

RESEARCH ARTICLE OPEN ACCESS

Predicting Subject Traits From Brain Spectral Signatures: An Application to Brain Ageing

Cecilia Jarne^{1,2,3}  | Ben Griffin⁴ | Diego Vidaurre^{3,4,5}

¹Departamento de Ciencia y Tecnología de la Universidad Nacional de Quilmes, Bernal, Buenos Aires, Argentina | ²CONICET, Buenos Aires, Argentina | ³Center of Functionally Integrative Neuroscience, Department of Clinical Medicine, Aarhus University, Aarhus, Denmark | ⁴Wellcome Centre for Integrative Neuroimaging, Oxford University, Oxford, UK | ⁵Oxford Centre for Human Brain Activity, Department of Psychiatry, Oxford University, Oxford, UK

Correspondence: Diego Vidaurre (dvidaurre@cfin.au.dk)

Received: 7 November 2023 | **Revised:** 29 October 2024 | **Accepted:** 26 November 2024

Funding: This work was supported by the European Research Council Starting Grant (ERC-StG-2019-850404), Novo Nordisk Foundation Emerging Investigator Fellowship (NNF19OC-0054895), Agencia Nacional de Promoción de la Investigación, el Desarrollo Tecnológico y la Innovación (PICT 2020-01413), Wellcome Trust (215573/Z/19/Z), Independent Research Fund of Denmark (2034-00054B).

Keywords: brain age | EEG | Kernel mean embedding regression | Kernel methods | machine learning | maximum mean discrepancy

ABSTRACT

The prediction of subject traits using brain data is an important goal in neuroscience, with relevant applications in clinical research, as well as in the study of differential psychology and cognition. While previous prediction work has predominantly been done on neuroimaging data, our focus is on electroencephalography (EEG), a relatively inexpensive, widely available and non-invasive data modality. However, EEG data is complex and needs some form of feature extraction for subsequent prediction. This process is sometimes done manually, risking biases and suboptimal decisions. Here we investigate the use of data-driven Kernel methods for prediction from single channels using the EEG spectrogram, which reflects macro-scale neural oscillations in the brain. Specifically, we introduce the idea of reinterpreting the spectrogram of each channel as a probability distribution, so that we can leverage advanced machine learning techniques that can handle probability distributions with mathematical rigour and without the need for manual feature extraction. We explore how the resulting technique, Kernel mean embedding regression, compares to a standard application of Kernel ridge regression as well as to a non-Kernelised approach. Overall, we found that the Kernel methods exhibit improved performance thanks to their capacity to handle nonlinearities in the relation between the EEG spectrogram and the trait of interest. We leveraged this method to predict biological age in a multinational EEG data set, HarMNqEEG, showing the method's capacity to generalise across experiments and acquisition setups.

1 | Introduction

Predicting behavioural and cognitive traits from brain data in a way that generalises to unseen subjects and is robust to acquisition idiosyncrasies is important because it can offer objective measures to otherwise elusive neurobiological constructs (Haynes and Rees 2006). A widely studied example is brain age. While measuring actual age is a straightforward task, the concept of brain age provides a marker of mental health by quantifying

how much a subject's brain appears to have aged with respect to the population average; that is, a predicted age that is lower than the individual's chronological age, for example, may indicate that the person has a brain that appears younger than expected for their actual age (Franke and Gaser 2019; Smith et al. 2019).

Much of the work on the prediction of subject traits (such as age) from brain data has been done with resting-state fMRI data (Dosenbach et al. 2010). Since we cannot

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2024 The Author(s). *Human Brain Mapping* published by Wiley Periodicals LLC.

straightforwardly predict from the raw data, an intermediate representation is typically used for prediction. In the case of fMRI, this is often a simple description of functional connectivity (Rosenberg et al. 2016) or some model of brain dynamics (Liegeois et al. 2019; Vidaurre et al. 2021; Ahrends, Woolrich, and Vidaurre 2024). Here, we focus on EEG, a considerably less costly technique. As an intermediate representation, we consider the EEG frequency spectrum, reflecting neural oscillations that are well-known correlates of different behavioural and cognitive states (Buzsáki and Draguhn 2004). How to meaningfully extract features for prediction from an EEG spectrum is an open question. Some current efforts are based on manual feature extraction (Al Zoubi et al. 2018; Engemann et al. 2022); or are based on models that need to be estimated from raw data and whose properties depend on the choice of configuration and hyperparameters (Vidaurre, Bielza, and Larrañaga 2013).

In this paper, we present a method to predict from single-channel spectrograms, that is, with no need of the raw data. In contrast to other studies that make predictions on whole brain signals (Dimitriadis and Salis 2017; Vandenbosch et al. 2019; Sabbagh et al. 2019), sometimes with complex methods such as deep learning (Khayretdinova et al. 2022), our approach has not only benefits in terms of computational simplicity but also for interpretation, as it allows us to compare the predictive power across sensors or brain regions. Specifically, with no prior predefinition of frequency bands or any manual feature engineering besides basic preprocessing, we propose the idea of interpreting the EEG spectrogram as a probability distribution. This way, we can fully leverage all the powerful machinery of Kernel learning on probability distributions. We consider Kernel mean embeddings (KMEs), a technique used to construct a representation of the data in a high-dimensional feature space (Smola et al. 2007; Iyer, Jagarlapudi, and Sarawagi 2014; Borgwardt et al. 2006). The KME technique maps joint, marginal and conditional probability distributions to vectors in a high (or even infinite) dimensional feature space that completely characterises the distribution (Fukumizu, Song, and Gretton 2011). Building upon KME, we used the maximum mean discrepancy (MMD) (Smola et al. 2007), a distance metric defined on the space of probability measures (here, EEG spectrograms) in combination with Kernel ridge regression (KRR) (Saunders, Gammernan, and Vovk (1998)). We refer to our approach as Kernel mean embedding regression (KMER), which we compare to KRR and ridge regression (RR).

We demonstrate the superior performance of the Kernel approaches (both KMER and KRR) for age prediction. Using a multisite, public resting-state EEG dataset with a very wide distribution of age (HarMNqEEG; Li et al. 2022), we show that these lead to better predictions, which can be projected on the EEG scalp for interpretation. Focusing on KMER, we show that parietal sensors are the most accurate in predicting age, with slightly greater accuracy in men. We also showed that the predictions generalise well across experiments and acquisition sites, even considering the large differences in age distribution across sites (a well-known problem in machine learning referred to as prior shift). Overall, by demonstrating its predictive capacity

and interpretability, Kernel methods can help unveil insights about brain age or other neurobiological constructs.

2 | Results

2.1 | Kernel Mean Embedding Regression

We developed KMER, a novel method to predict subject traits from EEG spectrograms; see Section 3 and Figure 1. In short, KMER is based on the idea of interpreting EEG spectrograms as probability distributions so that techniques based on Kernel learning of distributions are readily applicable. By forming a subject-by-subject Kernel matrix of similarities between subjects (using a Gaussian Kernel; see Section 3), this technique allowed us to derive predictions for each EEG channel that can naturally accommodate nonlinearities in a data-driven way without the need for manual feature extraction.

2.2 | Comparison of Performance Between Methods

First, compared KMER to: (i) RR, a regularised, linear regression method that uses the frequency bins' power as features; and (ii) KRR, the Kernelised version of RR, which allows the modelling of nonlinearities by the use of an appropriate Kernel. Our experiments were based on the prediction of age in the HarMNqEEG data set (Li et al. 2022), a multibatch, multicountry resting-state EEG dataset of 1966 subjects encompassing people across the entire lifespan (excluding the very youngest and using only subjects between 5 and 97 years old); see Figure S1 and Table S1 for some basic statistics about the dataset in terms of age, gender, and geography. The performance of the methods was assessed using cross-validation, such that the acquisition batches were never split across folds. The hyperparameters (described in Methods) were chosen using nested cross-validation for all approaches. We assessed the performance of the three methods per EEG channel using two measures of accuracy: prediction explained variance (R^2) and mean absolute error (MAE), which offers complementary information (Engemann et al. 2022).

Figure 2 summarises the results per EEG channel. Figure 2A illustrates an example of age vs. predicted age for a given channel (T3). Figure 2B,C shows a comparison of the three methods across channels in terms of MAE and explained variance, respectively. As shown, MAE oscillates between 11 and 12.5 years approximately for all channels. There is a larger variation in terms of explained variance, ranging from $R^2=0.24$ for sensor Fp1 to $R^2=0.44$ for sensor C3. KRR and KMER outperform RR for all channels, and KMER outperforms KRR only for some channels. The advantage of KRR and KMER over RR is highly significant ($p=1\times 10^{-5}$; permutation testing), but the advantage of KMER over KRR is not significant ($p=0.093$). This improvement of the Kernel methods over RR is likely due to the presence of nonlinearities in the relation between EEG spectrograms and age. Results for the linear and polynomial Kernel (referring to the Kernel used to construct the Kernel matrix, $f(\cdot)$; see Section 3) are shown in Figures S2 and S3; these performed worse than the radial basis function (Gaussian) Kernel.

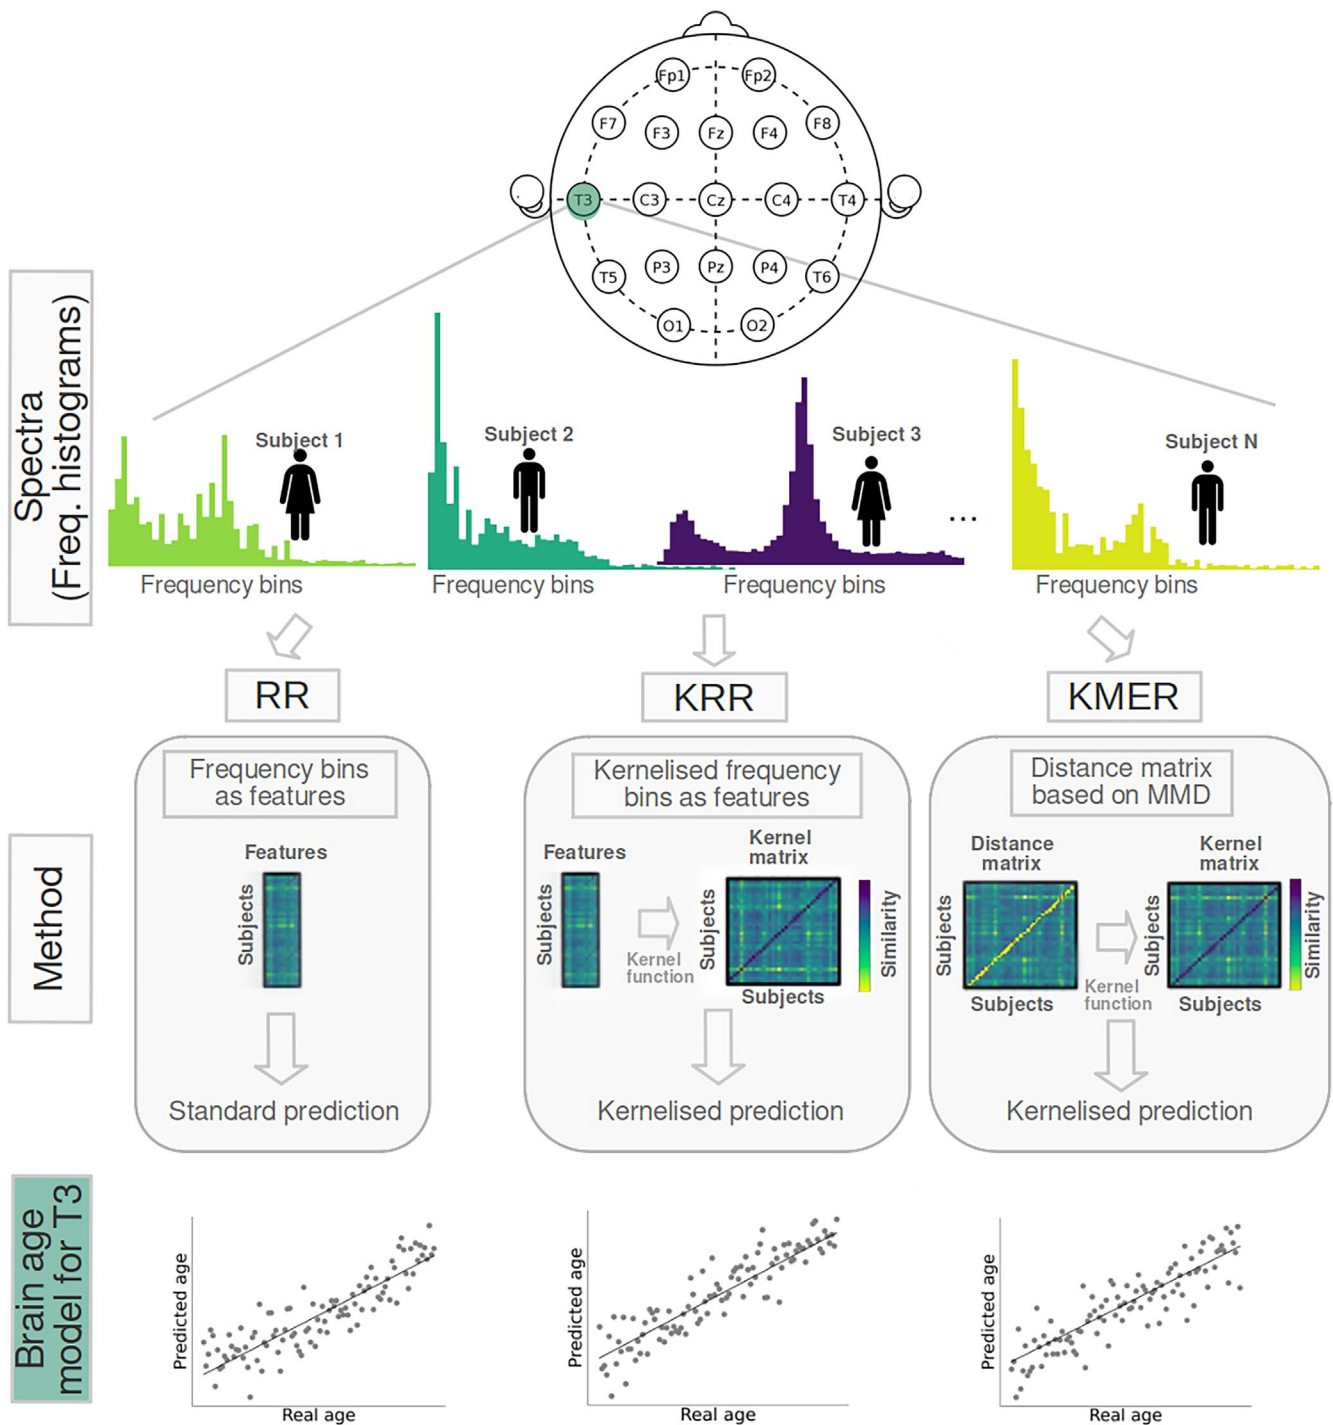


FIGURE 1 | General workflow of the analysis. We used the HarMNqEEG dataset collected across multiple sites (batches), which contains rescaled power densities per sensor and participant, as well as information on gender and age. We considered three prediction approaches, which we run separately per EEG sensor: (i) ridge regression (RR) using the power estimates at each frequency bin of the spectrogram as features; (ii) Kernel ridge regression (KRR), a Kernelised version of RR (based on a nonlinear radial basis function Kernel); and (iii) our Kernel mean embedding regression (KMER) approach, where we interpreted the power spectral estimates across bins as probability distributions so that we can leverage the mathematical machinery of Kernel learning on probability measures for prediction.

An important observation, as seen in Figure 2A, is that age differences are less well-predicted at older ranges. To try to address this problem, we considered the prediction of age in log space (such the differences in the younger are amplified). As observed in Figure S4, we obtained moderate improvements in R^2 and MAE.

A post hoc bias correction, where we removed the age dependence on the residual in a second step using linear regression (Smith et al. 2019), however, resulted in a more substantial improvement; see Figure S5. This correction is performed by training a linear regression model on the residuals of the training data to decrease the bias of the predictions in the testing data,

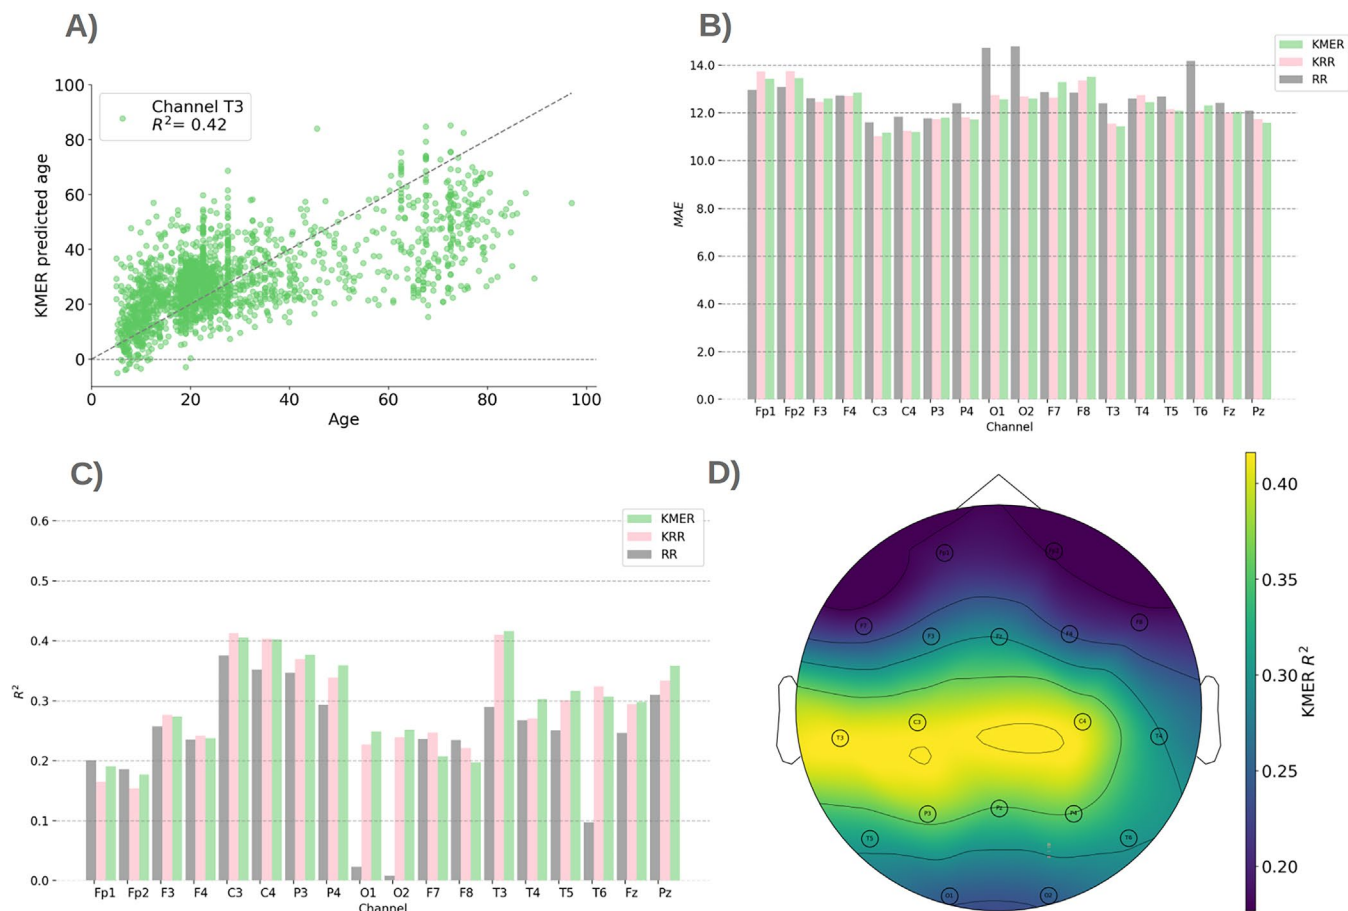


FIGURE 2 | Comparison between the methods. (A) Illustration of predicted versus real age for example channel T3 for KMER. (B) Mean absolute error (MAE) per channel for KMER, RR, and KRR. (C) Explained variance R^2 . (D) R^2 projected on the sensor space for KMER.

typically improving performance across metrics like R^2 and MAE. Specifically, we first calculate the residuals of the training data by subtracting the predictions from the actual training values. Then, we fit a linear regression model, where the chronological age is the independent variable and the residual is the dependent variable. This makes the (new) residual and age orthogonal (Smith et al. 2019). Here, the correction is incorporated into the original KRR/KMER predictions to generate bias-corrected predictions.

We also studied the effect of normalising each frequency bin across subjects, such that the variance is equal for all bins. As shown in Figure S7, the results did not change substantially.

Finally, Figure 2D shows the explained variance in sensor space for KMER (MAE has a similar but inverted topography; not shown). We can observe that the parietal sensors are the most predictive of chronological age, whereas prefrontal and occipital sensors are the least predictive.

2.3 | Sex Differences

Females and males have previously been shown to exhibit differences in their ageing trajectories (Hägg and Jylhävä 2021). Here, using KMER, we investigated the differences between sexes in prediction accuracy across channels for 884 females and 905 males (for 137 individuals, sex is non-specified).

Figure 3 presents the results. Figure 3A,B shows example scatter plots between age and predicted age per sex. Figure 3B,C shows R^2 across sensors in both bar plot and sensor space format. Analogously, Figure 3D,E shows MAE per sex. As observed, males generally exhibited a higher prediction accuracy than females ($p=0.002$), potentially suggesting that their pace of biological change corresponds more closely to chronological age; see Figure S6 for a side-by-side comparison.

2.4 | Performance Variations Across Scanning Sites

The HarMNqEEG dataset encompasses a wide range of age groups across various sites or geographical locations (14 different sites/experiments), each with a different number of individuals and specific age distributions; see Figure S1 and Table S1. Despite efforts made by the data curators to ensure homogeneity, it is possible that differences in practice and instrumentation between the sites leak into the data. In our reported results so far, sites were never split across cross-validation folds; that is, when predicting the age of subjects in a given site (e.g., Colombia), none of the subjects of that site were part of the training data. Given that substantial differences in age distribution between sites may potentially coexist with other between-site differences (e.g., instrumentation), this approach was performed to ensure that the reported prediction accuracies were purely reflective of age

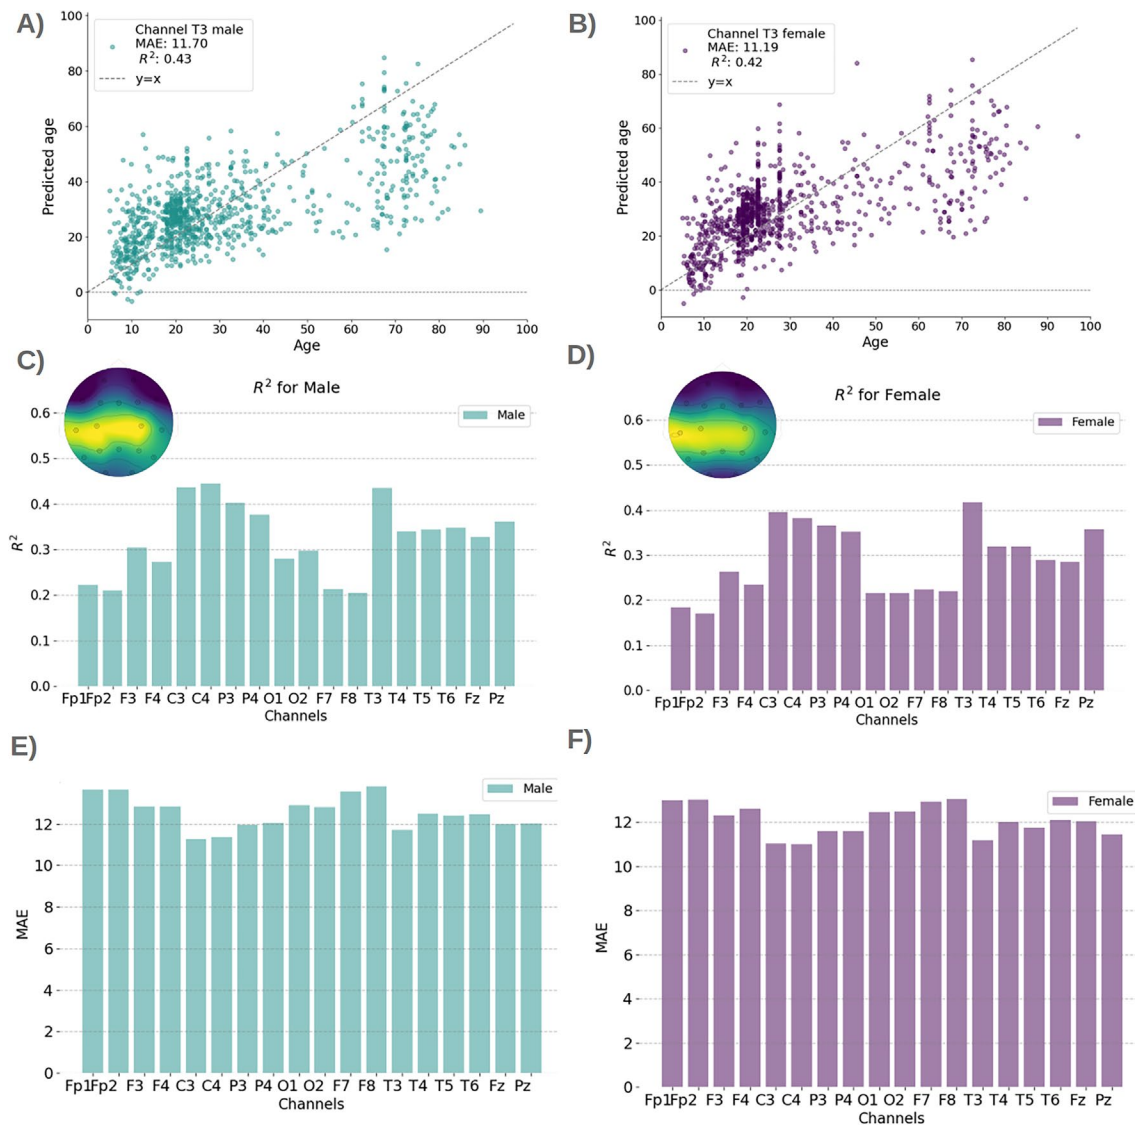


FIGURE 3 | Sex differences (left panels, male; right panels, female). (A, B) Predicted versus real age per sex for example channel T3; results are consistent for all channels (not shown). (C, D) Explained variance R^2 per sex and channel, in bar plot format and a topographic map. (E, F) MAE per sex and channel.

and not mixed with other factors. However, this conservative approach makes the prediction more difficult, inducing what is known in the machine learning literature as a prior shift (Kouw 2018), where the distribution of the dependent variable (here, age) changes between training and test.

To illustrate the issue, Figure 4A shows predicted age versus chronological age with colours indicating stratification by site, with two different perspectives: 3D to better appreciate each site individually and 2D to compare the sites side-by-side. Here, we can observe substantial differences between sites in both accuracy and age distribution. Figure 4B shows distributions of delta across sites, defined as the difference between predicted and chronological age (Franke and Gaser 2019; Smith et al. 2019). When compared with Figure S1, where we show the age distributions explicitly, we can observe that the distribution of errors is contingent on the age range within each group, underestimating the age of older individuals (see Switzerland as an example, which has the oldest population in the data set). The observed pattern is consistent across different channels (not shown).

We compared the cross-site to a within-site prediction, where we ran a cross-validated prediction within one given site without using data from the other sites. While this can only be done effectively for sites with a sufficient number of subjects, we observed substantial increases in accuracy in this type of prediction. Figure 4C,D shows the case of New York, where the explained variance went up to 0.63 and MAE became as small as 5.1.

Overall, these results demonstrate the challenges of predicting across sites when the age distributions are very different, but also that KMER and KRR can still produce reasonable accuracies even in this adverse situation.

2.5 | Comparison With Other Studies

To further compare with existing work on age prediction, Table 1 presents the prediction accuracies from other studies that performed age estimation from EEG data, using other methodological approaches and data sets. Considering that

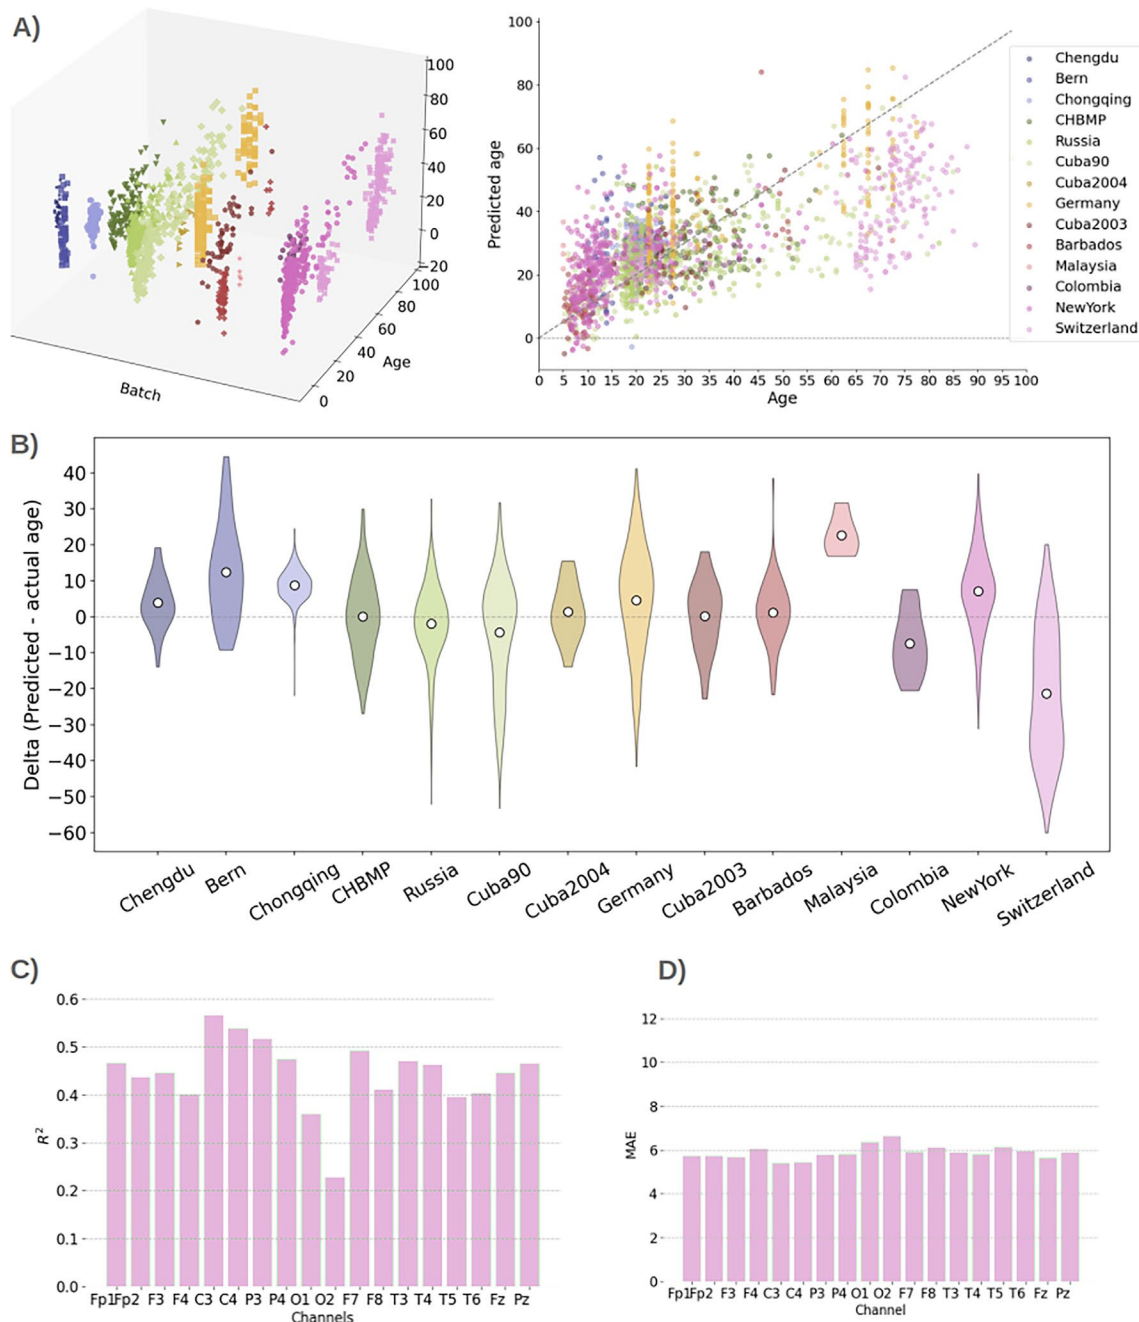


FIGURE 4 | Differences between sites and their effect on the prediction. (A) Predicted versus real age in 3D (left) and 2D (right), separated by site (batch). (B) Distribution of delta (defined as predicted minus real age) per site. (C, D) R^2 and MAE where training and testing come from a single site (New York).

our method operates at the single-channel level, that it only has access to the spectrograms, and that our reported accuracies are primarily across sites (which, as mentioned, is harder), the present accuracies fare well in relation to those reported elsewhere.

3 | Materials and Methods

The general workflow of our analysis is presented in Figure 1. In summary, we considered the spectrogram from every subject as the only input. Predictions of age were made per channel in a

cross-validation fashion, where the acquisition sites (e.g., Russia) were used as folds. We used three approaches to predict age from the EEG spectrograms: RR, KRR, and KMER. For scoring prediction performance, we used explained variance (R^2) and the MAE, which offers complementary information (Engemann et al. 2022). We also computed the so-called delta, defined as the signed difference between predicted and actual age, which is often used to quantify brain age. In this section, we describe the data set and KMER; RR and KRR are standard and their description can be found elsewhere. All code used in this paper is open and publicly available at: https://github.com/katejarne/Kernel_Max_mean_discrepancy_EEG_Age.

TABLE 1 | Other results from the literature of age estimation using EEG. These figures were obtained using different methods, age ranges and datasets. EO and EC, respectively, denote eyes open and eyes closed. The results are compared with ours for the best channel in age space, for NY, in log space and using bias correction (see Figure S5).

Work	Individuals	Range	MAE	R^2/R	Details
Khayretdinova et al. (2022)	1335	[5, 88]	6 ± 0.33 yr	$R^2 = 0.81$	(EO and EC) Neural networks
Vandenbosch et al. (2019)	702	[5, 18]	—	$0.53 < r < 0.74$	Random Forest
Dimitriadis and Salis (2017)	94	[19, 67]	—	$R^2 = 0.60/R^2 = 0.48$	EO/EC SV Regression + Classifier
Al Zoubi et al. (2018)	468	[18, 58]	6.68 ± 0.69	$R^2 = 0.37$	nested-cross-validation (NCV) and stack-ensemble
Sun et al. (2019) I	2535	[18, 80]	7.4yr	$R = 0.83$	(all electrodes)
Sun et al. (2019) II	1974	[40, 80]	Same I	Same I	Same I
Engemann et al. (2022)	2500	[25, 75]	Depends on a data set	R^2 in the range 0.60–0.74 (best)	3 EEG data sets: (LEMON, CHBP, TUAB)
HarMNqEEG data set (present work)	1926	[5, 97]	10.78yr./ 5.24yr. (NY best)	$R^2 = 0.44$ (best in age space) $R^2 = 0.59$ (NY best)	$R^2 = 0.53$ (best in log space) $R^2 = 0.73$ (best with Bias correction)

3.1 | The HarMNqEEG Dataset

We used EEG spectrograms from the HarMNqEEG dataset (Li et al. 2022), which originated from a legacy dataset associated with the Cuban Human Brain Mapping initiative (Valdes-Sosa et al. 2021). The data used in our study were collected from 9 countries, 12 EEG systems, and 14 experiments, which here we refer to as batches or sites. Overall, there are 1966 subjects, of which 40 subjects without a recorded age were excluded. Also, we discarded babies and toddlers, considering only subjects between 5 and 97 years old. For prediction, we used the rescaled raw spectrogram, such that (just like a probability distribution) the values integrate to 1.0. Additional details of the data set are described in Supporting Information, including batch names in the repository (Table S1 and Figure S1).

Recordings were taken from the 19 channels of the 10/20 International Electrodes Positioning System: Fp1, Fp2, F3, F4, C3, C4, P3, P4, O1, O2, F7, F8, T3/T7, T4/T8, T5/P7, T6/P8, Fz, Cz, and Pz, where Cz was the reference electrode in some batches, while in others average-referencing was applied. We therefore discarded Cz from the analysis and performed the predictions for the other 18 channels. Data was formatted as cross-spectral matrices sampled from 1.17 to 19.14 Hz, with a 0.39 Hz resolution. The scalp EEG cross-spectrum was calculated by the data set curators using Bartlett’s method (Møller (1986)), by averaging the periodograms of more than 20 consecutive and nonoverlapping segments. Because the spectra from different sites have different maximum cutoff frequencies, we used the lowest maximum cutoff frequency across sites for full compatibility. Thus, the histogram of each of the 18 channels has 49 bins corresponding to a maximum cutoff frequency of 19.13 Hz.

3.2 | Maximum Mean Discrepancy

This section describes the mathematical foundations of KME, which the next section will use to elaborate on how KME can be applied for prediction using EEG spectrograms.

KME fully characterise the distributions by mapping joint, marginal, and conditional probability distributions to vectors in a high (or infinite) dimensional feature space (Fukumizu, Song, and Gretton 2011; Smola et al. 2007). This implies that any two distributions P and Q with differences in any moment are mapped to separate points in a reproducing Kernel Hilbert space \mathcal{H} (Fukumizu, Song, and Gretton 2011), which is essentially a space of functions. In this space, we can for instance perform classification, regression or clustering of probability distributions, with no loss of information.

In practice, we do not have access to the distributions P or Q , so we use the available samples $\mathbf{X} = \{x_1, \dots, x_n\}$ and $\mathbf{Z} = \{z_1, \dots, z_m\}$. The empirical estimate of the KME μ_P (i.e., the projection of the probability distribution on the reproducing Kernel Hilbert space) is given by

$$\hat{\mu}_P = \frac{1}{n} \sum_{i=1}^n k(x_i, \cdot), \quad (1)$$

where $k(\cdot, \cdot)$ is a Kernel function.

Now, the MMD, which is a distance metric established on the space of probability measures (Gretton et al. (2012)), is defined as

$$\text{MMD}(P, Q) = \left\| \mu_P - \mu_Q \right\|_{\mathcal{H}}. \quad (2)$$

The empirical estimate of the MMD, which we use here, is given by:

$$D(P, Q) = \widehat{\text{MMD}}_k(P, Q) = \sqrt{\frac{1}{n^2} \sum_{i_1, i_2=1}^n k(x_{i_1}, x_{i_2}) + \frac{1}{m^2} \sum_{j_1, j_2=1}^m k(z_{j_1}, z_{j_2}) - \frac{2}{nm} \sum_{i,j=1}^{n,m} k(x_i, z_j)}. \quad (3)$$

where n and m are the number of samples in distributions P and Q , respectively. In other words, this formula is based on similarities between samples and quantifies the similarities within- (first two terms) versus between distributions (third term). In what follows, we considered this estimate of the MMD as our distance metric between distributions (EEG spectrograms). We used a radial basis function Kernel, lineal and polynomial as Kernel functions.

3.3 | Kernel Mean Embedding Regression

In this section, we link the above concepts to prediction using EEG spectrograms. We interpret the EEG spectrograms (rescaled to sum up to 1.0) as probability distributions (denoted above as P or Q). We expect these to vary across age so that we can leverage these differences for prediction. We deterministically generated samples (above, referred to as, e.g., \mathbf{X}) from each rescaled spectrogram (whose bin heights sum up to 1.0) so that we can compute the MMD following Equation (3). For example, for a given channel, if the bin “10Hz” has a height equal to 0.2, and we assume that we have $n = 1000$ samples, then we have $0.2 \times 1000 = 200$ samples equal to 10. This way, we can apply the empirical estimator $\widehat{\text{MMD}}_k$ defined in Equation (3) to each pair of spectrograms in the HarMNqEEG data set, considering different Kernel functions $k(\cdot, \cdot)$. This produced a $(N \times N)$ distance matrix \mathbf{D} per Kernel, where N is the number of subjects. This distance matrix needs then to be converted to a Kernel matrix \mathbf{K} , as required for any Kernelised algorithm. For this, we use a second Kernel function, denoted as $f(\cdot, \cdot)$, that takes a distance between two subjects to produce a quantification of similarity. Specifically, we use a radial basis function Kernel,

$$K(P, Q) = f(D(P, Q)) = \exp(\gamma D(P, Q)^2), \quad (4)$$

where the two subjects in question have been denoted as P and Q for continuity in the notation, and γ is a hyperparameter of the Kernel function $f(\cdot)$. The resulting Kernel matrix \mathbf{K} can then be used to estimate prediction weights as

$$\beta = (\mathbf{K} + \alpha \mathbf{I}) \mathbf{y} \quad (5)$$

where \mathbf{y} is the response variable (age), α is a regularisation hyperparameter (which we select using cross-validation) and \mathbf{I} is the identity matrix. A prediction for an incoming subject T is then made as

$$\hat{\mathbf{y}} = \sum_{i=1}^N \beta_i f(D(P_i, T)), \quad (6)$$

where P_1, \dots, P_N denote the training subjects.

4 | Discussion

In this paper, we have explored the use of Kernel methods for the prediction of individual traits from EEG spectrograms and we have introduced KMER, a prediction method based on the idea of interpreting channel spectrograms as probability distributions. By doing this, we could leverage mathematical principles from Kernel learning. Although not pursued in this paper, the same principles can also be applied to other problems, such as unsupervised clustering of subjects from EEG spectral information. Even though the mathematical foundations of the method are not particularly trivial, KMER is simple to implement and use, provides spatial interpretation at the sensor or source space, and is not computationally costly in comparison to more complex (e.g., deep learning) approaches that are estimated on the raw data. The fact that it only necessitates the spectrograms has additional practical benefits in terms of data sharing. These benefits are shared, nonetheless, by KRR.

We tested the methods in the prediction of age in the HarMNqEEG dataset, a recently published international initiative that made EEG spectrograms available across different cohorts and countries. In these data, we observed that KMER and KRR performed similarly, although they both outperformed non-Kernelised regression (RR), highlighting the benefits of Kernel methods in this context. We note that compared to other works that predicted age within a single cohort, predicting across sites (i.e., such that models were trained on some sites and tested in others) presents a more challenging problem because the distribution of age across sites varies substantially and potentially also because of protocol or infrastructure differences between the sites. Despite these difficulties, the presented results are comparable to previous EEG studies; see Section 2.5. Furthermore, when training and testing within one site only (New York, with a large number of subjects and the broadest age range), the accuracy for some of the individual sensors attained the state-of-the-art accuracies reported in the literature, some of which used more complex models and whole-brain data.

An overarching limitation of all the methods is that, while differences at younger ages are well-predicted, their predictive capacity is much reduced for the older age ranges. A conceivable possibility is that the ageing process is comprised of two separate components of change: a developmental and an ageing part, with more salient differences within development. To investigate this, we performed predictions on the logarithm of the age, which improved results moderately, and we carried out a post hoc bias correction (Smith et al. 2019), which did improve the results substantially. While the cause of this phenomenon may be important, it falls out of the scope of this work and will require further investigation.

Although we did not exhaustively explore it here, KMER could be further optimised by expanding the selection of the Kernel functions and hyperparameters. For instance, for KMER we optimised the Kernel hyperparameter that generates the Kernel matrix from the distance matrix ($f(\cdot)$, Equation 4), the Kernel used in the construction of the MMD metric ($k(\cdot, \cdot)$, Equation 3)

also has a Kernel hyperparameter which we set to 1.0 by default. Improving the hyperparameter tuning thus may be a potential avenue for KMER to outperform the simpler KRR approach. Future work will also explore extensions for multichannel predictions and prediction in source spaces.

Finally, it is worth noting that, although we here demonstrated the performance of the tested methods on EEG, KMER (as well as KRR) can be applied to other modalities such as MEG or ECoG; and can be used to predict other individual traits besides age, including cognitive and clinical variables.

Acknowledgments

D. Vidaurre is supported by a Novo Nordisk Foundation Emerging Investigator Fellowship (NNF19OC-0054895), an ERC Starting Grant (ERC-StG-2019-850404), and a DFF Project 1 from the Independent Research Fund of Denmark (2034-00054B). This research was funded in part by the Wellcome Trust (215573/Z/19/Z). For the purpose of Open Access, the author has applied a CC BY public copyright license to any Author Accepted Manuscript version arising from this submission. We acknowledge support from PICT 2020-01413. We also thank Sonsoles Alonso for her help.

Conflicts of Interest

The authors declare no conflicts of interest.

Data Availability Statement

The data that support the findings of this study are available in <https://www.synapse.org/> with id: syn26712693. Complete public access is available by registering and logging into the system.

References

- Ahrends, C., M. Woolrich, and D. Vidaurre. 2024. "Predicting Individual Traits From Models of Brain Dynamics Accurately and Reliably Using the Fisher Kernel." *eLife* 13: RP95125.
- Al Zoubi, O., C. Ki Wong, R. T. Kuplicki, et al. 2018. "Predicting Age From Brain EEG Signals—A Machine Learning Approach." *Frontiers in Aging Neuroscience* 10: 184.
- Borgwardt, K. M., A. Gretton, M. J. Rasch, H.-P. Kriegel, B. Schölkopf, and A. J. Smola. 2006. "Integrating Structured Biological Data by Kernel Maximum Mean Discrepancy." *Bioinformatics* 22, no. 14: e49–e57.
- Buzsáki, G., and A. Draguhn. 2004. "Neuronal Oscillations in Cortical Networks." *Science* 304, no. 5679: 1926–1929.
- Dimitriadis, S. I., and C. I. Salis. 2017. "Mining Time-Resolved Functional Brain Graphs to an EEG-Based Chronnectomic Brain Aged Index (CBAI)." *Frontiers in Human Neuroscience* 11: 423.
- Dosenbach, N. U. F., B. Nardos, A. L. Cohen, et al. 2010. "Prediction of Individual Brain Maturity Using fMRI." *Science* 329, no. 5997: 1358–1361.
- Engemann, D. A., A. Mellot, R. Höchenberger, et al. 2022. "A Reusable Benchmark of Brain-Age Prediction From M/EEG Resting-State Signals." *NeuroImage* 262: 119521.
- Franke, K., and C. Gaser. 2019. "Ten Years of *BrainAGE* as a Neuroimaging Biomarker of Brain Aging: What Insights Have We Gained?" *Frontiers in Neurology* 10: 789.
- Fukumizu, K., L. Song, and A. Gretton. 2011. "Kernel Bayes' Rule." In *Advances in Neural Information Processing Systems*, edited by J.

Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger, vol. 24. Red Hook, NY, USA: Curran Associates, Inc.

Gretton, A., K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. 2012. "A Kernel Two-Sample Test." *Journal of Machine Learning Research* 13, no. 1: 723–773.

Hägg, S., and J. Jylhävä. 2021. "Sex Differences in Biological Aging With a Focus on Human Studies." *eLife* 10: e63425.

Haynes, J.-D., and G. Rees. 2006. "Decoding Mental States From Brain Activity in Humans." *Nature Reviews Neuroscience* 7, no. 7: 523–534.

Iyer, A. S., S. Jagarlapudi, and S. Sarawagi. 2014. "Maximum Mean Discrepancy for Class Ratio Estimation: Convergence Bounds and Kernel Selection." In *Proceedings of the 31st International Conference on Machine Learning*, 530–538.

Khayretdinova, M., A. Shovkun, V. Degtyarev, A. Kiryasov, P. Pshonkovskaya, and I. Zakharov. 2022. "Predicting Age From Resting-State Scalp EEG Signals With Deep Convolutional Neural Networks on TD-Brain Dataset." *Frontiers in Aging Neuroscience* 14: 1019869.

Kouw, W. M. 2018. "An Introduction to Domain Adaptation and Transfer Learning." *CoRR*. <https://doi.org/10.48550/arXiv.1812.11806>.

Li, M., Y. Wang, C. Lopez-Naranjo, et al. 2022. "Harmonized-Multinational qEEG Norms (HarMNqEEG)." *NeuroImage* 256: 119190.

Liegeois, R., J. Li, R. Kong, et al. 2019. "Resting Brain Dynamics at Different Timescales Capture Distinct Aspects of Human Behavior." *Nature Communications* 10, no. 1: 2317.

Møller, J. 1986. "Bartlett Adjustments for Structured Covariances." *Scandinavian Journal of Statistics* 13: 1–15.

Rosenberg, M. D., E. S. Finn, D. Scheinost, et al. 2016. "A Neuromarker of Sustained Attention From Whole-Brain Functional Connectivity." *Nature Neuroscience* 19, no. 1: 165–171.

Sabbagh, D., P. Ablin, G. Varoquaux, A. Gramfort, and D. A. Engemann. 2019. "Manifold-Regression to Predict From MEG/EEG Brain Signals Without Source Modeling." In *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, vol. 32. Red Hook, NY, USA: Curran Associates, Inc.

Saunders, C., A. Gammerman, and V. Vovk. 1998. "Ridge Regression Learning Algorithm in Dual Variables." In *Proceedings of the Fifteenth International Conference on Machine Learning*, edited by J. Shavlik, 515–521. Morgan Kaufmann.

Smith, S. M., D. Vidaurre, F. Alfaro-Almagro, T. E. Nichols, and K. L. Miller. 2019. "Estimation of Brain Age Delta From Brain Imaging." *NeuroImage* 200: 528–539.

Smola, A., A. Gretton, L. Song, and B. Schölkopf. 2007. "A Hilbert Space Embedding for Distributions." In *Algorithmic Learning Theory*, edited by M. Hutter, R. A. Servedio, and E. Takimoto, 13–31. Berlin, Heidelberg: Springer Berlin Heidelberg.

Sun, H., L. Paixao, J. T. Oliva, et al. 2019. "Brain Age From the Electroencephalogram of Sleep." *Neurobiology of Aging* 74: 112–120.

Valdes-Sosa, P. A., L. Galan-Garcia, J. Bosch-Bayard, et al. 2021. "The Cuban Human Brain Mapping Project, a Young and Middle Age Population-Based EEG, MRI, and Cognition Dataset." *Scientific Data* 8, no. 1: 45.

Vandenbosch, M. M. L. J. Z., D. van't Ent, D. I. Boomsma, A. P. Anokhin, and D. J. A. Smit. 2019. "EEG-Based Age-Prediction Models as Stable and Heritable Indicators of Brain Maturation Level in Children and Adolescents." *Human Brain Mapping* 40, no. 6: 1919–1926.

Vidaurre, D., C. Bielza, and P. Larrañaga. 2013. "Classification of Neural Signals From Sparse Autoregressive Features." *Neurocomputing* 111: 21–26.

Vidaurre, D., A. Llera, S. Smith, and M. Woolrich. 2021. "Behavioural Relevance of Spontaneous, Transient Brain Network Interactions in fMRI." *NeuroImage* 229: 117713.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.