

Never a Dull Moment: Distributional Properties as a Baseline for Time-Series Classification

Trent Henderson¹[0009-0005-5467-9914], Annie G. Bryant¹[0000-0002-2217-6520],
and Ben D. Fulcher¹[0000-0002-3003-4055]

The University of Sydney, Camperdown NSW 2006, Australia.

Abstract. The variety of complex algorithmic approaches for tackling time-series classification problems has grown considerably over the past decades, including the development of sophisticated but challenging-to-interpret deep-learning-based methods. But without comparison to simpler methods it can be difficult to determine when such complexity is required to obtain strong performance on a given problem. Here we evaluate the performance of an extremely simple classification approach—a linear classifier in the space of two simple features that ignore the sequential ordering of the data: the mean and standard deviation of time-series values. Across a large repository of 128 univariate time-series classification problems, this simple distributional moment-based approach outperformed chance on 69 problems, and reached 100% accuracy on two problems. With a neuroimaging time-series case study, we find that a simple linear model based on the mean and standard deviation performs better at classifying individuals with schizophrenia than a model that additionally includes features of the time-series dynamics. Comparing the performance of simple distributional features of a time series provides important context for interpreting the performance of complex time-series classification models, which may not always be required to obtain high accuracy.

Keywords: Time series · Time-series classification · Benchmarking.

1 Introduction

Time-series classification is a key problem in the sciences and industry wherein time-varying data is used to distinguish labeled classes. The quantity and diversity of time-series classification algorithms is large and increasing, from simple linear decision boundaries in interpretable feature spaces [1] to complex methods based on deep neural networks, such as long short-term memory networks [2]. Complex new algorithms can yield impressive classification accuracy on challenging problems, but often at the expense of clear human interpretability [3].

The UEA/UCR univariate time-series classification repository, which currently contains 128 problems spanning a variety of domains [4], has been crucial for encouraging transparent reporting of the relative strengths and weaknesses of classification algorithms. It has also enabled the field to overcome key

limitations in prior standard practice of reporting and comparison, including avoiding cherry-picking of optimistic datasets when reporting algorithm performance [5]. Systematic comparisons of time-series classification algorithms across this database have been essential for benchmarking the accuracy of state-of-the-art algorithms [6]. However, in settings ranging from policy-making [7] to healthcare [8], deriving interpretable understanding that can guide subsequent decision-making can be more important than raw classification accuracy. In such settings, simpler methods that are faster to train and clearer to interpret, are often preferred.

Recent work has shown that parsimonious and interpretable methods can match (or even outperform) more sophisticated ones, in settings ranging from sleep-stage classification [9] to earthquake detection [10]. A particularly striking example is the recent demonstration of a two-parameter logistic regression model with equivalent performance in earthquake aftershock forecasting to a deep neural network with thousands of parameters [11]. For such settings, in which simple methods can perform well, the choice to instead use overly complex models risks over-fitting, leading to poor generalizability on out-of-sample data. This effect has been demonstrated in a recent meta-analysis of deep-learning models for classifying autism spectrum disorder from resting-state neuroimaging data, which exhibited inferior performance to linear support vector machine (SVM) models on unseen data [12].

Defaulting to complex classification models and comparing their performance to chance-level accuracy—or even other complex methods—can thus result in classifiers that are over-complicated and difficult to interpret. This issue is well-illustrated by the time-series classification task of distinguishing ‘epilepsy’ from ‘eyes-open’ states from electroencephalogram (EEG) data [13]. While state-of-the-art approaches had applied complex algorithmic approaches, ranging from independent components analysis and discrete wavelet transforms to multi-layer neural networks [14], it was found that the time-series standard deviation alone could completely separate the two classes [15]. The strong performance of a simple threshold classifier on standard deviation thus suggests it as an interpretable and parsimonious alternative to overly complicated algorithmic approaches to this problem, highlighting the utility of starting with a simple and parsimonious model and building in complexity only when it yields clearly demonstrable benefits.

Given that the characteristic complexity of time-series classification tasks (compared to classification using non-sequential data) relates to the challenge of quantifying class-informative dynamical patterns, here we aimed to investigate the performance of an extremely simple benchmark: a linear classifier in the two-dimensional space of two extracted features that ignore the sequential ordering of the data entirely—the mean and standard deviation. As these features instead characterize the distribution of time-series values, high performance of this simple benchmark indicates problems for which trivial properties of the distribution are already sufficient to perform well, undermining the need for more complex approaches that aim to quantify informative temporal patterns [16]. While some

previous work [1, 17] has z -scored all time series prior to analysis—thereby insuring mean and variance are uninformative, to focus on dynamical patterns—such normalization has not been applied consistently across all problems in the UEA/UCR Repository [6]. We also investigate the performance gains of additionally incorporating a simple set of time-series features that capture different types of dynamical patterns in the data, using the *catch22* feature set [17]. None of these features are sensitive to the mean or standard deviation of the data; they instead capture more subtle properties of the time series, including properties of distributional shape; periodic patterns; temporal spacing of outliers; and linear and nonlinear autocorrelation [17]. Finally, to demonstrate the practical implications of our findings, we present a case study of classifying individuals with schizophrenia using functional neuroimaging time series.

2 Methods

This study aims to evaluate the performance of a simple baseline linear classifier for time series in the two-dimensional space of their mean (μ) and standard deviation (σ)—features related to the First Two Moments (FTM) of the distribution. We first discuss methods related to fitting the FTM and *catch22* feature-based time-series classifiers across the UEA/UCR Repository [4] (in Sec. 2.1) and then describe specific methods related to our neuroimaging case study (in Sec. 2.2). Code for reproducing our results is available [18].

2.1 Feature-based classifier performance on the UEA/UCR Repository

In the UEA/UCR Time-Series Classification Repository [4], each of the 128 problems (as downloaded on 9 Feb 2023) are partitioned into designated train and test sets. These problems vary widely in the number of time series, the relative size of train and test set sizes, time-series lengths, and number of classes (see [4, 6] for details). For all time series, we computed the FTM feature set containing two features: the mean (μ) and standard deviation (σ). For comparison to more sophisticated features of the time-series dynamics, we used the *catch22* set of 22 time-series features [17]. Given a feature space, we fit and evaluated classification models following the resample-based procedure outlined in [6], using 30 resamples of train–test splits. In addition to the designated split, we generated 29 additional (seeded) resamples of the data that preserve the class proportions of the designated split, for a total of 30 train–test splits. Prior to fitting a linear SVM, features were normalized as a z -score by computing the mean and standard deviation of each feature in the train set, and using these values to rescale both the train and test sets (ensuring that test data was completely unseen). A linear SVM was fit on the train set and used to generate predictions for the test set. Classification accuracy was used as the performance metric, following prior work. To statistically compare the performance of FTM, and the union set of 24 features from both *catch22* and FTM (*catch22* + FTM), we implemented a correlated

test statistic [19] that corrects the traditional T -statistic to account for the violation of the assumption of independence incurred in the usage of resampling using the *correctR* package [20]. To compare the FTM and *catch22* + FTM feature sets, we removed four problems (`AllGestureWiimoteX`, `AllGestureWiimoteY`, `AllGestureWiimoteX`, and `PLAID`) due to presence of time series that had < 10 real values before containing all or mostly missing values—the minimum threshold for calculations in *catch22*.

2.2 Case study: Schizophrenia classification

As a case study, we investigated the performance of simple feature-based classifiers to distinguish adults with schizophrenia versus cognitively healthy controls based on their whole-brain activity dynamics. We obtained resting-state functional Magnetic Resonance Imaging (rs-fMRI) data from the University of California at Los Angeles Consortium for Neuropsychiatric Phenomics LA5c Study [21], which was pre-processed using the ICA-AROMA + 2P + GMR method, as described previously [22]. Blood Oxygen Level-Dependent (BOLD) signals consisting of 152 time samples were extracted for each of 68 cortical [23] and 14 subcortical [24] regions per participant. Any participants with constant time series across all brain regions after preprocessing were excluded, yielding a final sample of $N = 166$ participants ($N = 118$ control and $N = 48$ schizophrenia). The two groups differ in terms of age (Control = 31.5 ± 8.8 years; Schizophrenia = 36.6 ± 9.0 years; Wilcoxon rank sum test $p < 0.01$) and sex (Control = 46.6% Female; Schizophrenia = 25% Female; $X^2(1, N = 166) = 5.75, p < 0.05$).

After feature extraction (using *theft* [25]), the dataset was in the form of an $N \times R \times F$ matrix, for $N = 166$ subjects, $R = 82$ brain regions, and F features. We incorporated all combinations of 82 brain regions and features (either $F = 2$ FTM features or $F = 24$ FTM + *catch22* features) as inputs to a regularized linear SVM classifier (from *scikit-learn* [26]), setting the regularization parameter $C = 1$, disabling the shrinking heuristic, and using balanced class weights. In other words, the FTM-only model had a total of $2 \times 82 = 164$ SVM input features and the FTM + *catch22* model had a total of $24 \times 82 = 1968$ features.

Model performance was evaluated as balanced accuracy using 10-fold cross-validation (CV) with 10 repeats (setting the random state such that all compared models received the same resampled sets of test folds). For each k -fold, the same z -score feature-normalization was applied as described in Sec. 2.1 above. Balanced accuracy is reported as the mean \pm SD across the 10 repeats, in which each repeat contains the mean balanced accuracy across the 10 CV folds. Statistical significance of a given balanced accuracy value was assessed using permutation testing with 1000 null samples obtained from shuffling class labels, pooling null samples across all combinations of regions and features for each feature set [27]. False positives were controlled at the $\alpha = 0.05$ level using Bonferroni correction.

3 Results

Results are structured as follows. First, in Sec. 3.1, we describe findings across the UEA/UCR Repository, where we demonstrate that the FTM classifier outperforms chance on the majority of problems. We then highlight the practical ramifications of these results using the example of a neuroimaging biomarker classification task in Sec. 3.2, in which the mean and standard deviation exhibit strong performance that is weakened by adding dynamical properties of the functional neuroimaging time series.

3.1 Linear FTM classifier performance across the UEA/UCR Database

We first investigated how a linear classifier based on the two FTM features, (μ, σ) , performs across the 128 problems from the UEA/UCR Repository. To determine the problems on which FTM beats chance, we calculated a p -value of the chance probability against the distribution of 30 resampled FTM accuracy values for each problem. We found that the FTM-based classifier statistically outperformed chance ($p < 0.05$) on 69 of the 128 problems (i.e., on 53.9% of problems). FTM-based classification accuracies relative to chance level are plotted for these 69 problems in Fig. 1. We note that chance is often beaten by a substantial margin on these problems. Indeed, this simple classifier achieves 100% accuracy on two problems: `InsectEPGRegularTrain` and `GunPointOldVersusYoung`. For the 59 problems where the FTM-based classifier did not beat chance, such a finding could be due to two reasons: (i) the dataset has been correctly z -scored in the UEA/UCR Repository (meaning each time series has a mean of zero and standard deviation of one); or (ii) the FTM features are not informative of class differences for that problem. These results demonstrate that, even on a repository devoted to time-series classification tasks, strong performance can be obtained using the two simplest distributional statistics that are unrelated to the sequential ordering of the data, due to the database containing problems in which labeled classes have distinctive levels (means) or scales (variances).

To better understand how the FTM classifier is behaving, we examined the two problems for which it achieved 100% accuracy. For `GunPointOldVersusYoung`, the classes are clearly distinguished in the (μ, σ) feature space (Fig. 2A), while for `InsectEPGRegularTrain` each class has a characteristic mean level (Fig. 2B). For the binary (‘Young’ vs ‘Old’) `GunPointOldVersusYoung` task, the time series is the x -axis coordinate of the centroid of the actor’s right hand at each frame while moving it from rest position to a gun pose and back again [28]. Figure 2A shows that the ‘Young’ actor exhibits characteristically lower mean and variation of their hand coordinate, consistent with them being shorter and having a shorter arm than the ‘Old’ actor. It is worth noting that, in contrast to features of the time-series dynamics (many of which are invariant to linear rescalings of the time-series values), μ and σ are highly sensitive to the calibration of experimental measurement. For the `GunPoint` collection of problems, this makes these features less likely to generalize well to new examples, with small differences in




Fig. 1: Mean and standard deviation as features in a linear SVM classifier statistically outperforms chance-level accuracy on 69 of 128 problems in the UEA/UCR time-series repository. Classification accuracy (%) is displayed along the horizontal axis and the name of each problem is shown on the vertical axis. Chance accuracies are displayed as black crosses. Points indicate mean classification accuracy and error bars show ± 1 standard deviation (across train-test resamples).

the angle of the camera used to measure hand coordinates, or the distance from camera to actor, strongly affecting differences in μ and σ and hence accurate classification [29]. In `InsectEPGRegularTrain`, each class has a characteristic mean voltage level of the electrical circuit that connects insects with their food source, allowing the classes to be accurately distinguished by ignoring sequential patterns, and simply focusing on this mean voltage.

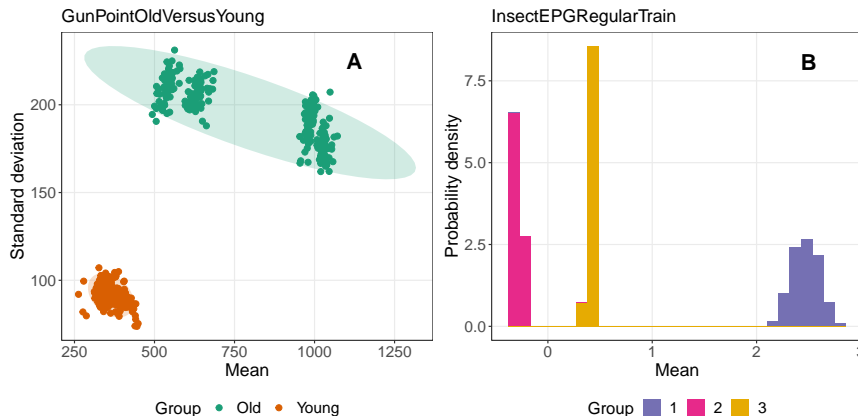


Fig. 2: Simple distributional moments can perfectly separate classes in the `GunPointOldVersusYoung` and `InsectEPGRegularTrain` datasets. **A** Individual time series in `GunPointOldVersusYoung` are represented in the two-dimensional feature space of mean and standard deviation. Points are colored by class (as labeled) and class-level covariances are shown as shaded ellipses to guide the eye. **B** Histogram of time-series means by class in `InsectEPGRegularTrain`.

Finally, we aimed to investigate the extent to which measuring the *catch22* features in addition to the FTM features would improve classification performance. We compared the performance of the FTM feature set (2 features) to that of the *catch22* + FTM feature set (24 features) across the 124 problems where *catch22* features could be successfully calculated (four problems were excluded due to a high number of missing values, cf. Methods). Adding *catch22* features resulted in an average absolute improvement in classification accuracy of 33.7% across the 124 problems, confirming the general importance of capturing dynamical properties for time-series classification problems. However, pairwise comparisons using the corrected test statistic revealed that there was no statistical difference between FTM and *catch22* + FTM on 45 of the 124 (or 36.4% of) problems. This demonstrates that, for many problems, simple distributional properties yield a surprisingly strong baseline against which to assess the benefits gained by using more complex approaches.

3.2 Neuroimaging Biomarker Case Study




Fig. 3: The FTM feature set shows strong performance at classifying individuals with schizophrenia from resting-state fMRI time series. The distribution of mean balanced accuracy (across 10 repeats), using the combination of all brain regions and either the two *FTM* features (green), or the 24 *FTM* + *catch22* features (orange), is shown as a combined spaghetti plot and boxplot. Each gray line indicates the same resampled set of $k = 10$ folds evaluated for each feature set along the x -axis. **: $p = 0.003$, corrected resampled $T = 3.63$.

We next extended our investigation of simple baseline time-series classifiers to a schizophrenia classification problem using rs-fMRI time series, a setting in which complex, opaque classifiers are common [30, 31]. Given the strong performance of simple distributional statistics for some time-series classification problems above, we evaluated the performance of FTM alone (and with the addition of the *catch22* feature set) in classifying individuals with schizophrenia from healthy controls.

As shown in Fig. 3, the classification model based on FTM features displayed a high balanced accuracy ($67.5 \pm 2.2\%$), which sits well within the range of a recent meta-analysis of schizophrenia classification studies using rs-fMRI [32]. Adding *catch22* features to the model decreased its performance to $63.6 \pm 1.0\%$ (corrected resampled T -statistic = 3.63, $p = 0.003$). Consistent with the findings on some problems in the UEA/UCR Repository above, these results demonstrate the surprisingly strong performance of basic properties of the distribution of time-series values for this fMRI classification task. Surprisingly, for this rich and complex whole-brain time-series dataset, incorporating information about subtle patterns in the data (using *catch22*) yields models with inferior performance than a simple linear model based on the mean and standard deviation.

4 Discussion

With the growing sophistication of machine-learning algorithms, modern data analysts have increasingly adopted a methodological approach that defaults to more complex statistical methods, which can obscure their clear interpretation. Presenting such complex approaches without direct comparison to simpler alternatives makes it difficult to discern whether the methodological complexity is beneficial relative to simpler approaches, an assumption that has been challenged in some recent studies [9–11]. In this work, focusing on time-series classification problems, we demonstrate that a method at the extreme end of simplicity (using just two distributional-moment-based features, μ and σ , and a linear classifier) performs surprisingly well on many such problems. Despite being insensitive to the unique property of sequential data relative to an unordered vector—its ordering (in time)—this simple benchmark classifier statistically outperformed chance on approximately half of the problems in a prominent time-series classification archive, the UEA/UCR Repository [4], and even achieved 100% accuracy on two problems. Our results emphasize the importance of carefully considering the factors that contribute to model performance, assessing model gains relative to simpler models, and favoring parsimony by carefully building up model complexity incrementally from simple baselines.

We demonstrated the applicability of our findings to a time-series classification problem in neuroimaging, a setting in which many prior models have been highly complex models [30] but can fail to generalize [12]. Importantly, classification performance was stronger with FTM than with a model that also included *catch22* features of the dynamics using a linear SVM classifier. The combination of σ and μ yielded a mean balanced accuracy of 67.5%, while being much clearer to interpret, both statistically and biologically.

The strong performance of mean and standard deviation across the UEA/UCR Repository demonstrates that time series in these problems have not been consistently normalized, as prior work has highlighted [6]. If all time series were individually z -score transformed, there would be no class differences in mean or standard deviation for any problem (and the FTM-based classifier would exhibit null performance). Our results have implications for comparing time-series classification algorithms on problems for which labeled classes can be distinguished based on distributional properties alone. For example, consider the *catch22* features, which are all insensitive to the mean and variance of the input time series [17]. Relative to *catch22*, the superior performance of an alternative time-series classifier (which *is* sensitive to μ or σ of the input series) could be driven entirely by a class-level difference in μ or σ : properties that are unrelated to dynamical patterns (and that *catch22* cannot access). One approach to testing the ability of different classification algorithms to capture properties related to the dynamics of a uniformly sampled univariate time series would be to normalize time series such that there are no class differences in basic distributional properties. A second approach would be to compare model performance to a benchmark in which only simple distributional properties are included, with gains relative to this benchmark then being attributable to class-relevant differences in more

complex properties. In this latter approach, there is scope for extending the two simple FTM features used here by adding higher-order moments, or other types of distributional features (such as those included in comprehensive time-series feature sets like *hctsa* [33]).

This work also highlights the need for careful consideration of the generalizability of models that use features sensitive to the calibration of time-series measurements (like μ and σ), relative to features that are invariant to linear rescalings of the input time series (like the *catch22* feature set). This is because such measurement-scale dependent features are highly sensitive to changes in the calibration of experimental measurements that may not be precisely maintained in new data. An illustrative example is shown here in the `GunPointOldVersusYoung` dataset, for which our FTM-based method achieved 100% accuracy. However, this may be a deceptively high value, as it is reliant on the precise calibration of new data. A prior analysis of this dataset, focusing on early classification, has shown that classification accuracy can plummet with only slight variability in experimental calibration (to changes as small as a $\approx 1.9^\circ$ tilt in the angle of the camera used to measure hand coordinates) [29]. In general, the decision of whether or not to include features that are sensitive to measurement scale (or focus on scale-invariant features of the dynamics or distribution shape, as in *catch22*) should be motivated by domain expertise to avoid overly optimistic classification results.

In summary, by highlighting many time-series classification problems for which simple distributional properties of a time series can achieve surprisingly high classification accuracy, our results raise important issues for the development and interpretation of time-series classification models. Future work on evaluating time-series classification algorithms may consider using simple benchmarks for comparison to aid interpretation and provide evidence for the contribution of model complexity to any performance advantage, particularly for problems highlighted here for which simple distributional features are highly informative of class differences.

Acknowledgements

The authors would like to thank Kevin Aquino for sharing preprocessed fMRI data [34] used in the case study, and Eamonn Keogh for providing useful feedback on a manuscript draft.

References

1. Fulcher, B.D., Jones, N.S.: Highly comparative feature-based time-series classification. *IEEE Transactions on Knowledge and Data Engineering* **26**, 3026 (2014)
2. Ismail Fawaz, H., Forestier, G., Weber, J., Idoumghar, L., Muller, P.A.: Deep learning for time series classification: A review. *Data Mining and Knowledge Discovery* **33**, 917 (2019)
3. Rudin, C.: Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence* **1**, 206 (2019)

4. Bagnall, A., Lines, J., Vickers, W., Keogh, E.: The uea & ucr time series classification repository, www.timeseriesclassification.com (2022)
5. Keogh, E., Kasetty, S.: On the Need for Time Series Data Mining Benchmarks: A Survey and Empirical Demonstration. *Data Mining and Knowledge Discovery* **7**, 349 (2003)
6. Bagnall, A., Lines, J., Bostrom, A., Large, J., Keogh, E.: The great time series classification bake off: A review and experimental evaluation of recent algorithmic advances. *Data Mining and Knowledge Discovery* **31**, 606 (2017)
7. Goodman, B., Flaxman, S.: European union regulations on algorithmic decision-making and a “right to explanation”. *AI Magazine* **38**, 50 (2017)
8. Caruana, R., Lou, Y., Gehrke, J., Koch, P., Sturm, M., Elhadad, N.: Intelligible models for healthcare: Predicting pneumonia risk and hospital 30-day readmission. In: *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. p. 1721. Association for Computing Machinery (2015)
9. Van Der Donckt, J., Van Der Donckt, J., Deprost, E., Vandebussche, N., Rademaker, M., Vandewiele, G., Van Hoecke, S.: Do not sleep on traditional machine learning: Simple and interpretable techniques are competitive to deep learning for sleep scoring. *Biomedical Signal Processing and Control* **81**, 104429 (2023)
10. Waheed, U., Afify, A., Fehler, M., Fulcher, B.: Winning with simple learning models: Detecting earthquakes in groningen, the netherlands. *EAGE 2020 Annual Conference & Exhibition Online* **2020**, 1 (2020)
11. Mignan, A., Broccardo, M.: One neuron versus deep learning in aftershock prediction. *Nature* **574**, E1 (2019)
12. Traut, N., Heuer, K., Lemaître, G., Beggiato, A., Germanaud, D., Elmaleh, M., Bethegnies, A., Bonnasse-Gahot, L., Cai, W., Chambon, S., et al.: Insights from an autism imaging biomarker challenge: promises and threats to biomarker discovery. *NeuroImage* **255**, 119171 (2022)
13. Andrzejak, R.G., Lehnertz, K., Mormann, F., Rieke, C., David, P., Elger, C.E.: Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state. *Physical Review E* **64**, 061907 (2001)
14. Subasi, A., Ismail Gursoy, M.: EEG signal classification using PCA, ICA, LDA and support vector machines. *Expert Systems with Applications* **37** (2010)
15. Fulcher, B.D., Little, M.A., Jones, N.S.: Highly comparative time-series analysis: The empirical structure of time series and their methods. *Journal of The Royal Society Interface* **10**, 20130048 (2013)
16. Fulcher, B.D.: *Feature-Based Time-Series Analysis*. In: *Feature Engineering for Machine Learning and Data Analytics*. CRC Press (2018)
17. Lubba, C.H., Sethi, S.S., Knaute, P., Schultz, S.R., Fulcher, B.D., Jones, N.S.: Catch22: CAnonical Time-series CHaracteristics. *Data Mining and Knowledge Discovery* **33**, 1821 (2019)
18. Henderson, T., Bryant, A.G.: *hendersontrend/mean-var-ts-classify: v0.1* (2023). <https://doi.org/10.5281/zenodo.7874699>
19. Nadeau, C., Bengio, Y.: Inference for the Generalization Error. *Machine Learning* **52**, 239 (2003)
20. Henderson, T.: *correctR: Corrected Test Statistics for Comparing Machine Learning Models on Correlated Samples* (2022), <https://CRAN.R-project.org/package=correctR>, r package version 0.1.2
21. Poldrack, R.A., Congdon, E., Triplett, W., Gorgolewski, K., Karlsgodt, K., Mumford, J., Sabb, F., Freimer, N., London, E., Cannon, T., et al.: A phenome-wide examination of neural and cognitive function. *Scientific Data* **3**, 1 (2016)

22. Aquino, K.M., Fulcher, B.D., Parkes, L., Sabaroedin, K., Fornito, A.: Identifying and removing widespread signal deflections from fmri data: Rethinking the global signal regression problem. *Neuroimage* **212**, 116614 (2020)
23. Desikan, R.S., Ségonne, F., Fischl, B., Quinn, B.T., Dickerson, B.C., Blacker, D., Buckner, R.L., Dale, A.M., Maguire, R.P., Hyman, B.T., et al.: An automated labeling system for subdividing the human cerebral cortex on MRI scans into gyral based regions of interest. *Neuroimage* **31**, 968 (2006)
24. Fischl, B., Salat, D.H., Busa, E., Albert, M., Dieterich, M., Haselgrove, C., Van Der Kouwe, A., Killiany, R., Kennedy, D., Klaveness, S., et al.: Whole brain segmentation: automated labeling of neuroanatomical structures in the human brain. *Neuron* **33**, 341 (2002)
25. Henderson, T., Fulcher, B.D.: Feature-Based Time-Series Analysis in R using the theft Package (Aug 2022). <https://doi.org/10.48550/arXiv.2208.06146>
26. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., Duchesnay, E.: Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* **12**, 2825 (2011)
27. Leek, J.T., Storey, J.D.: The joint null criterion for multiple hypothesis tests. *Statistical Applications in Genetics and Molecular Biology* **10** (2011)
28. Ratanamahatana, C.A., Keogh, E.: Everything you know about dynamic time warping is wrong. In: *Third Workshop on Mining Temporal and Sequential Data*. Citeseer (2004)
29. Wu, R., Der, A., Keogh, E.J.: When is Early Classification of Time Series Meaningful? *IEEE Transactions on Knowledge and Data Engineering* **35**(3), 3253–3260 (2023). <https://doi.org/10.1109/TKDE.2021.3108580>
30. Quaak, M., van de Mortel, L., Thomas, R.M., van Wingen, G.: Deep learning applications for the classification of psychiatric disorders using neuroimaging data: systematic review and meta-analysis. *NeuroImage: Clinical* **30**, 102584 (2021)
31. Zeng, L.L., Wang, H., Hu, P., Yang, B., Pu, W., Shen, H., Chen, X., Liu, Z., Yin, H., Tan, Q., et al.: Multi-site diagnostic classification of schizophrenia using discriminant deep learning with functional connectivity mri. *EBioMedicine* **30**, 74 (2018)
32. de Filippis, R., Carbone, E.A., Gaetano, R., Bruni, A., Pugliese, V., Segura-Garcia, C., De Fazio, P.: Machine learning techniques in a structural and functional mri diagnostic approach in schizophrenia: a systematic review. *Neuropsychiatric disease and treatment* p. 1605 (2019)
33. Fulcher, B.D., Jones, N.S.: Hctsa: A Computational Framework for Automated Time-Series Phenotyping Using Massive Feature Extraction. *Cell Systems* **5**, 527 (2017)
34. Aquino, K.M., Fulcher, B., Oldham, S., Parkes, L., Gollo, L., Deco, G., Fornito, A.: On the intersection between data quality and dynamical modelling of large-scale fMRI signals. *NeuroImage* **256**, 119051 (2022)