiological study. Nat. Neurosci. 19 (11), 1523–1536.

Moon, W.-J., Na, D.G., Kim, S.S., Ryoo, J.W., Chung, E.C., 2005. Diffusion abnormality of deep gray matter in external capsular hemorrhage. Am. J. Neuroradiol. 26 (2), 229–235.

Parra, L. C., 2018. Multi-set canonical correlation analysis simply explained. arXiv preprint arXiv:1802.03759.

Schaefer, A., Kong, R., Gordon, E.M., Laumann, T.O., Zuo, X.-N., Holmes, A.J., Eickhoff, S.B., Yeo, B.T., 2017. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity MRI. Cereb. Cortex 28 (9), 3095–3114.

Smith, S.M., Elliott, L.T., Alfaro-Almagro, F., McCarthy, P., Nichols, T.E., Douaud, G., Miller, K.L., 2020. Brain aging comprises many modes of structural and functional change with distinct genetic and biophysical associations. BioRxiv 802686.

Smith, S.M., Hyvärinen, A., Varoquaux, G., Miller, K.L., Beckmann, C.F., 2014. Group-PCA for very large fMRI datasets. Neuroimage 101, 738–749.

Smith, S.M., Jenkinson, M., Johansen-Berg, H., Rueckert, D., Nichols, T.E., Mackay, C.E., Watkins, K.E., Ciccarelli, O., Cader, M.Z., Matthews, P.M., et al., 2006. Tract-based spatial statistics: voxelwise analysis of multi-subject diffusion data. Neuroimage 31 (4), 1487–1505.

Smith, S.M., Jenkinson, M., Woolrich, M.W., Beckmann, C.F., Behrens, T.E., Johansen-Berg, H., Bannister, P.R., De Luca, M., Drobnjak, I., Flitney, D.E., et al., 2004. Advances in functional and structural mr image analysis and implementation as FSL. Neuroimage 23, S208–S219.

Smith, S.M., Nichols, T.E., 2018. Statistical challenges in æbig dataþ human neuroimaging. Neuron 97 (2), 263–268.

Smith, S.M., Nichols, T.E., Vidaurre, D., Winkler, A.M., Behrens, T.E., Glasser, M.F., Ugurbil, K., Barch, D.M., Van Essen, D.C., Miller, K.L., 2015. A positive-negative mode of population covariation links brain connectivity, demographics and behavior. Nat. Neurosci. 18 (11), 1565.

Sui, J., Adali, T., Yu, Q., Chen, J., Calhoun, V.D., 2012. A review of multivariate methods for multimodal fusion of brain imaging data. J. Neurosci. Methods 204 (1), 68–81.

Uludağ, K., Roebroeck, A., 2014. General overview on the merits of multimodal neuroimaging data fusion. Neuroimage 102, 3–10.

Van Essen, D.C., Smith, S.M., Barch, D.M., Behrens, T.E., Yacoub, E., Ugurbil, K., Consortium, W.-M.H., et al., 2013. The WU-Minn human connectome project: an overview. Neuroimage 80, 62–79.

Van Rooij, D., Anagnostou, E., Arango, C., Auzias, G., Behrmann, M., Busatto, G.F., Calderoni, S., Daly, E., Deruelle, C., Di Martino, A., et al., 2017. Cortical and subcortical brain morphometry differences between patients with autism spectrum disorder and healthy individuals across the lifespan: results from the ENIGMA ASD working group. Am. J. Psychiatry 175 (4), 359–369.

Wiberg, A., Ng, M., Al Omran, Y., Alfaro Almagro, F., McCarthy, P., Marchini, J., Bennett, D., Smith, S., Douaud, G., Furniss, D., 2019. Handedness, language areas and neuropsychiatric diseases: insights from brain imaging and genetics. Brain 142 (10), 2938–2947.

Zeighami, Y., Ulla, M., Iturria-Medina, Y., Dadar, M., Zhang, Y., Larcher, K.M.-H., Fonov, V., Evans, A.C., Collins, D.L., Dagher, A., 2015. Network structure of brain atrophy in de novo Parkinson's disease. Elife 4, e08440.

Zhang, H., Schneider, T., Wheeler-Kingshott, C.A., Alexander, D.C., 2012. NODDI: practical in vivo neurite orientation dispersion and density imaging of the human brain. Neuroimage 61 (4), 1000–1016.

Zou, H., Hastie, T., 2005. Regularization and variable selection via the elastic net. J. R. Stat. Soc. 67 (2), 301–320.