Can Machine Learning Help Identify Underperforming Schools?

James Tierney, Rhian Luard, Kitty McCarthy

Working Paper No. 004 August 2025

EMAP Working Paper Series

Queen Mary University of London

EXECUTIVE SUMMARY

(FOR NON-ECONOMISTS)

Using school-level data provided by Ofsted and the Department for Education, we explore and evaluate the use of supervised machine learning methods to predict the outcomes of school inspections for state secondary schools in England. To aid in identifying schools at risk of underperforming in future Ofsted inspections we compared and contrasted five commonly used classification algorithms from the machine learning literature.

We show that our models of choice perform significantly better at predicting the outcome of school inspections than random selection across all classification thresholds. We can therefore make the following policy recommendations:

- **Machine learning methods can help regulators identify underperforming schools –** school regulators looking to sample schools at risk of decline more often can make use of supervised machine learning methods to identify schools likely to be rated as underperforming so they can focus resources where they will likely have the greatest impact

- **Publicly available predicted scores can help fill informational gaps for parents –** school regulators do not have the capacity to inspect every school each year, meaning some schools can go many years without an inspection. Making the predicted outcomes publicly available will allow parents to make more informed decisions about current and prospective schools for their children

- **Model performance is high despite using only publicly available data –** we recommend regulators and government departments with access to secure microdata on schools and parental responses to Ofsted questionnaires look to further extend our work by including this data

CAN MACHINE LEARNING HELP IDENTIFY UNDERPERFORMING SCHOOLS?

Introduction

School inspections in their current form have been monitoring the quality of education provided by England's state-funded school system since the establishment of Section 9 of the Education (Schools) Act in 1992. They can be a time and resource intensive process and regulators do not typically have the capacity to inspect every school each year. According a 2018 report from the National Audit Office, England's school regulator, the Office for Standards in Education, Children's Services and Skills (Ofsted) had been struggling to meet its own inspection targets, in part due to shortages in available inspectors. It found the regulator had increased the average time between inspections, especially amongst schools it considered well-performing, to free up resources for other activities including increased focus on schools at risk of decline (National Audit Office, 2018). The same report had found that the average total cost of school inspections was £7200 as of 2018, meaning there are significant cost reductions to be had with better resource management (ibid., p.4). The success of this resource reallocation of course relies upon accurate identification of which schools are likely to be considered underperforming in future inspections.

It is also important to consider that there is an asymmetry of information between parents and schools on the quality of education provided, especially in years where the school is not inspected. With the average time between inspections sitting at four years (ibid., p.8) for secondary schools, this leaves parents close to a school without an inspection in many years with little understanding of how it would likely perform if it were to be rated today. Methods which better estimate how a school would perform in years it is not selected for inspection can therefore help plug some of the informational gaps produced by sampling.

Using school-level data provided by Ofsted and the Department for Education, we intend to explore and evaluate the use of supervised machine learning methods to predict the outcomes of school inspections for state secondary schools in England. In particular, we wish to understand if these methods could help regulators to identify underperforming schools such that resources can be focussed where they will likely have the greatest impact. We intend to compare and contrast the performance of commonly used classification methods in their ability to predict school inspections and provide guidance on which methods are likely to result in the greatest out of sample performance.

Machine learning in economics and education

Economists, particularly those in applied econometric fields (and especially in the economics of education), have historically focused on questions of causal inference. For example, much attention across both academic literature and government policy has been paid to understanding the relationship between educational attainment and earnings (Card, 1999; Department for Education, 2021). The most common tool used by economists to answer questions of cause and effect has been linear regression analysis and extensions designed to remove potential sources of selection bias, such as instrumental variables or regression discontinuity.

Techniques developed in the machine learning literature, although often based upon common statistical tools such as linear regression, are primarily designed to answer questions relating to prediction. Specifically, the goal is to find some function f which provides a good prediction (out of sample) of y given a set a features X – this has led to the adoption of processes such as *regularisation*, *bootstrapping* and *bagging* which help analysts address issues of overfitting which generate large differences in performance between in-sample and out-of-sample predictions. Using machine learning methods can therefore open up, to both economists and those working in education policy, a different set of tools which are optimised to solve a whole different set of questions.

To distinguish between the two sets of questions discussed above, we can consider the decision process around investment in supply of available teachers. The first question policymakers will want to understand is *'What is the impact on student performance of hiring an additional teacher?'* - this is a causal question for which there is a rich source of econometric literature to provide evidence (Thomas J. Kane, 2008). However, understanding the answer to this question does not provide insight into *which* teachers a local authority should hire to improve performance – forecasting teacher quality is a prediction problem. Chalfin et al. (2016) provide evidence on how machine learning tools can predict worker productivity using data from teacher tenure decisions. Below we provide two additional examples of how machine learning methods these have been applied in practice.

Using data on unauthorised absences and punctuality from administrative records in the Korean National Educational Information system (NEIS), Chung & Lee (2019) develop a random forest algorithm which can be used as an early warning system to identify students at risk of dropping out during high school level education. The optimal model performed well across classification thresholds with an AUC score of 0.97, indicating that the tool could feasibly be used to help schools identify students at risk of leaving high school before graduating.

In similar spirit to the aims of this report Kang, Kuznetsova, Luca, and Choi (2013) make use of customer reviews posted on the website *Yelp;* historic inspection results, and restaurant metadata to predict the outcome of hygiene inspections via the use of a machine learning technique called Support Vector Machines (SVM). Performance of the model was significantly better than random prediction with an accuracy score of 81.37%. The authors themselves suggest that the model could be used by inspectors to better allocate scarce resources and by customers wishing to visit restaurants without a recent inspection.

Background to Ofsted inspections

The Office for Standards in Education, Children's Services and Skills (Ofsted) are the non-ministerial department responsible for reporting on the quality of education across schools in England. They report directly to parliament and are required to inspect:

- Maintained and academy schools (mainstream and special schools)
- Pupil referral units
- Alternative provision academies
- Independent schools not affiliated to the Independent Schools Council (ISC)

Ofsted has two main types of routine inspections, ‘Section 5’ and ‘Section 8’ inspections, which are both named after the area of the Education Act 2005 from which they are derived. Section 5, or ‘full’ inspections are undertaken at prescribed intervals, with Ofsted inspectors asked to make graded judgements on the following criteria:

- Effectiveness of leadership and management
- Quality of teaching, learning and assessment
- Personal development, behaviour and welfare
- Outcomes for children and learners

These component criteria were used until September 2019 but have subsequently been updated as part of the Education Inspections Framework (Ofsted, 2015). Inspectors will also make comments on the quality of safeguarding provision, quality of adjustments for pupils with special educational needs or disabilities (SEND) and the effectiveness of pupil’s ‘spiritual, moral, social and cultural’ development. Together with the separated judgement scores, this is then aggregated into a single summary measure which Ofsted calls ‘Overall Effectiveness’. This is a four point grading system as follows:

- Grade 1 (Outstanding)
- Grade 2 (Good)
- Grade 3 (Requires improvement)
- Grade 4 (Inadequate)

According to the latest annual report of Her Majesty’s Chief Inspector, the proportion of state-funded secondary schools rating ‘good’ or ‘outstanding’ is currently at 62% (Ofsted, 2020). It is worth noting that schools graded as ‘Inadequate’ can also receive additional labels of ‘Serious Weakness’ or ‘Special Measures’ which help further differentiate schools which are considered to be highly underperforming. Schools rated as requires improvement or inadequate will usually receive a full re-inspection (under Section 5) within three years.

This differs from schools who are rated good or outstanding who are typically inspected every 4 years under Section 8. Section 8 inspections, or ‘short’ inspections do not result in a graded judgement – however if there is evidence gathered that the school’s performance has changed significantly, Ofsted will then upgrade the inspection to a full Section 5 procedure. In combination with Ofsted’s risk assessment process (of which further details are provided below) this results in a selection process where underperforming schools are more likely to undergo a full inspection in any given year.

How should we model inspection outcomes?

As outlined in the *Ofsted School Inspection Handbook* (Ofsted, 2022) inspectors are not only required to review classroom teaching and pupil reaction but will also gather a wealth of other detail including performance in statutory assessments and qualifications, cultural capital development, and pupil background. As these factors enter into the decision framework from which inspectors assess schools, we are able to make use of data collected directly on these measures or other indicators which might form part of the data generating process (or are correlated with these random variables) resulting in inspection outcomes. Broadly speaking, these factors can be either; related to the school itself (e.g., school type, previous inspection results), the pupils currently attending the school (e.g., attainment, socio-economic status) or the staff (e.g., teaching quality, pupil: teacher ratio, availability of support staff) both inside and outside the classroom.

Figure 1 displays the relationship between the percentage of students (in the year prior to the latest inspection) obtaining 5 A*-C qualifications in their end of KS4 assessments and their corresponding average performance at KS2, prior to joining the school. It is quite clear to see that using only these two indicators we are able to start separating schools likely to be considered underperforming by Ofsted from higher performing schools with cohorts which perform better in national examinations.

Figure 1: Latest Ofsted inspection outcome category by KS4 and KS2 exam performance

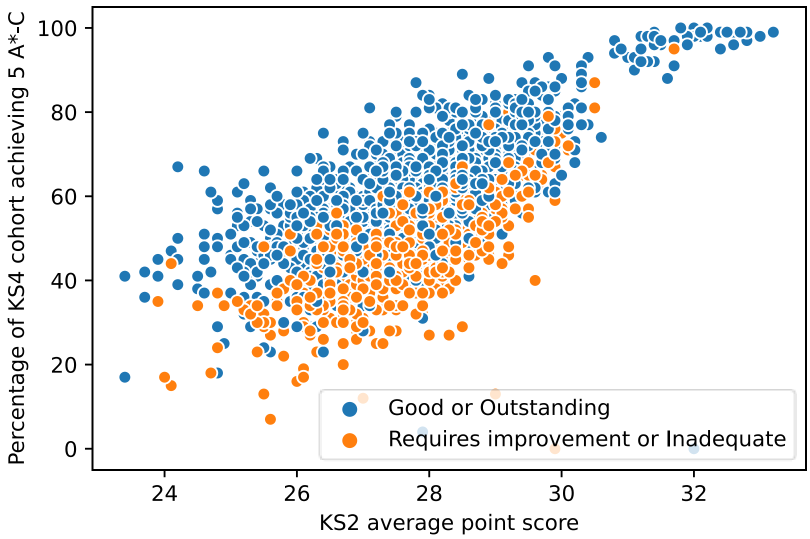

Source: Author plot using DfE and Ofsted data

In figure 2 we can also see evidence of persistence in a school's Ofsted rating over time – almost 9 in 10 schools rated as outstanding in their previous Ofsted inspection were given a good or outstanding rating in their latest inspection. Proportionately far fewer of schools given an Ofsted rating at or below grade 3 are considered good or outstanding according to their latest result, although a majority have still managed to see an improvement over time. The strong relationship between current and lagged inspection outcomes mean that we are likely to gain significant predicting from understanding both the overall outcome of previous inspections but also how the total score was derived.

Figure 2: Latest Ofsted inspection outcome category by Ofsted grade at previous inspection

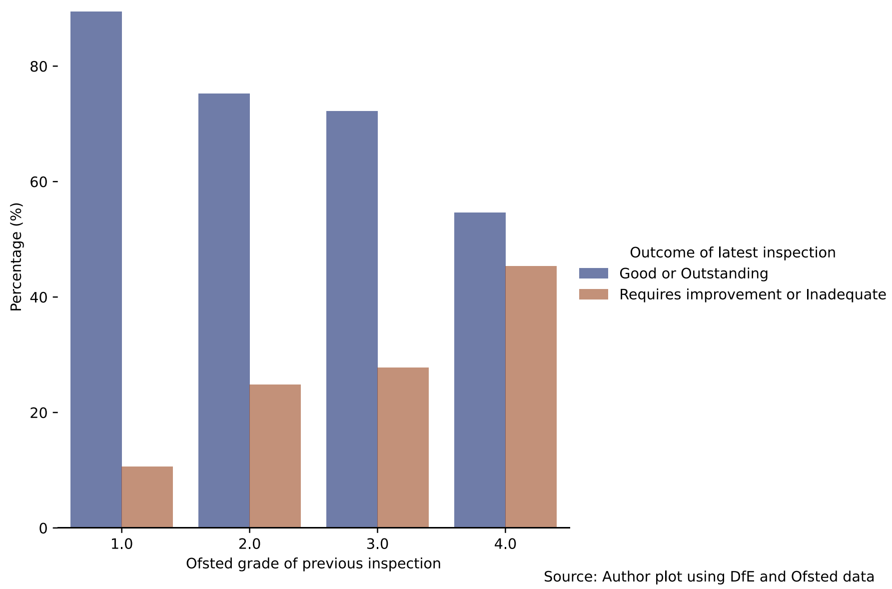

Source: Author plot using DfE and Ofsted data

Figure 3 shows a similar relationship between the Index of Multiple Deprivation Affecting Children (IDACI) quintile of the school and its latest Ofsted inspection outcome. Areas with the greatest levels of deprivation affecting children proportionately have a much larger share of underperforming schools than their more affluent counterparts.

Figure 3: Latest Ofsted inspection outcome category by IDACI quintile

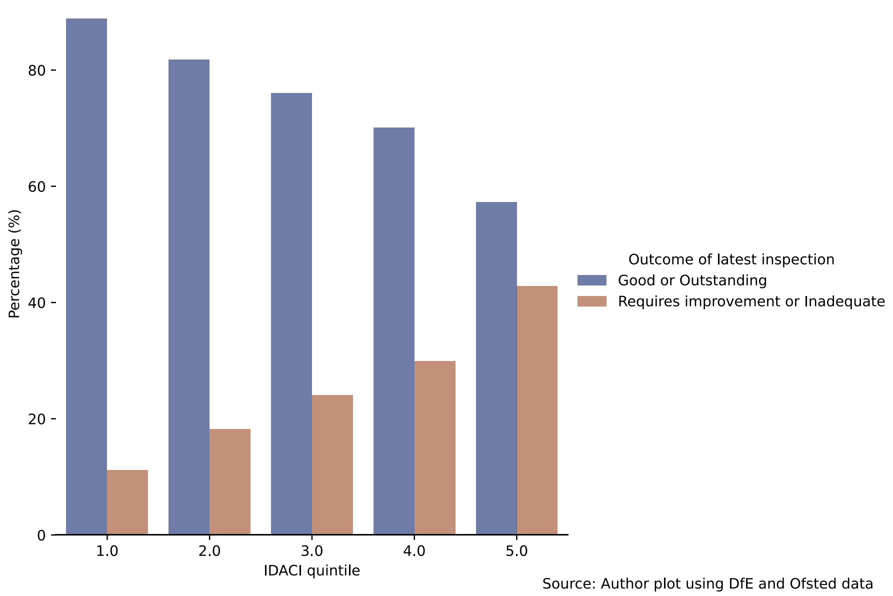

Source: Author plot using DfE and Ofsted data

Data & sample selection

To be included in the sample used to train and evaluate models, schools had to meet the following criteria:

- Be a state funded school within England providing secondary education (ages 11-16)
- Have available KS4 exam data for the KS4 cohort in the academic year prior to their latest inspection
- Have been subject to at least two full (section 5) Ofsted inspections since opening
- The school's latest inspection must have been on or after the academic year ending in 2011 and on or before the academic year ending 2019
- Have complete (no missing or suppressed) data for all variables under consideration

Of the 3448 state-funded English secondary schools open in the 18/19 academic year, we retained 2132 for whom we constructed a suitable linked dataset.

The main data sources used for this analysis are publicly available versions of Ofsted school inspections data and the National Pupil Database (NPD) aggregated to the school level, with full population coverage of schools across England.

Data provided by Ofsted includes information on both the latest and previous full (Section 5) inspection of each school with at least one completed inspection. Details are available on the 'Overall Effectiveness' grade awarded, component criteria scores, the type of inspection and flags for 'Special Measures' or 'Serious Weakness'. Limited data on school characteristics are provided but importantly include the Index of Multiple Deprivation Affecting Children (IDACI) quintile of the Super Output Area (SOA) where the school is located. This provides us with a small-area based proxy for socioeconomic deprivation which would otherwise only be captured via Free School Meal eligibility rates.

The National Pupil Database is a set of data collected by the Department for Education and contains rich details on pupil demographics, school characteristics and attainment data for national curriculum tests and other public examinations. For example, we are able to identify the gender composition, admissions policy and the proportion of students eligible for Free School Meals (FSM) at each school. All state-funded school types were included in the analysis and we have specifically flagged schools which changed school type between their previous and latest full inspection.

Outside of England, other devolved nations of the UK have their own school regulators, meaning use of the NPD captures data on all schools available to be inspected by Ofsted. We however focus our analysis on state-funded secondary schools who had their latest inspection between 2011-2019 as this allows us to make use of prior attainment data at KS2 which was not collected in a consistent format prior to this period. Independent schools are excluded on the basis that the majority are not required to be inspected by Ofsted.

Methodology

To aid in identifying schools at risk of underperforming in future Ofsted inspections we rely upon five commonly used classification algorithms from the machine learning literature (Zou & Hastie, 2005; Hoerl & W.Kennard, 1970; Ho, 1995):

- Logistic regression
- Ridge regression
- Lasso regression
- Elastic net regression
- Random forest

In this case, ridge, lasso and elastic net approaches are all extensions of the logistic regression specification; they are designed to prevent overfitting to the training data by introducing penalty parameters into the loss functions used to estimate regression coefficients. Overfitted models tend to capture not only the underlying data generating process but also noise included in the training data, thus leading to suboptimal predictive performance when presented with new, unseen testing data (Menelaos Pavlou, 2015). Regularisation processes attempt to achieve lower model error by reducing the variance of estimates at the cost of intentionally introducing bias.

The random forest algorithm is an extension of decision trees and is commonly used as a classification or regression method where prediction of an outcome is the primary goal. Principally it attempts to address issues of overfitting common to decision tree algorithms by aggregating predictions from a multitude of decorrelated trees. Key to the process is that each decision tree node uses a randomly selected subset of features from the original list (Oshiro T.M., 2012)

To indicate instances of underperformance during school inspections we construct a binary target variable using the 'Overall Effectiveness' summary measure provided at the latest Ofsted inspection of school i as follows:

$$Underperform_i = \begin{cases} 1 \; if \;\; 'Requires\ improvement' \, or \;\; 'Inadequate' \\ \qquad\qquad 0 \; if \;\; 'Good' \; or \; 'Outstanding' \end{cases}$$

A total of 55 predictor variables are included in each machine learning algorithm – many of which are one hot encoded dummy variables where one reference category has been excluded to avoid linear dependence of columns. A full list of features and an accompanying description is provided in Appendix A.

The approach taken to train, evaluate and compare each algorithm can be summarised by the following five step process:

1. Standardize and dummy code predictor variables
2. Split dataset into stratified **testing** and **training** samples (50:50)
3. Perform **5-fold cross-validation** process on training data to determine **optimal hyperparameters** for penalized logistic regression and random forest algorithms.
4. Use optimal hyperparameters to train models on full training data
5. Evaluate model performance by comparing Receiver Operator Characteristic (ROC) curves and Area Under the Curve (AUC) values using the unseen **testing** data

Table 1 summarises the required hyperparameters for each algorithm and the subset of possible options which were used as part of the cross-validation process. The resulting optimal hyperparameters selected for each algorithm can be found in Appendix B

Table 1: Hyperparameter selection criteria

Hyperparameter	Description	Algorithm(s)	Range of values
$\boldsymbol{\lambda}$	Regularisation penalty	Lasso, ridge and elastic net	$10^{-4} - 10^{4}$
$\boldsymbol{\alpha}$	Mixing parameter between ridge (α = 0) and lasso (α = 1)	Elastic net	0.2 – 0.8
N estimators	Number of trees in random forest	Random forest	100 - 500
Max features	Number of features to consider at each split	Random forest	$\sqrt{p}$
Max depth	Maximum number of levels in each tree	Random forest	5 - 45
Min sample split	Minimum number of samples required to split node	Radom forest	2, 5 & 10

Results

The selection criteria used to identify the optimal set of hyperparameters for each algorithm was the mean Area Under the Curve (AUC) score from a 5-fold cross validation process. The mean AUC score was selected as the scoring criteria across this analysis as it provides a summary measure of model performance across all possible classification thresholds. This is important in the context of school inspections as regulators should be clear how a model will perform holistically, in an environment where there is likely to be constant pressure to alter the rate of true positives or false positives. Using figure 3, which shows how this measure changes as a function of the regularisation penalty, we can see early indications that the introduction of penalised regression methods appears to provide a limited boost in predictive performance for school inspection outcomes.

Figure 3: Five fold cross-validation mean AUC score by penalty parameter value

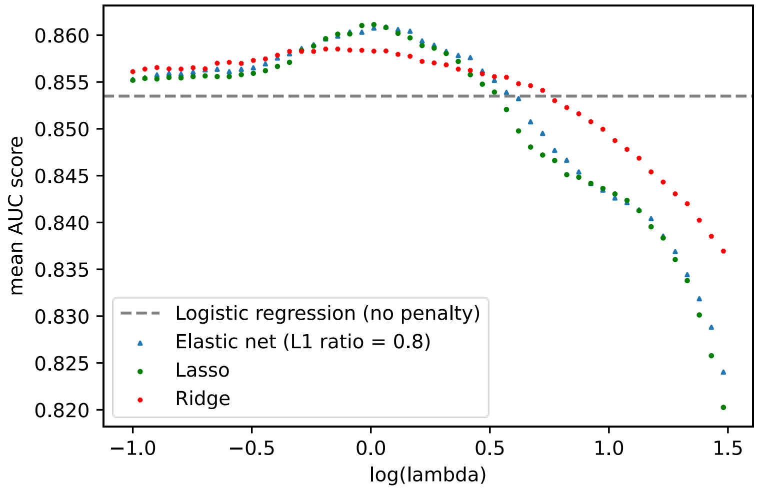

Source: Author plot using DfE and Ofsted data

Similar results are again shown when each algorithm is evaluated using new data. Figure 4 shows the Receiver Operator Characteristic (ROC) curve and corresponding AUC values for each algorithm when evaluated using the unseen testing data. To provide some context to these results, AUC scores can take on values from a range between 0 and 1, where a model with an AUC of 0.5 performing no better or worse than random prediction – this is represented by the 45 degree line. A model whose predictions are 100% accurate would have an AUC of 1. Another useful interpretation of AUC values is as the probability that a random positive unit is ranked higher than a random negative unit from the output of the model.

Figure 4: Receiver Operating Characteristic (ROC) curve

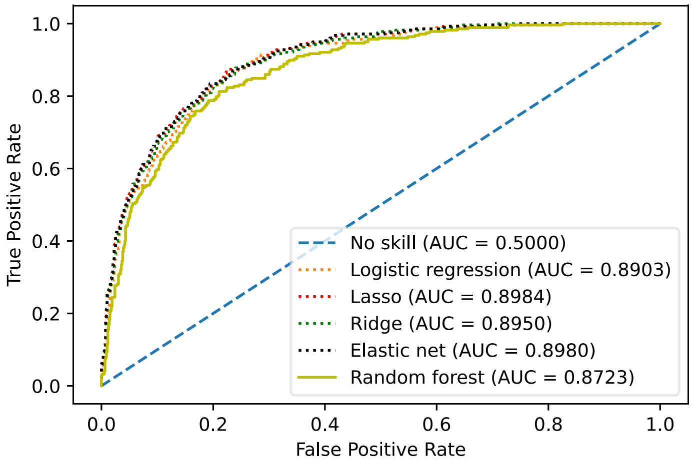

Source: Author plot using DfE and Ofsted data

ROC curves demonstrate the relationship between the true positive rate and false positive rate of an algorithm when the classification threshold is changed. All models perform significantly better at predicting the outcome of school inspections than random selection across all classification thresholds, as demonstrated by the position of all ROC curves to the left of the 'No Skill' line. More specifically, this demonstrates that for any given false positive rate (where our model predicts underperformance in schools when the actual outcome is 'Good' or 'Outstanding'), the true positive rate (where a prediction of underperformance coincides with an actual outcome of underperformance) is much higher when any of these machine learning methods are used, relative to random prediction. Penalised logistic regression methods provide the greatest predictive performance according out of sample AUC score, but the gains are relatively small when compared to standard logistic regression specifications. Random forest methods are outperformed by all logistic regression specifications (penalised or unpenalised) but still offer much greater predictive performance than a model with no skill. These results indicate that the underlying relationship between school inspection outcomes and our chosen features is likely to be linear.

Finally, figure 5 shows the distribution of permutation importance scores for the ten features found to have the greatest influence in predicting school inspection outcomes for our elastic net regression algorithm. These scores are derived by taking the difference in AUC scores when a single feature value is randomly shuffled. Our results show that the proportion of students obtaining 5 A*-C in the year prior to the latest inspection and their corresponding prior attainment at KS2 have by far the greatest feature importance. Indicators for the year of prior inspection are likely to have been influential due to the non-random sampling Ofsted employs to inspect schools. Readers should note that when predictors display multicollinearity, permutating one feature will have little effect on model performance because the model can derive the same information from features it is correlated with. This can lead to lower feature importance scores than you would expect given the high AUC scores found above.

Figure 5: Permutation Importances (Test Set)

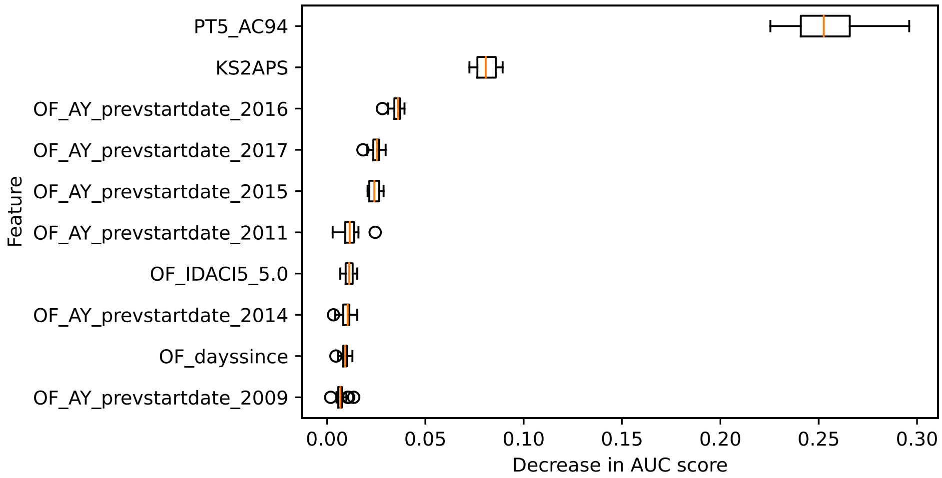

Source: Author plot using DfE and Ofsted data

Conclusion

Using publicly available data provided by Ofsted and the Department for Education we have explored and evaluated the use of common supervised machine learning methods to predict the outcomes of school inspections for secondary schools in England. Our results suggest that such methods provide much greater predictive performance compared to random prediction and could feasibly be used by regulators to better identify underperforming schools. Given the significant costs associated with inspection, use of these tools in combination with human oversight can help focus resources where they will likely have the greatest impact.

Even with better resource management, regulators are unlikely to ever have the capacity to inspect each school every year. We therefore recommend that predicted outcomes for schools are made publicly available, addressing the significant informational asymmetries during school selection and giving parents the opportunity to make more informed decisions around prospective schools for their children.

There remains a host of additional features which could be explored with access to secure microdata (e.g., school absences, parental views and expenditure per pupil) which we recommend are explored by regulators as possible extensions to further improve model performance. Our feature importance results also suggest that continued availability of KS4 and KS2 attainment in the public domain are key to successful use of the methods outlined in this report.

REFERENCES

Card, D., 1999. The Cause and Effect of Education on Earnings. *Handbook of Labour Economics,* Volume 3, pp. 1801-1863.

Chalfin, A. O. D. A. H. Z. J. M. L. J. L. S. M., 2016. Productivity and Selection of Human Capital with Machine Learning. *American Economic Review,* 5(106), pp. 124-127.

Department for Education, 2021. *GCSE attainment and lifetime earnings,* s.l.: s.n.

Hoerl, A. & W.Kennard, R., 1970. Ridge Regression: Applications to Nonorthogonal Problems. *Technometrics,* 12(1), pp. 69-82.

Ho, T. K., 1995. *Random Decision Forests.* Montreal, QC.

Menelaos Pavlou, G. A. S. D. I. Z. O., 2015. Review and evaluation of penalised regression methods for risk prediction in low-dimensional data with few events. *Statistics in medicine,* 35(7), pp. 1159-1177.

National Audit Office, 2018. *Ofsted's inspection of schools,* s.l.: s.n.

Ofsted, 2015. *The common inspection Framework: education, skills and early years,* s.l.: s.n.

Ofsted, 2020. *The Annual Report of Her Majesty's Chief Inspector of Education, Chlidren's Services and Skills 2019/20,* s.l.: s.n.

Ofsted, 2022. *Ofsted School Inspection Handbook.* [Online]
Available at: https://www.gov.uk/government/publications/school-inspection-handbook-eif/school-inspection-handbook
[Accessed 15 02 2022].

Oshiro T.M., P. P. B. J., 2012. How Many Trees in a Random Forest?. *Machine Learning and Data Mining in Pattern Recognition ,* pp. 154-168.

Thomas J. Kane, D. O. S., 2008. ESTIMATING TEACHER IMPACTS ON STUDENT ACHIEVEMENT:. *NATIONAL BUREAU OF ECONOMIC RESEARCH.*

Zou, H. & Hastie, T., 2005. Regularization and Variable Selection via the Elatic Net. *Journal of the Royal Statistical Society,* 2(67), pp. 301-320.

Appendix A

Table 2: Predictor variables and descriptions

Variable	Description	Data type
Previous overall effectiveness	Overall effectiveness grade awarded to each school in their previous Ofsted inspection	Integer
Previous concern 'Special Measures'	Binary indicator equal to 1 if the school received a 'Special Measures' label during their previous inspection	Integer
Previous concern 'Serious Weaknesses'	Binary indicator equal to 1 if the school received a 'Serious Weaknesses' label during their previous inspection	Integer
PNUMFSM	Percentage of total pupils who were eligible for and claimed Free School Meals in the academic year prior to the school's latest inspection	Float
PSENELSE	Percentage of total pupils who had a recorded special educational need or disability in the academic year prior to the school's latest inspection	Float
RATPUPTEA	Pupil: teacher ratio in the academic year prior to the school's latest inspection	Float
PNUMUNCFL	Percentage of total pupil whose first language was unclassified in the academic year prior to the school's latest inspection	Float
NOR	Total number of pupils enrolled in the academic year prior to the school's latest inspection	Float
ISPRIMARY	Binary indicator equal to 1 if the school provides primary education	Integer
SALARY	Mean salary of classroom teaching staff in the academic	Float

	year prior to the school's latest inspection	
NNONTEA	Total Number of Non Classroom-Based School Support Staff, Excluding Auxiliary Staff (Headcount) in the year prior to the school's latest Ofsted inspection	Integer
PNUMEAL	Percentage of students whose first language is not English, in the academic year prior to the school's latest inspection	Float
NFTETEAAS	Total Number of Teaching Assistants (Full-time Equivalent) in the year prior to the school's latest	Float
ISPOST16	Binary indicator equal to 1 if the school provides post-16 education	Integer
PT5_AC94	Percentage of KS4 pupils achieving 5 A*-C or equivalent (note these are graded as 9-4 in later years) in the academic year prior to the school's latest Ofsted inspection	Float
KS2APS	KS2 average point score of the KS4 cohort of the year prior to the school's latest Ofsted inspection	Float
OF_dayssince	Number of days elapsed between the previous and latest inspection	Float
Faith	Set of dummy variables indicating the official faith of the school (Christian, Jewish, Muslim, Non-faith & Other faith). 'Christian' is used as the reference category	Integer
IDACI quintile	Set of dummy variables indicating the IDACI quintile of	Integer

	the super output area where the school is located	
Event group	Set of dummy variables indicating the type of Section 5 inspection of the latest Ofsted inspection (Standard S5, S8 deemed S5 and S8 conversion)	Integer
School type	Set of dummy variables indicating school type in the year of the latest inspection (Community, Academy sponsor led, Academy conversion, Foundation, Free school, UTC, CTC, VA & VC)	Integer
School type change	Binary variable equal to 1 if the school changed type between the previous and latest Ofsted inspection	Integer
Admissions policy	Binary indicator equal to 1 if the school is selective in its admissions process	Integer
Gender	Set of dummy variables indicating the gender composition of the school in the year prior to the latest Ofsted inspection (girls, boys, mixed)	Integer
Previous academic year of inspection	Set of dummy variables indicating the academic year in which the previous inspection occurred	Integer

Appendix B

Table 3: Optimal hyperparameters for penalised logistic regression models

Penalty	**λ** ***(regularisation parameter)***	**α** ***(mixing parameter)***
Ridge	**0.725704**	--
Lasso	**1.029576**	--
Elastic net	**1.156888**	**0.8**

Table 4: Selected/optimal hyperparameters for random forest

Hyperparameter	***Selected/optimal value***
n estimators	**200**
max features	**sqrt(p)**
max depth	**45**
min sample split	**5**

School of Economics and Finance

Queen Mary
University of London

This working paper is based on project work
undertaken by EMAP apprentices

School of Economics and Finance
Queen Mary University of London
Mile End Road
London E1 4NS
Tel: +44 (0)20 7882 7356
Fax: +44 (0)20 8983 3580
Web: www.econ.qmul.ac.uk/research/workingpapers/