An application of machine learning techniques to galaxy cluster mass estimation using the MACSIS simulations

DOI:

[10.1093/mnras/stz039](https://doi.org/10.1093/mnras/stz039)

Document Version

Accepted author manuscript

[Link to publication record in Manchester Research Explorer](#)

Citation for published version (APA):

Armitage, T., Kay, S., & Barnes, D. (2019). An application of machine learning techniques to galaxy cluster mass estimation using the MACSIS simulations. *Monthly Notices of the Royal Astronomical Society*, 484(2), 1526-1537. <https://doi.org/10.1093/mnras/stz039>

Published in:

Monthly Notices of the Royal Astronomical Society

Citing this paper

Please note that where the full-text provided on Manchester Research Explorer is the Author Accepted Manuscript or Proof version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version.

General rights

Copyright and moral rights for the publications made accessible in the Research Explorer are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Takedown policy

If you believe that this document breaches copyright please refer to the University of Manchester's Takedown Procedures [<http://man.ac.uk/04Y6Bo>] or contact uml.scholarlycommunications@manchester.ac.uk providing relevant details, so we can investigate your claim.

An application of machine learning techniques to galaxy cluster mass estimation using the MACSIS simulations

Thomas J. Armitage,^{1*} Scott T. Kay¹ and David J. Barnes²

¹*Jodrell Bank Centre for Astrophysics, School of Physics and Astronomy, The University of Manchester, Manchester M13 9PL, UK*

²*Department of Physics, Kavli Institute for Astrophysics and Space Research, Massachusetts Institute of Technology, Cambridge, MA 02139, USA*

Accepted XXX. Received YYY; in original form ZZZ

ABSTRACT

Machine learning (ML) techniques, in particular supervised regression algorithms, are a promising new way to use multiple observables to predict a cluster’s mass or other key features. To investigate this approach we use the MACSIS sample of simulated hydrodynamical galaxy clusters to train a variety of ML models, mimicking different datasets. We find that compared to predicting the cluster mass from the σ – M relation, the scatter in the predicted-to-true mass ratio can be reduced by a factor of 4, from 0.130 ± 0.004 dex (≈ 35 per cent) to 0.031 ± 0.001 dex (≈ 7 per cent) when using the same, interloper contaminated (out to $5r_{200c}$), spectroscopic galaxy sample. Interestingly, omitting line-of-sight galaxy velocities from the training set has no effect on the scatter when the galaxies are taken from within r_{200c} . We also train ML models to reproduce estimated masses derived from mock X-ray and weak lensing analyses. While the weak lensing masses can be recovered with a similar scatter to that when training on the true mass, the hydrostatic mass suffers from significantly higher scatter of ≈ 0.13 dex (≈ 35 per cent). Training models using dark matter only simulations does not significantly increase the scatter in predicted cluster mass compared to training on simulated clusters with hydrodynamics. In summary, we find ML techniques to offer a powerful method to predict masses for large samples of clusters, a vital requirement for cosmological analysis with future surveys.

Key words: galaxies: clusters: general - galaxies: kinematics and dynamics - methods: general: numerical

1 INTRODUCTION

Galaxy clusters are the largest gravitationally bound structures in the Universe, tracing the high mass tail of the halo mass function, which makes them a valuable cosmological probe (see Allen et al. 2011; Kravtsov & Borgani 2012; Weinberg et al. 2013; Mantz et al. 2015). A key requirement for clusters to be used as a cosmological probe is a robust way to determine their masses. With impending surveys from the likes of eBOSS, DESI, eROSITA, SPT-3G, ActPol, and *Euclid*, the number of clusters with high quality data is set to increase substantially, creating the need for an automated way to estimate their masses. To date, significant work has been put into developing physically motivated methods to determine the mass of a cluster. These methods can be broadly split into several classes: X-ray, gravitational lensing, Sunyaev-Zel’dovich, and dynamical analysis of cluster galaxies.

Using data from X-ray observations of the diffuse intra-cluster medium (ICM) and assuming hydrostatic equilibrium, one is able to recover a measure of the cluster mass. However, sources of non-thermal pressure from active galactic nuclei (AGN), substructure, turbulence, magnetic fields, and cosmic rays introduce bias, estimated to be 10–40 per cent (Lau et al. 2009; Rasia et al. 2012; Nelson et al. 2014; Henson et al. 2017). A weak lensing analysis is likewise affected by cluster triaxiality and projection effects, but with a smaller bias of ~ 5 per cent (Okabe et al. 2010; Mahdavi et al. 2013; Hoekstra et al. 2015; Kettula et al. 2015; Oguri & Hamana 2011; Becker & Kravtsov 2011; Bahé et al. 2012; Henson et al. 2017). Dynamical mass estimates use galaxies as tracers of the underlying dark matter (DM) distribution. To what extent galaxies are fair tracers of the velocity distribution is disagreed upon, with the galaxy velocity bias being within ± 10 per cent of the DM dispersion (e.g. Munari et al. 2013; Elahi et al. 2017; Armitage et al. 2018b). Old et al. (2013) and Old et al. (2014) performed a comparison of different dynamical mass estimators, finding that all

* E-mail: thomas.armitage-3@postgrad.manchester.ac.uk

dynamical methods suffer from significant scatter, typically ~ 0.3 dex (i.e. a factor of two) when including projection effects and interloper galaxies.

Clusters are fundamentally complex systems, and many methods require strong assumptions to be made about the state of the cluster, such as spherical symmetry and virial or hydrostatic equilibrium. In addition, much of the information in the data is unused; for instance, in the case of the velocity dispersion - mass relation, the only information used is the tracer velocity dispersion. One would ideally build a model using all available information from the velocity, position and brightness distributions of the tracer population. Machine learning (ML) techniques offer a way of easily incorporating a set of cluster properties, referred to as *features* in ML, into a model which once trained, can then be used to predict the masses of clusters. The caveat is that physically understanding the model can often be difficult. The first application of ML techniques to the problem of cluster mass estimation was performed by Ntampaka et al. (2015) and Ntampaka et al. (2016), using the dark matter only (DMO) Multidark simulation. Ntampaka et al. (2015) use a support distribution machine (SDM, Sutherland et al. 2012) to characterise four feature distributions. Comparing to the velocity dispersion - mass relation, Ntampaka et al. (2015) found the ML implementation improves the scatter in the predicted mass by a factor of two. When the cluster galaxy samples are contaminated with interloper galaxies, the ML method still performed strongly, with less scatter in the contaminated sample for the ML model than for the uncontaminated $\sigma - M$ relation (Ntampaka et al. 2016). Ntampaka et al. (2018) has also explored the use of Convolutional Neural Networks (CNNs), often used in image processing, to predict cluster masses from mock X-ray observations of clusters. They find that using CNNs reduces the scatter in the predicted mass to 12 per cent, down from 15 – 18 per cent using traditional X-ray luminosity based methods.

In this work, we incorporate the impact of astrophysical processes by using the hydrodynamically simulated MACSIS (Barnes et al. 2017a) sample of 390 galaxy clusters, to develop a pipeline to predict cluster masses from galaxy data. The aim of this work is to show how one could apply ML techniques in an observational setting. This includes incorporating projection effects and training the ML models on imperfect masses derived from both X-ray and weak lensing analysis as well as quantifying the reduction in scatter when including additional data, such as the X-ray temperature and Sunyaev-Zel'dovich signal. We also test the impact of hydrodynamics by training ML models on both DMO and hydrodynamical simulations. The paper is organised as follows. Section 2 details the MACSIS clusters used in this analysis, while Section 3 outlines the analysis pipeline, including feature extraction. We compare several different ML algorithms in Section 4 to see which recover the mass of the clusters with the least scatter. In Section 5 we present our main results, comparing different cluster feature sets and target masses. We also discuss the systematics introduced when training using DMO simulations and applying the model to hydrodynamical simulations. Finally we summarise and discuss our findings, explicitly stating the caveats of our work, in Section 6.

2 SIMULATION DATASETS

The MACSIS sample consists of 390 galaxy clusters, with a mass range of $10^{14} < M_{200c}^1 / M_\odot < 4.75 \times 10^{15}$ at $z = 0$. The sample is described in detail in Barnes et al. (2017a); here, we provide a brief overview. The MACSIS clusters were selected from a large 3.2 Gpc DMO parent simulation. Clusters with a friends-of-friends mass between $10^{15} < M_{\text{FoF}} / M_\odot < 10^{16}$ at $z = 0$ were placed into logarithmic mass bins of width $\Delta \log_{10} M_{\text{FoF}} = 0.2$. The three lowest mass bins contained over 100 clusters each and so each bin was further subdivided into 10 logarithmically spaced sub-bins and 10 clusters randomly selected from each. For $M_{\text{FoF}} > 10^{15.6} M_\odot$ there were only 90 clusters, all of which were selected for the MACSIS sample.

The selected haloes were then re-simulated using the zoom technique (Katz & White 1993; Tormen et al. 1997), where the cluster region is simulated at a higher resolution than the parent box, with a lower resolution region beyond which accounts for the tidal forces from the surrounding large scale-structure. The MACSIS clusters were simulated using the *Planck* cosmology (Planck Collaboration et al. 2014a): with $\Omega_m = 0.3070$, $\Omega_b = 0.04825$, $\Omega_\Lambda = 0.6914$, $\sigma_8 = 0.8288$, $n_s = 0.9611$ and $h = 0.6777$. Two zoom simulations of each cluster were created: a DMO simulation, MACSIS-DMO, with a particle mass of $5.2 \times 10^9 M_\odot$, and a hydrodynamical simulation, MACSIS-GAS. For the latter, the initial gas and dark matter particle masses were $8 \times 10^8 h^{-1} M_\odot$ and $4.4 \times 10^9 h^{-1} M_\odot$ respectively. The Plummer equivalent gravitational softening length was fixed to $4h^{-1}$ kpc in physical units between $0 < z < 3$ and $16h^{-1}$ kpc in comoving coordinates at higher redshift.

The code and sub-grid physics model used in the MACSIS-GAS sample is the same as the BAHAMAS simulation (McCarthy et al. 2017). The simulations were performed using the smoothed particle hydrodynamics code GADGET-3, last publicly discussed in Springel (2005). GADGET-3 has been adapted to include the OWLS sub-grid physics model (Schaye et al. 2010), which incorporates the effects of radiative cooling (Wiersma et al. 2009), star formation (Schaye & Dalla Vecchia 2008), and feedback from supernovae (Dalla Vecchia & Schaye 2008) and AGN (Booth & Schaye 2009). The feedback calibration for BAHAMAS was done to match the observed gas mass fraction of groups and clusters (Vikhlinin et al. 2006; Maughan et al. 2008; Sun et al. 2009; Pratt et al. 2009; Lin et al. 2012) and the global stellar mass mass function (Li & White 2009; Baldry et al. 2012; Bernardi et al. 2013), as described in McCarthy et al. (2017). Barnes et al. (2017a) further showed that the MACSIS-GAS clusters reproduce the observed mass dependence of the hot gas mass, X-ray luminosity and Sunyaev-Zel'dovich signal at redshift $z = 0$.

In this paper we define any self-bound object as identified by SUBFIND (Springel et al. 2001; Dolag et al. 2009) with a stellar mass greater than $10^{10} M_\odot$ as a galaxy. Henson et al. (2017) found that due to the initial MACSIS cluster selection being based on M_{FoF} , the lower mass clusters in the sample had systematically low concentrations and high

¹ We define M_{200c} as the mass enclosed within a sphere of radius r_{200c} whose mean density is 200 times the critical density of the Universe.

spin parameters compared to the rest of the sample. To ensure that we do not use unrepresentative low-mass clusters, we perform the same mass cut as Henson et al. (2017), removing clusters below $M_{200c} = 10^{14.5} h^{-1} M_{\odot}$. This reduces the MACSIS sample to 357 clusters, with a mass range $14.6 \leq \log_{10} M_{200c}/M_{\odot} \leq 15.7$ at redshift $z = 0.25$, which is the sole redshift considered in our analysis.

3 MACHINE LEARNING METHODS

Currently, scaling relations are widely used to cheaply obtain cluster mass estimates (e.g. Arnaud et al. 2007; Evrard et al. 2008; Vikhlinin et al. 2009; Pratt et al. 2009; Zhang et al. 2011; Planck Collaboration et al. 2014b; Sereno & Ettori 2015; Mantz et al. 2016). These scaling relationships must either be calibrated using simulations or by using a high quality dataset where more advanced, and data heavy, mass estimation techniques can be applied (e.g. using weak lensing data). Likewise, supervised machine learning techniques require a calibration (training) dataset where the true value of a variable (feature) is known. A key advantage of ML methods, however, is that they are easily generalised to use many different features together, potentially making better use of all the available information.

Supervised machine learning techniques can be broadly classified into two types: classification and regression. Classifiers aim to place objects into discrete groups, e.g. whether a galaxy cluster is relaxed or not. Regression algorithms are for continuous distributions, and so are more suited for cluster mass estimation, which is the main aim of this study. One of the simplest regression models is the familiar least-squares linear regression, which can be generalised to fit n features to predict a desired observable. In this section we shall discuss how we extracted useful features from our cluster sample and the general principles of regression models before comparing the performance of different algorithms in Section 4².

3.1 Feature extraction and selection

Before we can apply the regression models we must first extract features from the MACSIS clusters. A feature must be a scalar value for each cluster, such as the number of galaxies present inside a given aperture, or the cluster velocity dispersion. In this work we initially focus on features that can be obtained from an optical survey of clusters, i.e. projected galaxy positions, line-of-sight velocities and stellar masses. We note that all values are taken directly from the simulation, i.e. we do not include measurement errors in this work. We project all cluster samples along a cylinder of length $10r_{200c}$ within a radius of either r_{200c} or 1.5 Mpc , to contaminate the sample with interloper galaxies. The first aperture, by definition, scales with cluster mass ($r_{200c} \propto M_{200c}^{1/3}$), while the second is chosen to assess the effect of choosing a fixed physical aperture. For each of these cluster projections we

Figure 1. Coefficient of determination against number of features in the training sample. Initially all features are used to train a regression model, the feature contributing the least information is then discarded. The feature discarded is shown at the top of the plot (with the most important features on the left). The process is repeated until only one feature remains.

calculate the mean, standard deviation, skewness and kurtosis of said galaxy properties. We also consider the total number of galaxies and the ellipticity (ratio of semi-minor to semi-major axes), e , using the method described in Bett et al. (2007). The position information is encoded as the projected cluster-centric radius of each galaxy. The features chosen for this work are by no means the only ones possible; one could incorporate relaxation criteria, or morphological indicators for example, as well as information from other datasets, such as weak lensing shear or X-ray surface brightness maps. In this paper, we focus on relatively direct observables, with more complex features left for future work. For conciseness we define the PHOT and SPEC feature sets: PHOT contains only projected positional galaxy information and their respective stellar masses, while SPEC includes the PHOT features and line-of-sight velocity information. Later, we also incorporate the X-ray temperature, X-ray (bolometric) luminosity and SZ flux as additional features. We summarise the quantities in each feature set in Table 1.

From Table 1 R , m and v are the projected radial distance, stellar mass and line-of-sight velocity of each galaxy in a cluster. We then take the mean, standard deviation, skewness and kurtosis of each distribution, say x , labelling x_{mean} , x_{std} , x_{skew} and x_{kurt} , respectively. Details for how the X-ray surface temperature, T_{500c} , and luminosity, $L_{X,500c}$ are calculated can be found in Barnes et al. (2017a). The ellipticity, e , is calculated following Bett et al. (2007), computing the mass distribution tensor and taking the ratios of the eigenvalues.

Having generated a number of features from the galaxy population of each cluster we must now determine which features contain the most information relevant to the target feature, in this case $\log_{10}(M_{200c})$. Including features that are uncorrelated with the target output can act to wash out the information contained in other features, as well as increase the computation time unnecessarily. One way to select

² We provide a basic template of the ML implementation used in this work, available at https://github.com/TomArmitage/ML_Template

Table 1. Summary of the quantities (features) in each feature set used throughout this paper. R , m and v are the projected radial distance, stellar mass and line-of-sight velocity of each galaxy in a cluster. We take the mean, standard deviation, skewness and kurtosis of these distributions as features for each cluster. The ratio e of the semi-minor and semi-major axes is calculated following Bett et al. (2007). All features in X are the core-excised variants found in Barnes et al. (2017a).

Feature set	Included Features
PHOT:	R_{mean} , std, skew, kurt, m_{mean} , std, skew, kurt, N_{gal} , e
SPEC:	PHOT + v_{mean} , std, skew, kurt
X:	T_{500c} , $L_{X,500c}$
SZ:	$Y_{\text{SZ},5r_{500c}}$

the best features is through Recursive Feature Elimination (RFE, Guyon et al. 2002). RFE is a brute force technique, that given a regression (or classifier) algorithm, trains using all features initially and then eliminates the least important feature according to some goodness of fit metric, proceeding recursively until a set number of features remain. The scoring metric used in this work is the coefficient of determination, r^2 , defined as

$$r^2(y, \hat{y}) = 1 - \frac{\sum_{i=0}^{n_{\text{samples}}-1} (y_i - \hat{y}_i)^2}{\sum_{i=0}^{n_{\text{samples}}-1} (y_i - \bar{y})^2} \quad (1)$$

where $\hat{y}_i$ is the predicted output for the i^{th} cluster and y is the set of true outputs, in this case cluster masses, with a mean value $\bar{y}$. n_{samples} is the total number of samples (galaxy clusters) used to train the ML model. The optimal r^2 score is 1, indicating a perfect fit to the data. Fig. 1 shows an example r^2 curve generated using the MACSIS sample, using galaxies within an aperture of r_{200c} and the SPEC feature set. One can see that beyond 7 features the curve is very flat, indicating that those features are contributing little additional information, while the first few features greatly improve the goodness of fit. The maximum value of r^2 is obtained with 10 features with the inclusion of the velocity dispersion, and so in the pipeline we develop only those features that would be used to train the regression model (the features eliminated are v_{kurt} , v_{mean} , v_{skew} and e). Note that the only velocity feature present after RFE is the velocity dispersion, which we explore further in Section 5. However, RFE is limited in a number of ways. Namely the initial weights on the importance of different features can be unreliable due to the nature of the problem, so the exact order of the features is unstable, however the overall trends are meaningful. It also ignores any non-linear relationships between features. The r^2 profiles do not qualitatively change with different combinations of features, e.g. PHOT+X, SPEC+YZ; as may be expected, there is a lot of redundant information in our initial feature set.

Another way to reduce the number of features is to perform a principal component analysis (PCA). Put briefly, PCA takes the multi-dimensional feature space (where each galaxy cluster has a position set [v_{std} , m_{skew} , N_{gal} , ...]) and re-projects it so that each new axis is independent of the others, i.e. the covariance matrix of the new features is diagonalised. One can then remove features based on their respective variance, allowing one to retain the maximum information after

removing a set number of features. A significant disadvantage with PCA is that each new feature is a linear combination of all the original features, making it difficult to determine what type of data is most useful. In practice, we found little difference in the estimated mass when pre-processing features using either RFE or PCA. Due to additional abstraction with PCA we opt to use RFE, selecting the feature set with the highest r^2 score, throughout the rest of the paper unless explicitly stated otherwise. We discuss which features contain the most useful information in Section 5.

A second important preprocessing step is standardisation of the data, as many ML algorithms assume the features have approximately zero mean and a standard deviation of one. To ensure our data are robust to outliers we instead subtract the median of each feature and divide by the 1σ (16th/84th) percentile range. For properties whose values span orders of magnitudes we take the base 10 logarithm before standardising (v_{std} , N_{gal} , $Y_{\text{SZ},5r_{500c}}$, T_{500c} and $L_{X,500c}$).

3.2 Model training and testing

The next stage of the analysis is to use the feature set, after RFE, to train a model to predict some desired quantity or target feature, in this case cluster mass. To train a model we require both a training and testing sample. The target feature is known for the training sample and is used to train the model. The model is then used to predict the same target feature in the test sample. It is vitally important that the training sample is a fair representation of the underlying distribution. In cases with only a few hundred samples one would ideally like to obtain estimated masses for all clusters. This can be achieved using k -fold cross validation. k -fold cross validation randomly splits the available data into k subsamples; $k - 1$ subsamples are then used to train a model which is then tested on the remaining subsample. The subsamples are then rotated with all galaxy clusters tested exactly once. In this work we take the commonly used value of $k = 10$, i.e. a 90-10 split (Ball & Brunner 2010). Once we have a predicted mass for each cluster we can then compute properties such as the scatter and bias between the predicted mass and the true mass. To prevent any duplication of data in the training and testing samples, we perform the k -fold cross validation completely separately for when we project clusters along the x , y and z axes i.e. three separate sets of 10-fold cross validation. Note that there should be negligible bias using this method, as long as the training set is a fair sample of the underlying distribution (we checked this and found this to be a good assumption in practice). We therefore focus on comparing results for the logarithmic scatter, δ , in $\log_{10}(M_{\text{pred}}/M_{\text{true}})$, where M_{pred} is the predicted cluster mass from the ML model and M_{true} is the true mass of the cluster.

We now discuss the implementation of the ML algorithms used in this work. There are a wide variety of regression algorithms available through software packages such as the Python based SCIKIT-LEARN used here (Pedregosa et al. 2011). We shall now describe two of the simpler, yet most relevant, procedures that are widely used throughout this work. For a brief overview of the other methods we considered and their performance, see Section 4.

The simplest regression algorithm is ordinary linear regression (OLR), which is used ubiquitously. To summarise,

if one has a set of n independent observables of i systems, $\mathbf{X} = \{\mathbf{x}_0^T, \mathbf{x}_1^T, \dots, \mathbf{x}_n^T\}$, where $\mathbf{x}$ is a vector of length i with corresponding dependent variables $\mathbf{y}$, then one can relate the two with an unknown vector $\boldsymbol{\beta}$. By finding the value $\hat{\boldsymbol{\beta}}$ that minimises a loss function, in this case the sum of least squares

$$\mathcal{L}_{\text{RSS}} = \sum_i (y_i - \mathbf{x}_i^T \boldsymbol{\beta})^2, \quad (2)$$

one tends to obtain a good estimate of the true value of $\boldsymbol{\beta}$. However, if the columns of $\mathbf{X}$ are highly correlated then the above procedure can yield highly unstable estimates of $\boldsymbol{\beta}$. For a simple example of how this can arise, if one has just two features, $\mathbf{x}_0$, and $\mathbf{x}_1 = \mathbf{x}_0 + \delta$, where δ is a small random number, then there will be many values of $\hat{\boldsymbol{\beta}}$ which produce low values of RSS. This can result in over-fitting of the data, as the value of $\hat{\boldsymbol{\beta}}$ becomes very sensitive to the specific training values of $\mathbf{X}$.

A solution to this is to introduce a term which penalises large values of $\boldsymbol{\beta}$. In the case of ridge regression (Hoerl & Kennard 1970) the quantity to minimise is

$$\mathcal{L}_{\text{Ridge}} = \mathcal{L}_{\text{RSS}}(\boldsymbol{\beta}) + \lambda \sum_{j=1}^n \beta_j^2, \quad (3)$$

where the second term prefers values of β_j closer to zero. If one has already standardised the feature set by subtracting the mean and normalising by the standard deviation then this is good assumption. To find $\hat{\boldsymbol{\beta}}$ in the case of ridge regression one computes

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}^T \mathbf{X} + \lambda \mathbf{I})^{-1} \mathbf{X}^T \mathbf{y}, \quad (4)$$

where λ is a hyper parameter that needs to be chosen beforehand. In the case of $\lambda = 0$ then ridge regression becomes identical to OLR and as $\lambda \rightarrow \infty$, $\hat{\boldsymbol{\beta}} \rightarrow \mathbf{0}$. The selection of λ is often made during the training process by testing a range of different values and comparing the residuals of $y - \hat{y}$. In this work we perform this search of λ values via a basic grid search of the parameter space. During training, multiple values of λ are tested on a separate subset of the training data to see which value of λ minimizes the sum of residuals.

4 PERFORMANCE OF ML ALGORITHMS

In this section we compare a selection of regression algorithms to determine the most suitable algorithm for predicting the masses of our clusters. All methods were trained using the SPEC feature set to predict M_{200c} . If any cluster contained fewer than 10 galaxies we rejected it from the sample. Since we had already applied the mass cut of [Henson et al. \(2017\)](#) the latter condition only removed around one cluster for any given projection. We now provide a brief³ overview of each ML algorithm considered:

- OLR - Ordinary Linear Regression. This method was already described in Section 3.2.

³ For a more detailed description of each method and their implementation, see [Pedregosa et al. \(2011\)](#) and the documentation included with SCIKIT-LEARN.

Figure 2. Comparison of the scatter between the predicted and true mass, δ , for the ten different machine learning algorithms, using the SPEC feature set. The solid portion of each bar illustrates the 1σ scatter and the hatched section 2σ . Similar methods are coloured as such, with acronyms defined in the text.

- RR, BRR, KRR - three variants of ridge regression, the basic implementation of which is also described in Section 3.2. The Bayesian ridge regression (BRR) variant incorporates Gaussian priors over the model weights. Kernel ridge regression (KRR) combines ridge regression with the kernel trick, used in SVR below.

- SVR - Support vector regression: The continuous form of support vector machines (SVMs), which classifies data by embedding the features into a higher dimensional space and separating different classes using a hyper-dimensional plane, which can also be thought of as re-projecting the data using some kernel function (this is referred to as the ‘kernel trick’), in this paper we use the radial basis function (rbf) as the kernel. SVR outputs continuous values rather than discrete classes as for SVM.

- DTR - decision tree: A decision tree is created using a set of *if-then-else* rules. The depth of the tree can be adjusted; for DTR5 we use a maximum depth of 5, which can help prevent the model from over-fitting the training data. We also allow the tree to go deep enough that all leaves only contain one sample (DTR).

- GBDT - gradient boosted regressor trees: Gradient boosting involves sequentially adding trees of increasing complexity in order to minimise some loss function.

- **ADA** - AdaBoosted decision trees: AdaBoosted trees are an ensemble of small decision trees, which are weighted based upon their accuracy, with additional trees added to improve poor estimates. The final prediction is a weighted sum of all decision trees (Freund & Schapire 1997).

- **k-NN** - *k* Nearest Neighbours: After computing the *k* nearest neighbours to the desired point in feature space, the average target value of the *k* neighbours is taken as the estimate.

For our main results, we present the scatter in the predicted-to-true mass ratio, δ , which we define as half the 16th – 84th percentile range of

$$\mathcal{R} = \log_{10} (M_{\text{pred}}/M_{\text{true}}), \quad (5)$$

where M_{pred} is the predicted cluster mass and M_{true} is the mass we are attempting to recover, usually M_{200c} . In general, we find the distribution of $\mathcal{R}$ to be approximately Gaussian, and we show the 1σ (68 per cent) and 2σ (95 per cent) percentile scatter in Fig. 2 for each method as the solid and hatched bars respectively. For each model, we predict M_{200c} for each cluster via *k*-fold cross validation. The analysis was repeated for three orthogonal projections of the clusters and the predicted masses of the tested clusters (around 1,000 in total) were used to calculate the scatter, δ in dex.

We can see from Fig. 2 that the best performing models belong in the ridge regression family: RR, BRR, and KRR. However, KRR suffers from a reasonable number of catastrophically-underestimated outliers. These occur at the edges of the cluster mass distribution and are the primary cause of the slightly increased scatter seen in Fig. 2. Due to these outliers we discard KRR. Surprisingly, ordinary linear regression is also one of the best performing methods, even though it is the simplest algorithm tested.

All of these models have a 1σ scatter of less than 0.075 dex with the best having $\delta \sim 0.03$ dex. Compared to the results of Armitage et al. (2018a), who quantified the scatter of traditional dynamical mass estimators using the high resolution hydrodynamical C-EAGLE clusters (Barnes et al. 2017b), such as the caustic and Jeans methods, all of the machine learning models outperform the traditional estimators under comparable conditions which produce an average scatter of 0.1 dex, i.e. a factor of 3 reduction. Furthermore, when using a stellar mass limited sample of $10^9 M_{\odot}$, with perfect knowledge of 3D galaxy positions, velocities and cluster membership Armitage et al. (2018a) finds the scatter is only reduced to 0.07 dex, significantly higher than the 0.03 dex obtained with the projected and contaminated MACSIS sample.

The OLR, RR and BRR methods produce the least scatter with little to distinguish them. Both the RR and BRR methods require hyper-parameters to be tuned, whereas OLR has none. RR has one hyper-parameter compared to four for BRR, which reduces the computational time required to optimise the parameters via a grid-search. We performed all of the analysis in Section 5 using both RR and OLR, finding little difference between the two algorithms, with OLR suffering from slightly more scatter. As a result we only present results obtained using the ridge regression method (RR) from now on.

Figure 3. Top: Velocity dispersion, calculated within r_{200c} , against M_{200c} for all MACSIS clusters along three separate lines-of-sight. The dashed red line shows the best power law fit to eq. (6). Bottom: Same as above, though the velocity dispersion is calculated within a fixed aperture of 1.5 Mpc, rather than r_{200c} .

5 FEATURE SET COMPARISON

Having selected ridge regression as the optimal algorithm to predict the masses of our clusters, we now consider the performance of different feature sets and the effect of baryonic physics. We first discuss the use of the $\sigma - M$ relation, which we use as a baseline comparison, before we consider the performance of different optical feature sets and aperture sizes, namely the difference between a photometric and spectroscopic galaxy sample, when the galaxies are selected inside r_{200c} or within a fixed aperture of 1.5 Mpc. These two choices of aperture are intended to show a ‘best’ case and realistic scenario, respectively. We then add X-ray and SZ properties, before training the models to fit either estimated weak lensing or hydrostatic masses. Finally, we test the sensitivity of our ML model to baryonic physics by training on a DMO simulation and applying that model to hydrodynamic simulations.

5.1 Velocity dispersion - mass baseline model

Throughout this section we shall use the velocity dispersion - mass relation ($\sigma - M$) as a baseline point of comparison. The velocity dispersion, like optical richness (e.g. Yee & Ellingson 2003; Simet et al. 2017), and the Sunyaev-

Table 2. The 1σ percentile scatter, δ , in $\log_{10}(M_{\text{est}}/M_{200c})$ for different feature sets. Table 1 shows the features included in each row. Galaxies were either taken within a cluster aperture of r_{200c} or 1.5 Mpc. For comparison, we provide the scatter when obtaining masses via the $\sigma - M$ relation, also calibrated using the MACSIS clusters. The errors in δ are the 1σ percentile scatter obtained via bootstrap resampling of the clusters.

Data set	$\delta_{r_{200c}}$	$\delta_{1.5 \text{ Mpc}}$
PHOT	0.032 ± 0.001	0.065 ± 0.002
SPEC	0.031 ± 0.001	0.064 ± 0.002
$\sigma - M$	0.130 ± 0.004	0.122 ± 0.003

Zel’dovich (SZ) flux (e.g. Planck Collaboration et al. 2014b; Saliwanchik et al. 2015) is used as a mass proxy (e.g. Zhang et al. 2011; Bocquet et al. 2015; Sereno & Ettori 2015). Like all mass proxies one must calibrate the relation using either numerical simulations or clusters with masses already derived, somewhat analogous to the use of training data in ML. While Evrard et al. (2008) showed that σ scales tightly with cluster mass and is insensitive to the background cosmology, the restriction to line-of-sight velocities has been found to increase the scatter from ~ 13 per cent up to $\sim 30 - 40$ per cent (Saro et al. 2013).

We calculate the velocity dispersion using the ‘gapper’ technique in Beers et al. (1990) and use the whole MACSIS sample to fit a relation of the form

$$\log(\sigma/\text{km s}^{-1}) = A \log(M_{\text{pred}}/M_{\odot}) + B, \quad (6)$$

representing the best possible case for $\sigma - M$, zeroing any bias. Fig. 3 shows the obtained $\sigma - M$ relation with the best fit parameters, taking galaxies within a projected aperture of r_{200c} (top panel) and 1.5 Mpc (bottom panel). Using the fitted relation, we calculate the expected cluster mass, given the velocity dispersion throughout the rest of the paper.

5.2 Dependence on optical features and aperture size

We first compare how constraining the different galaxy features are individually. We train an RR model to predict M_{200c} from a single feature using k -fold cross validation with 10 bins. Fig. 4 shows the scatter in the resulting mass estimates, where the galaxies were selected from inside a projected aperture of r_{200c} . We find that R_{STD} , m_{skew} and N_{gal} , are the most constraining galaxy-based features while e , an indicator of cluster morphology, is one of the least constraining features with a scatter of 0.23 dex, as the shape of a cluster is weakly correlated with mass, i.e. both small and large clusters are equally likely to be extended. We note that if we train on completely random noise we obtain a scatter of ~ 0.24 dex, as the model effectively outputs the mean mass of the training sample, so no feature should ever result in a scatter > 0.24 dex. One caveat is that MACSIS does not include low mass clusters, which may be more likely to be dynamically relaxed and spherical. As might be expected, the velocity dispersion is the single most constraining velocity feature ($\delta \sim 0.143$ dex), however, it is significantly less constraining than R_{STD} (0.075 dex), m_{skew} (0.086 dex), N_{gal} (0.089 dex), R_{mean} (0.094 dex) and m_{kurt} (0.103 dex).

If instead we take galaxies within a fixed aperture of 1.5 Mpc, we obtain the results shown in Fig. 5. Selecting a fixed physical aperture removes the intrinsic mass depen-

Figure 4. Comparison of the scatter, in dex, of each individual feature considered in this work. The scatter is defined as in Fig. 2. The cluster sample was selected within an aperture of r_{200c} . The different colours correspond to the physical feature; green, red and purple denote features derived from the galaxy radial, mass and velocity distribution of each cluster, respectively. The number of galaxies, ellipticity, X-ray features and the SZ flux are shown in brown, blue, yellow and cyan respectively. Features derived from the cluster ICM (galaxies) are above (below) the black dashed line.

dence that is included when using r_{200c} . The most marked change is the increased scatter for R_{mean} and R_{STD} , by a factor of 2 and 3, respectively. Many of the other galaxy-derived features also suffer an increase in scatter, for similar reasons. However, both v_{STD} and m_{mean} improve with the fixed aperture, with their scatter decreasing to $\delta = 0.87$ and $\delta = 0.8$ of the r_{200c} values, respectively. This is likely due to the fraction of interlopers present in the 1.5 Mpc aperture being half that of the r_{200c} sample on average (we note that the mean r_{200c} of the MACSIS clusters is ~ 1.9 Mpc).

We now combine the features into two sets that represent different observational sets, PHOT and SPEC, using the two different aperture sizes (see Table 1). We remind the reader that the features were first filtered using RFE and then used to train ridge regression models using the true M_{200c} cluster mass. Table 2 shows the scatter, δ , for

Figure 5. Comparison of the scatter, in dex, of each individual feature considered in this work measured within a fixed aperture of 1.5 Mpc for each cluster, instead of r_{200c} as in Fig. 4.

each feature set. We see that there is remarkably little difference between PHOT and SPEC results, although both are significantly better than using the $\sigma - M$ relation, which is included in Table 2 for comparison. (We note that we fit the $\sigma - M$ relation using the MACSIS galaxies, and then use that relation to predict the mass of MACSIS clusters from their velocity dispersion.)

Interestingly the fixed 1.5 Mpc aperture results in an increased scatter (by ~ 0.03 dex) for the PHOT and SPEC set, but decreases the scatter by ~ 0.01 dex for the $\sigma - M$ relation. The decrease in scatter for the $\sigma - M$ relation is far less statistically significant, at around the 2σ level. The increase in scatter for the machine learning methods comes from the loss of information from galaxies beyond 1.5 Mpc, particularly for the more massive systems. In addition to this, r_{200c} implicitly contains information about M_{200c} , which would also reduce the scatter.

The main implication of Table 2 is that, given an ideal training sample, the addition of galaxy line-of-sight velocities does not improve the accuracy over a purely photometric data set, i.e. one can obtain the same degree of scatter in the predicted mass with a photometric sample as a spectroscopic sample. Ntampaka et al. (2016) found a similar result

using support distribution machines (SDM), where the scatter using the projected radial distribution of the galaxies was less than using the line-of-sight velocity distribution for their high-mass ($M > 10^{15} M_{\odot}$) cluster sample, and the same for their overall sample. One notable caveat to this statement is that our training sample does not contain any interlopers beyond $5r_{200c}$ of the cluster, due to the limited size of the high resolution volume in the simulation. Spectroscopic data may still be required to accurately remove more distant interloper galaxies, however Rozo et al. (2015) found when using redMaPPer that they could achieve cluster membership probabilities with photometric data that matched spectroscopic estimates at the per cent level.

5.3 Inclusion of X-ray and SZ features

In addition to features derived from the galaxy population, one can include features from other observations to predict M_{200c} . We include the additional X-ray and SZ features in Figs. 4 and 5, shown above the black dashed line. The Y_{SZ} flux inside $5r_{500c}$ is the single most constraining feature, better than the galaxy-based features. The X-ray derived features, $L_{X, 500c}$ and T_{500c} , are taken within an observing aperture of r_{500c} , which will contain, implicitly, information about M_{500c} .

Fig. 6 shows the scatter when each feature is added in order of the amount of information contributed, according to RFE, using an aperture of r_{200c} . Starting from the bottom of the plot, each feature is added sequentially and the regression model is trained using it and all the features below. The black dashed line marks the point where there is no further improvement to the scatter (0.03 dex) through the addition of features. At the top of the plot, above the solid black line, we show the scatter for the velocity dispersion and number of galaxies when used individually (i.e. the same values as shown in Fig. 4) for comparison.

We stress that the exact order of features in Fig. 6 is somewhat unstable due to the limitations of RFE as stated previously, particularly above the black dashed line. However, we can infer information from the general trends. As expected, the SZ flux and the projected radial distribution of galaxies both contribute a significant amount of information, as well as the higher order moments of the galaxy stellar mass distribution and the total number of galaxies present (richness). We can also see that the majority of the features are unnecessary; in the case where we use a fixed aperture, the number of rejected features decreases, though the majority only yield marginal improvements in the scatter.

As previously stated, there are some surprising results, particularly the galaxy velocity distribution not contributing much additional information. This leads to the question: given photometric observations of a galaxy cluster, how much is the scatter improved by including additional data/features? Fig. 7 shows how different combinations of observational feature sets affect the scatter in the predicted mass. SPEC and PHOT are as defined previously, SZ includes $Y_{SZ, 5r500c}$, while X includes $L_{X, 500c}$ and T_{500c} as stated in Table 1. We use a shorthand notation to represent combinations of these feature sets, e.g. PHOT X includes the features in both PHOT and X together. Each feature set is optimised using RFE before training on M_{200c} .

Figure 6. Comparison of the scatter, in dex, using different combinations of features. The ranges shown are the same as in Fig. 2 and the colours match those in Fig. 4. The order is determined by recursive feature elimination, where each row shows the scatter obtained from training a ridge regression model using all features below and at that point. The dashed horizontal black line shows the point where the minimum scatter (0.031 dex) is obtained; above the line the scatter is either the same or greater. Above the solid horizontal line, we show the scatter for two common observables: the number of galaxies in the cluster and their velocity dispersion, which were computed individually as a point of comparison. The cluster sample was selected within an aperture of r_{200c} .

We see there is very little absolute difference between any combination of features. In general, the information loss from any singular feature is compensated for by other features contributing comparatively more unique information compared to when the removed feature was included. Discarding measures of the galaxy distribution does however lead to increased scatter. For example, removing the skew and kurtosis increases the scatter by 68 per cent, from 0.031 dex to 0.052 dex, with an error obtained via bootstrap resampling the data 1000 times of $\sim \pm 0.001$ dex for each.

Figure 7. Comparison of the bias and scatter, in dex, using different feature sets, with galaxy derived features obtained within an aperture of r_{200c} . The ranges shown are the same as in Fig. 2. PHOT refers to a photometric dataset of galaxy positions and stellar mass; SPEC additionally includes line-of-sight velocity information; SZ marks the inclusion of the SZ flux within $5r_{500c}$; X denotes the inclusion of ICM derived properties, the bolometric luminosity and the surface temperature

5.4 Training on mock X-ray and weak lensing masses

A significant issue with the ML implementation as presented thus far is that it relies on knowing the true masses for a training set of clusters. This problem could be alleviated by training the model on a set of simulated clusters and then applying it to real data. Alternatively, one could train on a feature that is directly observable, such as the velocity dispersion (Caldwell et al. 2016), or on a set of clusters with a known mass, estimated via some other means. The first option could suffer from mis-calibrated subgrid models, cluster selection effects or instrumental limitations. At the very least for a new method, such as this, there needs to be some independent validation. In this spirit, we test the performance of our regression models when trained to reproduce masses obtained from mock weak lensing and X-ray observations. The driving philosophy behind this is that one could train on a high quality set of data and then apply that model to a larger, poorer quality, dataset, fulfilling a

Table 3. The 1σ percentile scatter, δ , for different feature sets. Table 1 summarises the quantities in each feature set. The uncertainties are the 1σ percentiles computed from 1000 bootstrap samples of the clusters. The left and right hand side of the table shows the scatter when taking galaxies within an aperture of r_{200c} and 1.5 Mpc respectively. The three columns in each are for the different target features, the true M_{200c} , M_{200c} inferred from weak lensing, $M_{200c,WL}$, and the hydrostatic mass inside r_{500c} obtained from X-ray spectra $M_{500c,X}$. For comparison we include the scatter obtained when finding the mass using the $\sigma - M$ relation, calibrated using the MACSIS clusters on the relevant target mass.

Data set	$\delta_{r_{200c}}$			$\delta_{1.5 \text{ Mpc}}$		
	M_{200c}	$M_{200c,WL}$	$M_{500c,X}$	M_{200c}	$M_{200c,WL}$	$M_{500c,X}$
$\sigma - M$	0.130 ± 0.004	0.117 ± 0.004	0.182 ± 0.006	0.122 ± 0.003	0.117 ± 0.004	0.177 ± 0.006
PHOT	0.032 ± 0.001	0.071 ± 0.002	0.127 ± 0.004	0.065 ± 0.002	0.072 ± 0.002	0.133 ± 0.004
SPEC	0.031 ± 0.001	0.064 ± 0.002	0.130 ± 0.004	0.064 ± 0.002	0.066 ± 0.002	0.137 ± 0.004
PHOT X	0.033 ± 0.001	0.067 ± 0.002	0.088 ± 0.003	0.065 ± 0.002	0.065 ± 0.002	0.090 ± 0.003
SPEC X	0.032 ± 0.001	0.061 ± 0.002	0.087 ± 0.003	0.064 ± 0.002	0.060 ± 0.002	0.090 ± 0.003
PHOT SZ	0.031 ± 0.001	0.069 ± 0.003	0.123 ± 0.005	0.052 ± 0.002	0.070 ± 0.002	0.131 ± 0.005
SPEC SZ	0.031 ± 0.001	0.064 ± 0.002	0.128 ± 0.005	0.051 ± 0.002	0.065 ± 0.002	0.133 ± 0.006
PHOT X SZ	0.031 ± 0.001	0.063 ± 0.002	0.086 ± 0.003	0.051 ± 0.002	0.066 ± 0.002	0.089 ± 0.004
SPEC X SZ	0.031 ± 0.001	0.060 ± 0.002	0.087 ± 0.003	0.052 ± 0.001	0.062 ± 0.002	0.089 ± 0.003

role similar to current scaling relations. While we have only trained to reproduce X-ray and weak lensing masses in this work, in principle one could train a model to reproduce any mass proxy, such as L_X or SZ signal.

Our results are summarised in Table 3, which shows the scatter for a variety of apertures, feature sets and target masses. Models are trained using either M_{200c} , the estimated value of M_{500c} from X-ray analysis ($M_{500c,X}$) or M_{200c} from weak lensing ($M_{200c,WL}$). The $M_{500c,X}$ masses are from the work of Barnes et al. (2017a) and are computed by fitting the density and temperature profiles, parametrised by Vikhlinin et al. (2006), and used to produce a hydrostatic mass profile to obtain $M_{500c,X}$. The weak lensing masses are from Henson et al. (2017), who fit the shear profile assuming an NFW mass distribution. Both $M_{500c,X}$ and $M_{200c,WL}$ are obtainable from observational data. This represents the scenario where one has, for example, a high quality X-ray sample of clusters, with overlapping optical data. One can then train a model using the smaller sample, before applying it to the larger optical catalogue. In addition the models were trained only using galaxies within a projected radial distance of 1.5 Mpc from the cluster centre, removing any prior knowledge of the cluster size.

We see a stark difference in the scatter depending on whether the model was trained using the true, X-ray or weak lensing mass. As we discussed above, training on the true value of M_{200c} results in a uniform scatter of ~ 0.03 dex, irrespective of the feature set, within an aperture of r_{200c} . Training on $M_{200c,WL}$ increases the scatter to ~ 0.065 dex. We also see that the addition of spectroscopic data yields a 0.07 dex improvement in the scatter over PHOT. Training on $M_{500c,X}$ results in a significant increase in scatter, with the SPEC feature set resulting in a scatter of 0.13 dex, compared to 0.064 dex or 0.031 dex when training with $M_{200c,WL}$ and M_{200c} respectively. While it is not surprising that the scatters qualitatively increase, a jump of ~ 19 per cent to ~ 35 per cent is rather large. This indicates that, while in principle one could train using any mass estimate or cluster property, it must be chosen with care. Notably, fixing the aperture to 1.5 Mpc increases the scatter by only a small amount when training on $M_{500c,X}$ or $M_{200c,WL}$, but the scatter around the true mass nearly doubles in dex for many cases.

Table 4. The 1σ percentile scatter, δ , for different feature sets. This shows a subset of training masses, inside 1.5 Mpc. The formatting is the same as for Table 3.

Data set	$\delta_{1.5 \text{ Mpc}}$	$\delta_{1.5 \text{ Mpc}}$
	$M_{500c,X}$	M_{500c}
$\sigma - M$	0.177 ± 0.006	0.130 ± 0.004
PHOT	0.133 ± 0.004	0.072 ± 0.002
SPEC	0.137 ± 0.004	0.071 ± 0.002
PHOT X	0.090 ± 0.003	0.056 ± 0.002
SPEC X	0.090 ± 0.003	0.053 ± 0.001
PHOT SZ	0.131 ± 0.005	0.060 ± 0.002
SPEC SZ	0.133 ± 0.006	0.058 ± 0.002
PHOT X SZ	0.089 ± 0.003	0.048 ± 0.002
SPEC X SZ	0.089 ± 0.003	0.048 ± 0.002

Table 4 shows the scatter in dex when training to reproduce $M_{500c,X}$ and M_{500c} with a fixed aperture around each cluster of 1.5 Mpc. For reference the average true value of r_{500c} for the MACSIS clusters is 1.23 Mpc, and 1.93 Mpc for r_{200c} . The only difference between the two columns in Table 4 is the target feature, removing the possibility that the increase in scatter with $M_{500c,X}$ is due to predicting a quantity based on M_{500c} rather M_{200c} . We can see that the scatter improves significantly with the addition of more features. The PHOT feature set has a scatter of 0.072 ± 0.002 dex compared to the SPEC X SZ set, 0.048 ± 0.002 . Interestingly there is never a significant difference between PHOT and SPEC when training on the true value of M_{500c} . While there is an improvement from adding additional features to reproduce $M_{500c,X}$ it is much less significant. The intrinsic scatter in $M_{500c,X}$ appears to be the limiting factor in this case.

5.5 Dependence on baryonic physics

Finally, we now consider the effects of baryonic physics on the model accuracy. We are interested in whether the inclusion of baryonic physics alters the scatter and how robust the ML models are to a change in subgrid physics (i.e. trained on DMO simulations and then applied to hydrodynamical clusters). The MACSIS clusters were simulated both with

Figure 8. Top: Comparison of the predicted mass using a ridge regression model trained on DMO simulations of the MACSIS clusters, which is then tested using the hydrodynamical versions of the clusters. The dashed red line shows the best power law fit to the predicted masses, $\log(M_{\text{pred}}/M_{\odot}) = A \log(M_{200c}/M_{\odot}) + B$. Bottom: Same as above, though the aperture is fixed to 1.5 Mpc, rather than r_{200c} .

and without baryonic physics. To quantify the effect of baryonic physics on our feature set we train a ridge regression model using the MACSIS-DMO clusters and then apply that model to the MACSIS-GAS sample. We also computed ‘GAS to GAS’ and ‘DMO to DMO’ estimates for comparison. Throughout this section we use the the SPEC feature set exclusively, using subhaloes to represent the galaxies. In order to avoid any unrealistic correlations we train the model using 90 per cent of the DMO clusters and test on the remaining 10 per cent of clusters using the MACSIS-GAS sample. This way the DMO variant of a cluster will never be in the training sample when the MACSIS-GAS version is used to test the model. Similarly, we project the clusters along three orthogonal axes, but perform the training and testing analysis entirely separately for each projection. We only combine the projected samples once all predicted masses have been obtained.

Table 5 shows the scatter in dex for the six cases, note that the GAS-GAS values differ from those in Table 2 due to the cut in total rather than stellar mass. We applied a mass cut to the subhaloes of $5 \times 10^{11} M_{\odot}$, which approximately corresponds to a stellar mass cut of $10^{10} M_{\odot}$. The scatter is the same as for a model trained on the MACSIS-GAS clusters for an aperture of r_{200c} , however, scatter for the

Table 5. The 1σ percentile scatter, δ , for when training and applying ML models to different simulations. The formatting is the same as for Table 3. The first part of the row label denotes whether the model was training using a DMO or GAS simulation, the latter half the type of simulation the model was applied too. The two columns are using galaxies within r_{200c} and 1.5 Mpc respectively.

Train-Test	$\delta_{r_{200c}}$	$\delta_{1.5 \text{ Mpc}}$
GAS-GAS	0.037 ± 0.001	0.055 ± 0.002
DMO-DMO	0.037 ± 0.001	0.062 ± 0.002
DMO-GAS	0.037 ± 0.001	0.064 ± 0.002

1.5 Mpc fixed aperture GAS-GAS case is smaller than for either DMO-DMO or DMO-GAS by ~ 0.007 dex. This is a small difference corresponding to ~ 1.5 per cent, this is likely due to the MACSIS-GAS clusters containing 1.32 times more galaxies within 1.5 Mpc than MACSIS-DMO but only 1.15 times more at r_{200c} .

Fig. 8 shows the predicted M_{200c} masses against the true values for clusters in the MACSIS-GAS runs, for the two apertures (r_{200c} and 1.5 Mpc). The DMO training set recovers the true mass of the hydrodynamical clusters as accurately as for the DMO clusters, and introduces no significant bias (~ 1 per cent). For the 1.5 Mpc aperture we see that the high mass clusters are underestimated.

We note a few important caveats: selecting lower total mass cuts introduces a bias into the recovered MACSIS-GAS mass estimates. This is primarily driven by the increased number of galaxies in the MACSIS-GAS sample, for a mass cut of $> 10^{10} M_{\odot}$ there are twice as many galaxies per cluster in the GAS sample. The difference is likely caused by the complex interplay of subgrid physics and limited resolution. We find the bias has converged by $M_{\text{cut}} > 5 \times 10^{11} M_{\odot}$, though it is difficult to validate as above this limit many clusters do not contain the minimum 10 galaxies. We therefore present this as an indication that DMO simulations can be used to train mass estimation models to be used in systems with baryonic physics, but more work is needed to validate this claim.

6 SUMMARY AND DISCUSSION

In this work we have presented a novel way to use machine learning techniques to recover galaxy cluster masses. Features for each galaxy cluster are created from the observational properties of the clusters, such as the line-of-sight velocity dispersion, or the number of galaxies within an aperture. The features containing the most information are then identified via recursive feature elimination and this subsample of features is then used to train a ridge regression model. In this work we train the model on 90 per cent of the clusters, where the mass (or some proxy) of each cluster is known and then predict the mass of the remaining 10 per cent. When comparing the ML result to the $\sigma - M$ relation with a fixed aperture, we find the scatter in the predicted mass is reduced by a factor of 2, similar to that found in Ntampaka et al. (2016), who make use of support distribution machines which use the distribution of the line-of-sight velocities and

projected radial position of galaxies as features. Our main findings can be summarised as follows:

(i) We find ridge regression to be the optimal regression algorithm, although the simple ordinary linear regression method performs similarly well when combined with recursive feature elimination (Figs. 1 & 2).

(ii) Compared to the $\sigma - M$ relation the scatter in the recovered mass is reduced from 0.130 ± 0.004 dex (~ 35 per cent) to 0.031 ± 0.001 dex (~ 7 per cent), for clusters contaminated by interlopers out to $5r_{200c}$ (Table 2).

(iii) The addition of line-of-sight galaxy velocities does not improve the scatter in recovered mass when training on M_{200c} (Table 2).

(iv) The skewness and kurtosis of the galaxy distributions contribute a significant amount of information and should be included for optimal results. The single most powerful feature is the SZ flux inside $5r_{500c}$ (Figs. 4 & 6).

(v) We find no significant improvement over the PHOT feature set when including additional features from other observations, such as SZ or X-ray, when fitting for the true value of M_{200c} (Fig. 7)

(vi) Choosing to fit the estimated M_{200c} mass from weak lensing or the hydrostatic M_{500c} results in increased scatter of 0.063 ± 0.002 dex and 0.130 ± 0.004 dex relative to fitting the true M_{200c} mass, 0.031 ± 0.001 . In general we find the hydrostatic M_{500c} a difficult feature to reproduce and would recommend training a model using $M_{200c, WL}$. This is likely to be due to the intrinsic scatter in the X-ray observations, not the change in aperture from M_{200c} to M_{500c} (Tables 3 & 4).

(vii) We find that training a model using DMO simulations and then estimating the masses of clusters run with hydrodynamics (GAS) gives the same scatter as estimating the masses of DMO clusters (Fig. 8 & 5). The scatter is slightly larger than for GAS-GAS training and testing, though the galaxy selection criteria differ. There is a bias introduced when estimating the mass of MACSIS-GAS clusters using a MACSIS-DMO trained model which is dependent on the mass cut used for the galaxies; this is likely due to the complex interplay between the subgrid model and the mass resolution.

The ML techniques that we have presented in this paper have several caveats that we shall explicitly state here. The first is that we present models that have been trained using the true value of M_{200c} . Clearly, in an observational setting one would never have access to the true mass of the cluster and so the models would need to be trained using either a mass proxy or some other mass estimate, such as the hydrostatic mass. We mimic such a scenario in Section 5.4, with promising results when using mass estimates derived from weak lensing analysis. However, if the training mass is biased, then the predicted masses from the ML model will also share that bias. One could attempt to circumvent this by training a model using simulations and then applying that to observations directly, perhaps with a validation sample of clusters with masses estimated independently. This would require all of the observational systematics to be properly modelled and accounted for as well as the explicit assumption that the simulated clusters accurately represent reality.

We also note that while our sample of clusters include projection effects and interlopers, due to the nature of ‘zoom’

simulations only relatively nearby ($< 5r_{200c}$) interlopers are included due to the limited simulation volume. Planned simulations in the near future will allow for the creation of light cones to properly mimic the inclusion of interlopers. In the meantime, we have also performed preliminary work, artificially injecting interlopers out to ± 50 Mpc. We created the interloper sample by sampling the galaxy stellar mass function (GSMF), using the single Schechter function parameters found in Table A1 of Furlong et al. (2015). The galaxies were distributed uniformly ± 50 Mpc from the cluster, though none were added in the region less than $5r_{200c}$ away from the cluster, with line of sight velocities set to the Hubble flow. We found little change in the predicted mass scatter.

The strength of the ML approach suggested in this work is that it can easily include multiple features at once to create a pseudo-scaling relation. This allows one to recover one feature, in this case the cluster mass, from a set of available features. This has the effect of drastically reducing the scatter compared to, say, the velocity dispersion mass relation.

In conclusion, ML algorithms, and in particular ridge regression, are a promising new approach to cluster mass estimation. The models can either be trained using simulated galaxy clusters, or by using observational data. To train a model using observational data one requires a good sample of galaxy clusters with estimated masses obtained via weak lensing for example. This would form the training set used to train the ML model, which would then be applied to a larger cluster catalogue which overlaps with the initial training data. Future work should focus on making more accurate mock observations to replicate the identification of clusters and member selection. This would likely cause a difference in the PHOT and SPEC feature sets once the entire observing pipeline is replicated. Identification of additional features, such as substructure and colour measurements could also improve the scatter.

ACKNOWLEDGEMENTS

This work used the DiRAC Data Centric system at Durham University, operated by the Institute for Computational Cosmology on behalf of the STFC DiRAC HPC Facility (www.dirac.ac.uk). This equipment was funded by BIS National E-infrastructure capital grant ST/K00042X/1, STFC capital grants ST/H008519/1 and ST/K00087X/1, STFC DiRAC Operations grant ST/K003267/1 and Durham University. DiRAC is part of the National E-Infrastructure. DJB and STK acknowledge support from STFC through grant ST/L000768/1. TJA is supported by an STFC studentship. The authors thank Monique Henson for providing the weak lensing masses for the MACSIS clusters. This work benefited greatly from the availability of the SCIKIT-LEARN python package (Pedregosa et al. 2011).

REFERENCES

- Allen S. W., Evrard A. E., Mantz A. B., 2011, *ARA&A*, **49**, 409
- Armitage T. J., Kay S. T., Barnes D. J., Bahé Y. M., Dalla Vecchia C., 2018a, preprint, ([arXiv:1809.01704](https://arxiv.org/abs/1809.01704))
- Armitage T. J., Barnes D. J., Kay S. T., Bahé Y. M., Dalla Vecchia C., Crain R. A., Theuns T., 2018b, *MNRAS*, **474**, 3746
- Arnaud M., Pointecouteau E., Pratt G. W., 2007, *A&A*, **474**, L37

- Bahé Y. M., McCarthy I. G., Crain R. A., Theuns T., 2012, *MNRAS*, **424**, 1179
- Baldry I. K., et al., 2012, *MNRAS*, **421**, 621
- Ball N. M., Brunner R. J., 2010, *International Journal of Modern Physics D*, **19**, 1049
- Barnes D. J., Kay S. T., Henson M. A., McCarthy I. G., Schaye J., Jenkins A., 2017a, *MNRAS*, **465**, 213
- Barnes D. J., et al., 2017b, *MNRAS*, **471**, 1088
- Becker M. R., Kravtsov A. V., 2011, *ApJ*, **740**, 25
- Beers T. C., Flynn K., Gebhardt K., 1990, *AJ*, **100**, 32
- Bernardi M., Meert A., Sheth R. K., Vikram V., Huertas-Company M., Mei S., Shankar F., 2013, *MNRAS*, **436**, 697
- Bett P., Eke V., Frenk C. S., Jenkins A., Helly J., Navarro J., 2007, *MNRAS*, **376**, 215
- Bocquet S., et al., 2015, *ApJ*, **799**, 214
- Booth C. M., Schaye J., 2009, *MNRAS*, **398**, 53
- Caldwell C. E., McCarthy I. G., Baldry I. K., Collins C. A., Schaye J., Bird S., 2016, *MNRAS*, **462**, 4117
- Dalla Vecchia C., Schaye J., 2008, *MNRAS*, **387**, 1431
- Dolag K., Borgani S., Murante G., Springel V., 2009, *MNRAS*, **399**, 497
- Elahi P. J., Power C., Lagos C. d. P., Poulton R., Robotham A. S. G., 2017, preprint, ([arXiv:1712.01989](https://arxiv.org/abs/1712.01989))
- Evrard A. E., et al., 2008, *ApJ*, **672**, 122
- Freund Y., Schapire R. E., 1997, *Journal of Computer and System Sciences*, **55**, 119
- Furlong M., et al., 2015, *MNRAS*, **450**, 4486
- Guyon I., Weston J., Barnhill S., Vapnik V., 2002, *Machine Learning*, **46**, 389
- Henson M. A., Barnes D. J., Kay S. T., McCarthy I. G., Schaye J., 2017, *MNRAS*, **465**, 3361
- Hoekstra H., Herbonnet R., Muzzin A., Babul A., Mahdavi A., Viola M., Cacciato M., 2015, *MNRAS*, **449**, 685
- Hoerl A. E., Kennard R. W., 1970, *Technometrics*, **12**, 55
- Katz N., White S. D. M., 1993, *ApJ*, **412**, 455
- Kettula K., et al., 2015, *MNRAS*, **451**, 1460
- Kravtsov A. V., Borgani S., 2012, *ARA&A*, **50**, 353
- Lau E. T., Kravtsov A. V., Nagai D., 2009, *ApJ*, **705**, 1129
- Li C., White S. D. M., 2009, *MNRAS*, **398**, 2177
- Lin Y.-T., Stanford S. A., Eisenhardt P. R. M., Vikhlinin A., Maughan B. J., Kravtsov A., 2012, *ApJ*, **745**, L3
- Mahdavi A., Hoekstra H., Babul A., Bildfell C., Jeltama T., Henry J. P., 2013, *ApJ*, **767**, 116
- Mantz A. B., et al., 2015, *MNRAS*, **446**, 2205
- Mantz A. B., et al., 2016, *MNRAS*, **463**, 3582
- Maughan B. J., et al., 2008, *MNRAS*, **387**, 998
- McCarthy I. G., Schaye J., Bird S., Le Brun A. M. C., 2017, *MNRAS*, **465**, 2936
- Munari E., Biviano A., Borgani S., Murante G., Fabjan D., 2013, *MNRAS*, **430**, 2638
- Nelson K., Lau E. T., Nagai D., 2014, *ApJ*, **792**, 25
- Ntampaka M., Trac H., Sutherland D. J., Battaglia N., Póczos B., Schneider J., 2015, *ApJ*, **803**, 50
- Ntampaka M., Trac H., Sutherland D. J., Fromenteau S., Póczos B., Schneider J., 2016, *ApJ*, **831**, 135
- Ntampaka M., et al., 2018, preprint, ([arXiv:1810.07703](https://arxiv.org/abs/1810.07703))
- Oguri M., Hamana T., 2011, *MNRAS*, **414**, 1851
- Okabe N., Zhang Y.-Y., Finoguenov A., Takada M., Smith G. P., Umetsu K., Futamase T., 2010, *ApJ*, **721**, 875
- Old L., Gray M. E., Pearce F. R., 2013, *MNRAS*, **434**, 2606
- Old L., et al., 2014, *MNRAS*, **441**, 1513
- Pedregosa F., et al., 2011, *Journal of Machine Learning Research*, **12**, 2825
- Planck Collaboration et al., 2014a, *A&A*, **571**, A1
- Planck Collaboration et al., 2014b, *A&A*, **571**, A20
- Pratt G. W., Croston J. H., Arnaud M., Böhringer H., 2009, *A&A*, **498**, 361
- Rasia E., et al., 2012, *New Journal of Physics*, **14**, 055018
- Rozo E., Rykoff E. S., Becker M., Reddick R. M., Wechsler R. H., 2015, *MNRAS*, **453**, 38
- Saliwanchik B. R., et al., 2015, *ApJ*, **799**, 137
- Saro A., Mohr J. J., Bazin G., Dolag K., 2013, *ApJ*, **772**, 47
- Schaye J., Dalla Vecchia C., 2008, *MNRAS*, **383**, 1210
- Schaye J., et al., 2010, *MNRAS*, **402**, 1536
- Sereno M., Ettori S., 2015, *MNRAS*, **450**, 3675
- Simet M., McClintock T., Mandelbaum R., Rozo E., Rykoff E., Sheldon E., Wechsler R. H., 2017, *MNRAS*, **466**, 3103
- Springel V., 2005, *MNRAS*, **364**, 1105
- Springel V., White S. D. M., Tormen G., Kauffmann G., 2001, *MNRAS*, **328**, 726
- Sun M., Voit G. M., Donahue M., Jones C., Forman W., Vikhlinin A., 2009, *ApJ*, **693**, 1142
- Sutherland D. J., Xiong L., Póczos B., Schneider J., 2012, preprint, [p. arXiv:1202.0302](https://arxiv.org/abs/1202.0302) ([arXiv:1202.0302](https://arxiv.org/abs/1202.0302))
- Tormen G., Bouchet F. R., White S. D. M., 1997, *MNRAS*, **286**, 865
- Vikhlinin A., Kravtsov A., Forman W., Jones C., Markevitch M., Murray S. S., Van Speybroeck L., 2006, *ApJ*, **640**, 691
- Vikhlinin A., et al., 2009, *ApJ*, **692**, 1033
- Weinberg D. H., Mortonson M. J., Eisenstein D. J., Hirata C., Riess A. G., Rozo E., 2013, *Phys. Rep.*, **530**, 87
- Wiersma R. P. C., Schaye J., Smith B. D., 2009, *MNRAS*, **393**, 99
- Yee H. K. C., Ellingson E., 2003, *ApJ*, **585**, 215
- Zhang Y.-Y., Andernach H., Caretta C. A., Reiprich T. H., Böhringer H., Puchwein E., Sijacki D., Girardi M., 2011, *A&A*, **526**, A105

This paper has been typeset from a $\text{\TeX}$ / $\text{\LaTeX}$ file prepared by the author.