



















# ABCD Neurocognitive Prediction Challenge 2019: Predicting Individual Fluid Intelligence Scores from Structural MRI Using Probabilistic Segmentation and Kernel Ridge Regression

Agoston Mihalik<sup>1,2</sup> , Mikael Brudfors<sup>1,3</sup> , Maria Robu<sup>1,4</sup> ,  
Fabio S. Ferreira<sup>1,2</sup> , Hongxiang Lin<sup>1</sup> , Anita Rau<sup>1,4</sup> , Tong Wu<sup>1,2</sup> ,  
Stefano B. Blumberg<sup>1</sup> , Baris Kanber<sup>5</sup> , Maira Tariq<sup>1</sup> ,  
Mar Estarellas Garcia<sup>1</sup> , Cemre Zor<sup>1,2</sup> , Daniil I. Nikitichev<sup>1,4</sup> ,  
Janaina Mourão-Miranda<sup>1,2</sup> , and Neil P. Oxtoby<sup>1</sup>  

<sup>1</sup> Centre for Medical Image Computing (CMIC), Department of Computer Science  
and Department of Medical Physics and Biomedical Engineering,  
University College London, Gower Street, London WC1E 6BT, UK  
[n.oxtoby@ucl.ac.uk](mailto:n.oxtoby@ucl.ac.uk)

<sup>2</sup> Max Planck UCL Centre for Computational Psychiatry and Ageing Research,  
University College London, Gower Street, London WC1E 6BT, UK

<sup>3</sup> The Wellcome Centre for Human Neuroimaging,  
University College London, Gower Street, London WC1E 6BT, UK

<sup>4</sup> Wellcome/EPSRC Centre for Interventional and Surgical Sciences (WEISS),  
University College London, Gower Street, London WC1E 6BT, UK

<sup>5</sup> Department of Clinical and Experimental Epilepsy,  
Queen Square Institute of Neurology, University College London,  
Gower Street, London WC1E 6BT, UK

**Abstract.** We applied several regression and deep learning methods to predict fluid intelligence scores from T1-weighted MRI scans as part of the ABCD Neurocognitive Prediction Challenge 2019. We used voxel intensities and probabilistic tissue-type labels derived from these as features to train the models. The best predictive performance (lowest mean-squared error) came from kernel ridge regression ( $\lambda = 10$ ), which produced a mean-squared error of 69.7204 on the validation set and 92.1298 on the test set. This placed our group in the fifth position on the validation leader board and first place on the final (test) leader board.

**Keywords:** Kernel ridge regression · Fluid intelligence · PRoNT0 · MRI · Convolutional neural networks · Hackathon

---

A. Mihalik, M. Brudfors, J. Mourão-Miranda and N. P. Oxtoby—These authors contributed equally to this work.

## 1 Introduction

Establishing the neurobiological mechanisms underlying intelligence is a key area of research in neuroscience [1, 2]. General intelligence at a young age is predictive of later educational achievement, occupational attainment, and job performance [3–6]. Moreover, intelligence in childhood or early adulthood is associated with health outcomes later in life as well as mortality [5, 7–10]. Thus, understanding the mechanisms of cognitive abilities in children potentially has important implications for society and can be used to enhance such abilities, for example through targeted interventions such as education and the management of environmental risk factors [4, 11].

Neuroimaging can play a key role in advancing our understanding of the neural mechanisms of cognitive ability. Several brain-imaging studies have shown that total brain volume is the strongest brain imaging derived predictor of general intelligence [12–14] ( $r \approx 0.3 - 0.4$ ). To a somewhat lesser degree, regional cortical volume and thickness differences in the frontal, temporal, and parietal lobes have also been linked to intelligence [12, 13, 15–17]. Converging neuroimaging evidence led to the proposal of the parieto-frontal integration theory [18] whereby a distributed network of brain regions is responsible for the individual variability in cognitive abilities. This theory is also supported by human lesion studies [19, 20].

The ABCD-NP Challenge 2019 asked the question “*Can we predict fluid intelligence from T1-weighted MRI?*” We took an exploratory, data-driven approach to answering this question—a hackathon organised by our local research centres: the UCL Centre for Medical Image Computing (CMIC) and Wellcome/EPSRC Centre for Interventional and Surgical Sciences (WEISS). Our centres aim to address key medical challenges facing 21<sup>st</sup>-century society through world-leading research in medical imaging, medical image analysis, and computer-assisted interventions. Our expertise extends from feature extraction/generation through to image-based modelling [21], machine learning [22–24], and beyond. The hackathon involved researchers across research groups in our centres, in addition to colleagues from the affiliated Wellcome Centre for Human Neuroimaging and the Department of Clinical and Experimental Epilepsy at UCL. The hackathon took place on an afternoon in February 2019, after which we followed up with regular progress meetings.

In this paper we report our findings for predicting fluid intelligence in 9/10-year-olds from T1-weighted MRI using machine learning regression and deep learning methods (convolutional neural networks—CNNs). Our paper is structured as follows. The next section describes the challenge data and our methods. Section 3 presents our results, which we discuss in Sect. 4 before concluding.

## 2 Methods

### 2.1 Data


The ABCD-NP Challenge data consists of pre-processed T1-weighted MRI scans and fluid intelligence scores for children aged 9–10 years. The imaging protocol

can be found in [25]. Pre-processing included skull-stripping, noise removal, correction for field inhomogeneities [26, 27], and affine alignment to the SRI24 adult brain atlas [28]. SRI24 segmentations and corresponding volumes were also provided.

The cohort was split into training ( $N = 3739$ ), validation ( $N = 415$ ), and test ( $N = 4515$ ) sets. The training and validation sets also include scores of fluid intelligence, which are measured in the ABCD Study using the NIH Toolbox Neurocognition battery [29]. For the challenge, fluid intelligence was residualised to remove linear dependence upon brain volume, data collection site, age at baseline, sex at birth, race/ethnicity, highest parental education, parental income and parental marital status.

## 2.2 Features Derived from the Data

We trained the models to predict fluid intelligence both from the provided T1-weighted images (voxel intensity) as well as voxel-wise feature maps generated from these images using a probabilistic segmentation approach. There are many different methods to extract various features from T1-weighted images [30], such as tissue-type labels obtained from probabilistic segmentations. These segmentations can be constructed in a way to capture not only the relative tissue composition in a voxel, but also information about shape differences between individuals. This requires mapping each subject to a common template—a fundamental technique of computational anatomy [31]. Here we constructed such a template from all available T1-weighted MRI scans ( $S = \textit{train} + \textit{validation} + \textit{test} = 3,739 + 415 + 4,515 = 8,669$ ), which generated normalised (non-linearly aligned to a common mean) tissue segmentations for each subject.



**Fig. 1.** Template generated from fitting the generative model to all of the subjects in the ABCD population ( $S = 8,669$ ). Green corresponds to grey matter tissue, blue to white matter, and red to other. (Color figure online)

We used a generative model [32] to probabilistically segment each T1-weighted MRI in the challenge data set into three tissue types: grey-matter,

white-matter, other—see Fig. 1. The in-house model<sup>1</sup> used here contains key improvements over the one in [32]: (1) we place a smoothing prior on the template; (2) we obtain better initial values by first working on histogram representations of the images; (3) we normalise over population image intensities in a principled way, within the model; (4) we place a prior on the proportions of each tissue, which is also learned during training. There are two types of normalised segmentations: non-modulated and modulated. Modulated segmentations include the relative shape change when aligning to the common template.

Seven features per voxel were considered: T1-weighted intensity plus our six derived features corresponding to modulated and non-modulated probabilities for each tissue type. All images, including the feature maps from probabilistic segmentations, were spatially smoothed with a Gaussian kernel of 12 mm FWHM [30] and masked to remove voxels outside of the brain.

### 2.3 Predicting Fluid Intelligence: Machine Learning Regression

We explored several machine learning regression algorithms of varying complexity, including Multi-Kernel Learning (MKL) [33], Kernel Ridge Regression (KRR) [34], Gaussian Process Regression (GPR) [35] and Relevance Vector Machines [36]. The inputs to these models consisted of different concatenated combinations of our seven voxel-wise features described in Sect. 2.2. Analyses were run in *PRoNT*o version 3 [22] (<http://www.mnl.cs.ucl.ac.uk/pronto>), a software toolbox of pattern recognition techniques for the analysis of neuroimaging data, as well as custom-written code.

In our preliminary analyses, we trained different combinations of regression algorithms and features using 5-fold cross-validation within the training set to select the best combination of algorithm and features as measured by lowest cross-validated mean-squared error (MSE). We then retrained the best-performing model using the entire training set and used this trained model to generate predictions of fluid intelligence scores for the validation and test sets. Our best-performing model was KRR using all six voxel-wise derived features (tissue-type probabilities) concatenated into an input feature vector of length  $\sim 1.7$  million per individual. We set the regularization hyperparameter to  $\lambda = 10$  [34], which was optimised through 5-fold nested cross-validation within the training set in preliminary analyses.

We investigated robustness/stability of our KRR model using modified jackknife resampling (80/20 train/test split). Explicitly, we trained the model on a random subsample of 80% of the training set and generated predictions for both the held out 20% of the training set and the full validation set. We repeated this procedure 1000 times to generate confidence bounds on performance (MSE).

### 2.4 Predicting Fluid Intelligence: Convolutional Neural Networks

Separately, we explored the use of CNNs. The motivation was to incorporate spatial information that is not explicitly modelled as features. The CNNs were

<sup>1</sup> Available from <https://github.com/WCHN/segmentation-model>.

trained directly on the pre-processed T1-weighted images. Similarly to previous work that predicted brain age, Alzheimer’s disease progression, or brain regions from MRI scans [37–39] we applied various layers of 3D-convolutional kernels with filter size of  $3 \times 3 \times 3$  voxels on down-sampled images with dimensions of  $61 \times 61 \times 61$  voxels. We trained and validated multiple neural networks including those in [24]. Our best performing CNN (lowest MSE) consisted of four convolutional layers and three fully-connected layers followed by dropout layers with a probability of 0.5. The first six layers were activated by rectified linear units. The convolutional layers were followed by batch normalization and max-pooling operations. We used the Adam optimizer with an initial learning rate of  $10^{-5}$  and a decay of  $10^{-5}$ , and we stopped training at the 100<sup>th</sup> epoch. To evaluate the network, we randomly sampled 10 subsets of 1870 subjects from the training set, applied them to train 10 CNN models and evaluated performance on the validation set. We report MSE averaged over the ten passes.

### 3 Results

Our best-performing regression model was KRR using our six derived voxel-wise features as input. This produced  $\text{MSE} = 69.72$  on the validation set. By comparison, our best-performing CNN achieved  $\text{MSE} = 70.82$  (average  $\text{MSE} = 73.40$ ) on the validation set. However, we observed that the better-performing CNNs on the training set did not generalise well to the validation set, possibly due to over-fitting, or mismatch between training and validation data (as suggested by the considerable difference in variances).

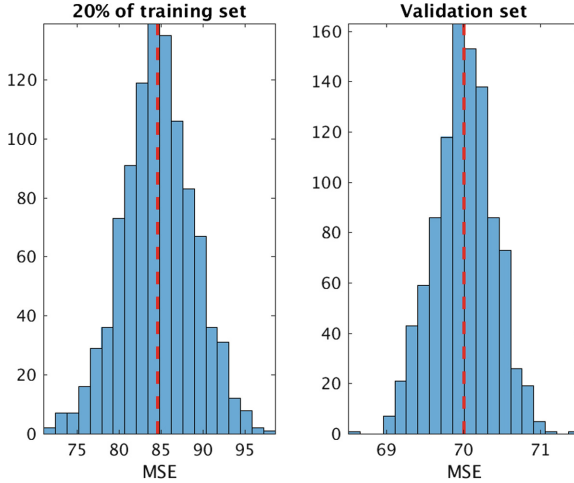
Table 1 shows MSE and Pearson’s correlation coefficient (mean  $\pm$  std) for training and validation predictions from our top-performing methods. We found that KRR performed the best and we used this model to generate our challenge submission, which resulted in  $\text{MSE} = 92.13$  on the test set.

**Table 1.** Model performance for the training and validation sets.

Method	Training set		Validation set		Test set
	MSE	Correlation	MSE	Correlation	MSE
KRR	77.64	0.1427	<b>69.72</b>	0.0311	<b>92.13</b>
CNN	$69.10 \pm 3.13$ (best: 64.39)	$0.1913 \pm 0.0423$ (best: 0.2542)	$73.40 \pm 2.13$ (best: 70.82)	$-0.0202 \pm 0.0296$ (best: 0.0157)	N/A

Figure 2 shows model robustness results from modified jackknife resampling (1000 repetitions, 80/20 split). Confidence in our predicted MSE on the validation set is within  $\pm 1$  residual IQ point.

Figure 3 shows the residual fluid intelligence score prediction errors for KRR on the training and validation sets. The V shape of the curves shows that smaller residual scores have lower errors, since the model is predicting close to the mean value.



**Fig. 2.** Model robustness for KRR using modified jackknife re-sampling. MSE distributions for 1000 predictions: 20% of the training set (left); full validation set (right); for models trained on the other 80% of the training set. Red line indicates the mean. (Color figure online)

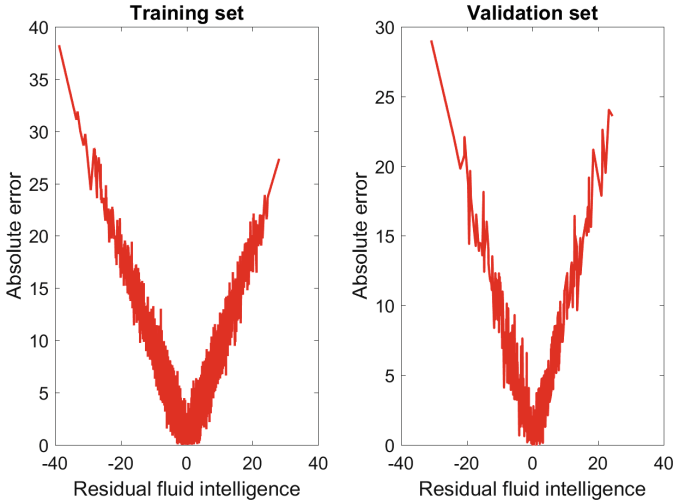
## 4 Discussion

We found that predicting residual fluid intelligence from structural MRI images is challenging. The correlation between predicted and actual intelligence scores was low for all methods we tested ( $r \approx 0.05-0.15$ ). This contrasts with previous studies for predicting (non-residualised) fluid intelligence, which have demonstrated that both total brain volume and regional cortical volume/thickness differences are relatively strong predictors ( $r \approx 0.3-0.4$ ) [12–14].

The lower predictive performance we observed might be influenced by the residualisation, which prevents modelling of covariance between the residualisation factors and the image-based features. Moreover, there is evidence that including variables in the residualisation procedure that are correlated with the regression targets/labels is likely to remove important variability in the data leading to predictive models with low performance [40].

We note that previous studies used small sample sizes (of order 10–100), whilst the ABCD-NP Challenge dataset comprised a very large-scale dataset (of order 1000–10000). Whereas subject recruitment for small samples tend to be well controlled, resulting in homogeneous sample characteristics, large samples are more heterogeneous by nature, thus predictive models are more challenging to build. Accordingly, a recent study has demonstrated that the accuracy of classification results tends to be smaller for larger sample sizes [41].

Our image-based features were voxel-wise probabilistic tissue-type labels: grey matter, white matter, and other. Beyond tissue-type labels, there might be value in investigating other features generated by the generative segmentation



**Fig. 3.** Absolute prediction errors against residual fluid intelligence scores for KRR (training set on the left, validation set on the right).

model. One such interesting feature is scalar momentum [30], which has been shown to be predictive for a range of different problems [30].

Finally, we mention a possible limitation of our approach. In recent years, it has become standard practice to create study-specific group templates in neuroimaging [31], especially in Voxel-Based Morphometry (VBM) analyses. Originally, this approach was proposed for group analysis using mass univariate statistics (e.g., statistical parametric mapping). However, one needs to exercise caution when applying such an approach in machine learning, as it might lead to slightly optimistic predictions by creating dependence across the overall dataset. In order to avoid this potential issue, one would need to create templates based only on the training set. This might be computationally challenging when cross-validation strategies are used. To the best of our knowledge, no studies have investigated whether study-specific templates indeed result in slightly inflated predictions, and it remains an interesting question for future work.

## 5 Conclusion

Our paper presents the winning method for the ABCD Neurocognitive Prediction Challenge 2019. We found that kernel ridge regression outperformed more complex models, such as convolutional neural networks, when predicting residual fluid intelligence scores for the challenge dataset using our custom tissue-type features derived from the preprocessed T1-weighted MRI. The correlation between the predicted and actual scores is very low ( $r = 0.03$  for the KRR on the validation set), implying that the association between structural images and

residualised fluid intelligence scores is low. It may be that structural images contain very little information on residualised fluid intelligence, but further study is warranted.

**Acknowledgements.** This study was funded by the UCL Centre for Medical Image Computing and UK EPSRC platform grant “Medical image computing for next-generation healthcare technology” (EP/M020533/1) and supported by researchers at the National Institute for Health Research University College London Hospitals Biomedical Research Centre. FSF is funded by a PhD scholarship awarded by Fundacao para a Ciencia e a Tecnologia (SFRH/BD/120640/2016). NPO acknowledges support from the NIHR UCLH Biomedical Research Centre and the [EuroPOND project](#)— This project has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No. 666992. AM, CZ, TW, and JM-M acknowledge funding from the Wellcome Trust under grant number WT102845/Z/13/Z.

## References

1. Goriounova, N.A., Mansvelder, H.D.: Genes, cells and brain areas of intelligence. *Front. Hum. Neurosci.* **13**, 44 (2019). <https://doi.org/10.3389/fnhum.2019.00044>
2. Deary, I.J., Penke, L., Johnson, W.: The neuroscience of human intelligence differences. *Nat. Rev. Neurosci.* **11**(3), 201–211 (2010). <https://doi.org/10.1038/nrn2793>
3. McCall, R.B.: Childhood IQ’s as predictors of adult educational and occupational status. *Science* **197**(4302), 482–483 (1977). <https://doi.org/10.1126/science.197.4302.482>
4. Gottfredson, L.S.: Why g matters: the complexity of everyday life. *Intelligence* **24**(1), 79–132 (1997). [https://doi.org/10.1016/S0160-2896\(97\)90014-3](https://doi.org/10.1016/S0160-2896(97)90014-3)
5. Deary, I.J., Strand, S., Smith, P., Fernandes, C.: Intelligence and educational achievement. *Intelligence* **35**(1), 13–21 (2007). <https://doi.org/10.1016/j.intell.2006.02.001>
6. Johnson, W., McGue, M., Iacono, W.G.: Genetic and environmental influences on academic achievement trajectories during adolescence. *Dev. Psychol.* **42**(3), 514–32 (2006). <https://doi.org/10.1037/0012-1649.42.3.514>
7. Batty, G.D., Deary, I.J., Gottfredson, L.S.: Premorbid (early life) IQ and later mortality risk: systematic review. *Ann. Epidemiol.* **17**(4), 278–288 (2007). <https://doi.org/10.1016/j.annepidem.2006.07.010>
8. Batty, G.D., et al.: IQ in early adulthood and mortality by middle age. *Epidemiology* **20**(1), 100–109 (2008). <https://doi.org/10.1097/ede.0b013e31818ba076>
9. Lam, N.H., et al.: Effects of Altered Excitation-Inhibition Balance on Decision Making in a Cortical Circuit Model. *bioRxiv*, p. 100347 (2017). <https://doi.org/10.1101/100347>
10. Deary, I.J., Pattie, A., Starr, J.M.: The stability of intelligence from age 11 to age 90 years: the lothian birth cohort of 1921. *Psychol. Sci.* **24**(12), 2361–2368 (2013). <https://doi.org/10.1177/0956797613486487>
11. Fors, S., Torssander, J., Almquist, Y.B.: Is childhood intelligence associated with coexisting disadvantages in adulthood? Evidence from a Swedish cohort study. *Adv. Life Course Res.* **38**, 12–21 (2018). <https://doi.org/10.1016/J.ALCR.2018.10.005>



12. MacLulich, A.M.J., Ferguson, K.J., Deary, I.J., Seckl, J.R., Starr, J.M., Wardlaw, J.M.: Intracranial capacity and brain volumes are associated with cognition in healthy elderly men. *Neurology* **59**(2), 169–174 (2002). <https://doi.org/10.1212/WNL.59.2.169>
13. McDaniel, M.A.: Big-brained people are smarter: a meta-analysis of the relationship between in vivo brain volume and intelligence. *Intelligence* **33**(4), 337–346 (2005). <https://doi.org/10.1016/j.intell.2004.11.005>
14. Rushton, J.P., Ankney, C.D.: Whole brain size and general mental ability: a review. *Int. J. Neurosci.* **119**(5), 692–732 (2009). <https://doi.org/10.1080/00207450802325843>
15. Andreasen, N.C., et al.: Intelligence and brain structure in normal individuals. *Am. J. Psychiatry* **150**(1), 130–4 (1993). <https://doi.org/10.1176/ajp.150.1.130>
16. Narr, K.L., et al.: Relationships between IQ and regional cortical gray matter thickness in healthy adults. *Cereb. Cortex* **17**(9), 2163–2171 (2007). <https://doi.org/10.1093/cercor/bhl125>
17. Karama, S., et al.: Cortical thickness correlates of specific cognitive performance accounted for by the general factor of intelligence in healthy children aged 6 to 18. *NeuroImage* **55**(4), 1443–1453 (2011). <https://doi.org/10.1016/j.neuroimage.2011.01.016>
18. Jung, R.E., Haier, R.J.: The parieto-frontal integration theory (P-FIT) of intelligence: converging neuroimaging evidence. *Behav. Brain Sci.* **30**(2), 135–154 discussion 154–187 (2007). <https://doi.org/10.1017/S0140525X07001185>
19. Gläscher, J., et al.: Lesion mapping of cognitive abilities linked to intelligence. *Neuron* **61**(5), 681–91 (2009). <https://doi.org/10.1016/j.neuron.2009.01.026>
20. Woolgar, A., et al.: Fluid intelligence loss linked to restricted regions of damage within frontal and parietal cortex. *Proc. Nat. Acad. Sci.* **107**(33), 14899–14902 (2010). <https://doi.org/10.1073/pnas.1007928107>
21. Oxtoby, N.P., Alexander, D.C., for the EuroPOND consortium: Imaging plus X: multimodal models of neurodegenerative disease. *Curr. Opin. Neurol.* **30**(4), 371–379 (2017). <https://doi.org/10.1097/WCO.0000000000000460>
22. Schrouff, J., et al.: PRoNTTo: pattern recognition for neuroimaging toolbox. *Neuroinformatics* **11**(3), 319–37 (2013). <https://doi.org/10.1007/s12021-013-9178-1>
23. Schrouff, J., Monteiro, J.M., Portugal, L., Rosa, M.J., Phillips, C., Mourão-Miranda, J.: Embedding anatomical or functional knowledge in whole-brain multiple kernel learning models. *Neuroinformatics* **16**(1), 117–143 (2018). <https://doi.org/10.1007/s12021-017-9347-8>
24. Blumberg, S.B., Tanno, R., Kokkinos, I., Alexander, D.C.: Deeper image quality transfer: training low-memory neural networks for 3D images. In: Frangi, A.F., Schnabel, J.A., Davatzikos, C., Alberola-López, C., Fichtinger, G. (eds.) *MICCAI 2018*. LNCS, vol. 11070, pp. 118–125. Springer, Cham (2018). [https://doi.org/10.1007/978-3-030-00928-1\\_14](https://doi.org/10.1007/978-3-030-00928-1_14)
25. Casey, B.J., et al.: The adolescent brain cognitive development (ABCD) study: imaging acquisition across 21 sites. *Dev. Cogn. Neurosci.* **32**, 43–54 (2018). <https://doi.org/10.1016/j.dcn.2018.03.001>
26. Hagler, D.J., et al.: Image processing and analysis methods for the Adolescent Brain Cognitive Development Study. *bioRxiv*, p. 457739 (2018). <https://doi.org/10.1101/457739>
27. Pfefferbaum, A., et al.: Altered brain developmental trajectories in adolescents after initiating drinking. *Am. J. Psychiatry* **175**(4), 370–380 (2018). <https://doi.org/10.1176/appi.ajp.2017.17040469>

28. Rohlfing, T., Zahr, N.M., Sullivan, E.V., Pfefferbaum, A.: The SRI24 multichannel atlas of normal adult human brain structure. *Hum. Brain Mapp.* **31**(5), 798–819 (2010). <https://doi.org/10.1002/hbm.20906>
29. Akshoomoff, N., et al.: VIII. NIH toolbox cognition battery (CB): composite scores of crystallized, fluid, and overall cognition. *Monogr. Soc. Res. Child Dev.* **78**(4), 119–132 (2013). <https://doi.org/10.1111/mono.12038>
30. Monté-Rubio, G.C., Falcón, C., Pomarol-Clotet, E., Ashburner, J.: A comparison of various MRI feature types for characterizing whole brain anatomical differences using linear pattern recognition methods. *NeuroImage* **178**, 753–768 (2018). <https://doi.org/10.1016/j.neuroimage.2018.05.065>
31. Ashburner, J.: A fast diffeomorphic image registration algorithm. *NeuroImage* **38**(1), 95–113 (2007). <https://doi.org/10.1016/J.NEUROIMAGE.2007.07.007>
32. Blaiotta, C., Freund, P., Cardoso, M.J., Ashburner, J.: Generative diffeomorphic modelling of large MRI data sets for probabilistic template construction. *NeuroImage* **166**, 117–134 (2018). <https://doi.org/10.1016/j.neuroimage.2017.10.060>
33. Rakotomamonjy, A., Bach, F.R., Canu, S., Grandvalet, Y.: SimpleMKL. *J. Mach. Learn. Res.* **9**, 2491–2521 (2008). <http://www.jmlr.org/papers/v9/rakotomamonjy08a.html>
34. Shawe-Taylor, J., Cristianini, N.: *Kernel Methods for Pattern Analysis*. Cambridge University Press, Cambridge (2004)
35. Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge (2006). <http://www.GaussianProcess.org/gpml>
36. Tipping, M.E.: Sparse Bayesian learning and the relevance vector machine. *J. Mach. Learn. Res.* **1**, 211–244 (2001). <http://www.jmlr.org/papers/v1/tipping01a.html>
37. Sturmfels, P., Rutherford, S., Angstadt, M., Peterson, M., Sripada, C., Wiens, J.: A domain guided CNN architecture for predicting age from structural brain images. In: Doshi-Velez, F., et al. (eds.) *Proceedings of the 3rd Machine Learning for Healthcare Conference*. *Proceedings of Machine Learning Research*, Palo Alto, California, vol. 85, pp. 295–311. PMLR (2018). <http://proceedings.mlr.press/v85/sturmfels18a.html>
38. Payan, A., Montana, G.: Predicting Alzheimer’s disease: a neuroimaging study with 3D convolutional neural networks. arXiv e-prints [arXiv:1502.02506](https://arxiv.org/abs/1502.02506), February 2015. Preprint
39. Milletari, F., et al.: Hough-CNN: Deep Learning for Segmentation of Deep Brain Regions in MRI and Ultrasound. arXiv e-prints [arXiv:1601.07014](https://arxiv.org/abs/1601.07014) (2016)
40. Rao, A., Monteiro, J.M., Mourao-Miranda, J.: Alzheimer’s disease initiative: predictive modelling using neuroimaging data in the presence of confounds. *NeuroImage* **150**, 23–49 (2017). <https://doi.org/10.1016/j.neuroimage.2017.01.066>
41. Varoquaux, G.: Cross-validation failure: small sample sizes lead to large error bars. *NeuroImage* **180**, 68–77 (2018). <https://doi.org/10.1016/j.neuroimage.2017.06.061>