Applied Soft Computing Journal 151 (2024) 111121

Contents lists available at ScienceDirect

Applied Soft Computing

journal homepage: www.elsevier.com/locate/asoc

An approach to investigate fairness using Dominance-based Rough Sets Analysis—How fair were the COVID-19 restriction decisions in the UK?

Edward Abel [a,*], Sajid Siraj [b,c]

[a] University of Southern Denmark, Denmark
[b] Centre for Decision Research, Leeds University Business School, Leeds, UK
[c] COMSATS University Islamabad, Wah Campus, Pakistan

ARTICLE INFO

Keywords:
Multi-criteria data analysis
Dominance-based Rough Sets
COVID-19
Data wrangling
Data-driven decision making

ABSTRACT

Fairness is a crucial aspect to consider within decision support systems, to seek to strive for equitable decision outcomes. Therefore, in this work, we introduce an approach to investigate fairness in data-driven decisions. The Dominance-based Rough Sets Approach (DRSA) has been widely used to extract a single set of if-then types of rules from data. Conversely, our approach investigates fairness by extracting multiple separate if-then rule sets for separate groups. The proposed approach facilitates fairness analysis to be performed amongst groups represented by these rule-sets.

During the COVID-19 pandemic, several countries have taken the approach of tiered restrictions, which has remained a point of debate due to a lack of transparency. Using our proposed approach, we explore fairness analysis with regards to the UK government's COVID-19 tiered restrictions allocation system. These insights from the analysis are translated into "if-then" type rules, which can easily be interpreted by policy makers. The differences in the rules extracted from different geographical areas suggest inconsistencies in the allocations of tiers in these areas. We found that the differences delineated an overall north south divide in England, however, this divide was driven mostly by London. Such analysis could provide a more transparent approach to localised public health restrictions, which can help ensure greater conformity to the public safety rules. Our analysis demonstrates the usefulness of our approach, to explore fairness analysis in terms of equal-treatment within data-driven decisions, which could be applied in numerous other domains, for investigating the fairness and explainabilty of decisions.

1. Introduction

Within decision support systems, fairness considerations are important to help ensure decision outcomes are equitable, transparent and trustworthy. This research proposes a novel way to investigate fairness in data-driven decisions by identifying potential inconsistencies with the help of the Dominance-based Rough Set Approach (DRSA). DRSA has been widely researched [1] and applied in various domains [2,3]. The main benefit of using DRSA is its ability to extract if-then rules from data that can easily be presented and interpreted by non-technical people (such as managers and decision makers). Whilst DRSA has been widely used to extract rules for descriptive insights or predictive purposes; in this work we propose an approach that focuses on extending this technique for exploring fairness analysis within data driven decision making. Fairness is a crucial aspect to be considered in decision support systems to seek to avoid generating inequitable decision outcomes. It addresses the challenges of potential biases in data or algorithms, that might result in unfair discrimination against any particular group [4]. Numerous definitions of fairness have been proposed [5,6], looking to detect biases in data and/or algorithms in different nuanced ways. Within our proposed approach, fairness is explored in terms of equal treatment within data-driven decisions [7], that implies that groups should not be unjustly advantaged or disadvantaged based on inappropriate attributes. Our approach facilitates such fairness analysis through extracting multiple sets of rules, where each set is extracted from different disjoint dataset slices, from which fairness analysis through comparative analysis between the multiple sets of rules is then performed. Through identification of potential inconsistencies present between different rule sets our approach is able to assess the fairness of decision outcomes.

The proposed approach is illustrated by investigating inconsistencies in decisions related to the pandemic caused by Coronavirus disease 2019 (COVID-19). The issue of public health and safety, looking to dampen the impacts of COVID-19, forces policy makers to restrict a lot

* Correspondence to: University of Southern Denmark, Room 41.58, Universitetsparken 1, Kolding, Denmark.
E-mail addresses: abel@sdu.dk (E. Abel), s.siraj@leeds.ac.uk (S. Siraj).

https://doi.org/10.1016/j.asoc.2023.111121
Received 19 June 2023; Received in revised form 24 November 2023; Accepted 26 November 2023
Available online 5 December 2023

of business and social activities involving physical interactions. Several governments tried to address this issue by introducing region-level restrictions based on the health and safety risk faced by each region. For example, the United Kingdom (UK) government announced a tiered restriction system in September 2020 [8], as opposed to its previous country-wide restrictions policy. The allocation of tiers during the pandemic has been questioned due to a lack of transparency, leading to accusations of unfair decision making.

We demonstrate the usefulness of our proposed approach by investigating the COVID-19 tiers allocation system in the UK. After acquiring and preprocessing data from various disparate sources, we extract distinct if-then rule-sets for different geographical areas and compared these rule-sets to identify inconsistencies. The analysis identifies that there were inconsistencies amongst geographical areas in the UK suggesting an issue of fairness.

The structure of the rest of the paper is as follows. A literature review is presented in Section 2, then, Section 3 provides background to the UK's tiered response to COVID-19, along with discussions of the system's fairness. Section 4 details our methodology for collecting, preparing and processing data, and performing fairness analysis. The findings and analyses are then presented and discussed in Section 5, and Section 6 concludes.

2. Literature review

Dominance-based Rough Set Approach (DRSA) is a well-known method for multi-criteria classification, which is used to extract understandable If-Then decision rules from analysing historical data [9]. DRSA has been employed in many applications, and has been extended in numerous ways to enhance its performance and/or applicability within different situations.

2.1. DRSA rule-sets for gaining insights

DRSA has been applied to determine a single set of rules for utilisation in unearthing descriptive insights. For example, service quality in various airports has been investigated by extracting a set of rules using DRSA [10]. Data were collected from an international airport in Taiwan, which were then processed to understand passengers' perceptions of airport service levels. The problem of Multi-attribute group decision-making, is areas such as labour management negotiation problems, is explored in [11]. Here, an extended rough set model for multi-attribute group decision-making is introduced, that presents optimistic/pessimistic multi-granulation approximations for conflict analysis. The problem of effective budget allocations for pavement maintenance activities within budget constraints is tackled in [12]. Here, an Interactive Multiobjective Optimisation Dominance Rough Set Approach (IMO-DRSA) is utilised to derive a decision-rule preference model, that translates preference information into decision rules, facilitating interaction between analysts and decision-makers to allocate funds more effectively. The use of DRSA to aid the identification of areas vulnerable to COVID-19 infections in Brazil was proposed in [13]. Within the study, a number of criteria and risk factors for COVID-19 were considered. DRSA has been used to analyse socio-demographic, environmental, economic, and accessibility dimensions in two northern Italian rural provinces [14]. Here, from the analysis, hidden patterns and decision rules are discovered, revealing how economic growth often conflicts with environmental sustainability. Analysis of the relationships in multiple-unit pellet systems (MUPS) is explored using DRSA in [15]. Here, extracted decision rules indicate properties to ensure consistent desired results, aiding technologists in optimal formulation development.

2.2. DRSA as a predictive tool

Various work has explored utilising DRSA to extract a set of rules, to then utilise for predictive purposes. For example, DRSA has been tasked to tackle the problem of detecting auto loans frauds in [16]. The method was reported to be effective and to outperform traditional techniques for fraud prediction. Similarly, in the healthcare domain, DRSA has been used to improve prediction accuracy of seminal quality [17]. Here, DRSA is explored to address the problem of the global decline in male fertility, exploring lifestyle factors looking to uncover insights for improved early detection prediction. For the problem of looking to predict symptomatic cases of COVID-19 patients the use of rough sets was explored for making such predictive classifications in [18]. Here, by analysing symptom data, the approach looks to identify key predictive indicators of the disease. The approach enhances patient assessment accuracy, so as to aid in shaping medical guidelines and policies. The issue of managing spare parts was addressed in [19], where DRSA was used to extract rules involving multiple criteria, such as criticality, price, demand, delivery time, to be utilised to aid future predictions. The proposed framework was tested in practice and was validated by the feedback received from industry experts. To address the problem of high false alarms in Tyre Pressure Monitoring Systems (TPMS), a decision rule-based tyre monitoring system using DRSA was introduced in [20]. Through considering the dynamic relationship between tyre pressure and temperate a significant reduction in false alarm predictions was achieved. An auction approach combining Agent systems with a dominance relation-based ranking was proposed in [21]. The approach looks to automate preference-based auctions without requiring extensive preference parameters from the buyer, thereby reducing their cognitive effort.

In all of the applications of DRSA above, whether used for descriptive or prediction purposes, the focus is on extracting a single set of rules from a single dataset. Conversely, our approach looks to utilised DRSA to extract and compare multiple separate rule sets from disjointed slices of a dataset.

2.3. DRSA & fuzzy applications

Work has also investigated extending DRSA so as to expand its capabilities and/or applicability within different situations. One extension area of DRSA has been to extend it to problems where a dataset contains fuzzy or uncertain information [22,23]. Such fuzzy approaches provide DRSA with a flexible and nuanced way to handle uncertainties and ambiguities in data, which in turn can enhance insights and decision-making within more complex and uncertain scenarios.

2.4. Reducts & bireducts

Some recent works have investigated having different perspectives on a dataset such as looking to prune a dataset, to achieve a balance between reduction of attributes and characterisation of objects without loss of its discernibility before extracting rules. Within rough set theory, a reduct is a minimal subset of attributes that retains the discernibility power of the original set of attributes. In this way, a reduct is the smallest set of attributes that can still distinguish between objects in the same way as the full attribute set [24]. An expansion of the concept of reducts is bireducts. While a reduct focuses on reducing the number of attributes, a bireduct looks to minimise both the number of attributes and the number of objects simultaneously [25]. Therefore, Bireducts explore calculating minimal subsets of criteria and data points that retain the discernibility power of the full dataset. This can be valuable to identify and remove redundancy, which can help reduce computational complexity and improve performance. Bireducts also can provide opportunities to enhance descriptive insights to a decision maker through facilitating the potential for extraction of more concise rules [25]. Multiple different bireducts can be derived from a single dataset which could then be utilised in combination to construct families of rule-based classifiers, looking to leverage the potential enhancement that ensemble learning provides upon classification performance. Work has explored how such bireducts can be determined, such as through tackling the Simplest Correct Decision Bireduct Ensemble Problem (SCDBEP), so as to look to determine a robust ensemble of decision bireducts [26]. As well as the potential to enhance predictive performance and stability, bireducts can also enhance descriptive insights and explainablility. Through reduction of the number of dimensions under consideration more flexible and easy to interpret data based knowledge representations can be derived, helping explainability within classification results [27], Moreover, bireducts can also provide insights regarding the relative importance of attributes [28], helping decision makers comprehend the most significance attributes within their data.

Bireducts facilitate determining multiple slices of a dataset. However, these represent reduced versions of the same data set, where different bireducts derived from a dataset may have overlapping data and thus not be disjoint. Whilst this has a number of advantages in different applications, this is quite different from our approach which looks to generate multiple disjoint subsets of the whole dataset, to then compare the separate rules extracted from these different subsets with respect to fairness analysis.

2.5. Data perspectivism and multi-calibration

Another recent employment of the notion of having different perspectives on a dataset is that of Data perspectivism, which emphasises the importance of considering multiple viewpoints when analysing data [29]. One scenario could be, given a dataset for which certain values have been labelled, say by domain experts, potentially multiple subjective opinions (perspectives) will need to be aggregated in some way to arrive at single values. Alternatively, the different perspectives could be maintained, resulting in multiple versions of a dataset so as to preserve different perspectives. Analysis and insights can then look to maintain and utilise these perspectives, to promote more nuanced decision-making, and to mitigate potential bias. In a similar vain, Multi-calibration looks to explore the potential of data bias, and looks to mitigate inadvertent or malicious discrimination that may be introduced at training time within machine learning models [30]. Potentially output may be discriminatory due to training data containing biases, which through identification could be corrected, or result in different versions of a dataset which could be analysed as part of the calibration process, to seek predictions of greater accuracy and consistency.

Such notions concerning perspectives and calibration of data are mainly focused upon potential data veracity and erroneousness issues within a dataset, from which multiple versions of a dataset may be derived, which will invariably contain various overlapping elements. Conversely, our approach's focus is on taking a dataset and carving multiple disjointed slices out of it, so as to then compare separate rule sets extracted separately from our slices.

2.6. DRSA for fairness assessment

Although DRSA has been widely researched and practically used in numerous application domains, the focus invariably remains on using a set of extracted rules to identifying patterns or to apply on new/unseen cases. Whilst exploration of DRSA extensions with result in multiple separate data-sets invariably result in separate datasets that contain data overlap, and are tasked for creating more robust classification, or to maintain alternative perspectives within the data.

However, we argue that DRSA can be used as a tool to assess fairness which is an unexplored area of its application. Therefore, in this research, we propose an approach to fairness analysis, that looks to extract separate sets of rules using DRSA for different segments of a data set. For example, considering the classical school grading example in DRSA, imagine a situation where a school administration applies DRSA to understand the grading of their students. If applying DRSA on the whole data-set, we can extract rules on how performance in different subjects translate into an overall grade. However, considering DRSA as a comparison tool, we can split the records for each gender, and DRSA can also be used to assess any inconsistencies (or similarities) in grades assigned for each gender. In this way, we could identify any inconsistency that might link to preferential treatment. Our proposed approach can be an effective tool to investigate issues that emerged during the COVID-19 pandemic, for example, to investigate fairness within a country having tiered restrictions of movement.

In the next section, we develop a case study from the United Kingdom during the COVID-19 pandemic, wherein their government introduced a system of tiered restrictions that was questioned for equal treatment. We then apply our proposed approach to demonstrate its usefulness in assessing fairness.

3. DRSA and the COVID-19 tiered system in the UK

The UK government applied country-wide lockdown restrictions when the first wave of COVID-19 pandemic hit the nation in March 2020. However, as the data collection for COVID-related metrics matured, disparities within different geographic areas became more visible. It could be sub-optimal to keep the same level of restrictions to the whole country when there are rising numbers of COVID-19 positive cases (hereafter referred as 'cases' only) only in a specific local area. Hence, when facing the second wave, the UK government announced a tiered restriction system in September 2020 [8], as opposed to a country wide restrictions policy. The tiered approach may mitigate the severity of the impacts of restrictions on economic activities, such as disruptions in the critical supply chains [31].

By monitoring certain measurable indicators, decisions can be made to move areas up a tier (if they are not improving) or down (if the trajectory improves). We discuss this in detail below, in order to make an argument that this is a multi-criteria decision making problem.

3.1. The tiered system

The tiered system, set out in the UK Government's 2020 Winter Plan [32], facilitated a more systematic and data driven approach to decision making. They proposed a set of factors (criteria) to determine what level of restrictions (Tier) should be imposed on different geographical areas in England. The restriction were tiered from Tier-1 to Tier-4,[1] where Tier-1 was the most relaxed set of restriction while Tier-4 represented the most constrained level of restrictions.

The UK government chose a set of five criteria to allocate tiered restrictions. The rationale for the choice was given as "*The indicators have been designed to give the government a picture of what is happening with the virus in any area so that suitable action can be taken* [33]". The set of criteria, to determine which Tier from 1 to 4 an area should be placed in, were (C1) case detection rate in all age groups, (C2) case detection in people aged 60 or above, (C3) how quickly case rates were rising or falling, (C4) ratio of positive cases in the general population, (C5) pressure on the healthcare service. In addition to these five criteria, further consideration pertaining to the local context and exceptional circumstances could also be considered, such as a local but contained outbreak. The set of five criteria are defined and discussed below.

Case detection rate in all age groups (C1): This criterion gives a measure of the number of cases within a given 24 h time frame for a given geographical area. The cases were detected as positive based on the tests recorded by the UK health system. From this criteria, an indication of how many people are catching COVID-19 in an area can be gleamed. However, the measure is only based on tests that are chosen

[1] Initially only Tiers 1–3 were used and then Tier 4 came into effect later.

to be taken, therefore, many cases may go undetected, and differences between areas' residents reluctance to go for tests could impact the measure.

Case detection rate in the over 60 year olds (C2) This criterion gives a measure of the number of cases within a given 24 h time frame (for a given geographical area) where the patient was over sixty years old. Due to the risk of serious illness from COVID-19 increasing with age this criterion is able to differentiate better the severity of the outbreak in an area than just a overall number of cases, which cannot determine if, the majority of cases are over 60 and thus more serious than if the majority of cases are under 60. Like C1, this criterion is only based on tests that are chosen to be taken and so this self selecting sample group could impact the accuracy of the measure.

The rate at which cases are rising or falling (C3) This criterion gives a measure of the rate of change of the number of cases between one 24 h time frame and the next. From this an indication of whether the number of cases is growing or receding can be determined. As this criterion denotes a rate of change, and not an absolute values, it can be the case that high volatility can be present when the number of cases are small. Moreover, like C1 and C2, this criterion is also based on tests that are chosen to be taken so again represents a self selecting sample group.

Positivity rate (C4): The positivity rate criterion is a measure of the number of positive cases that are detected as a percentage of all the tests taken within a given 24 h time frame. From this an indication of the general prevalence of COVID-19 can be extrapolated, however, it is calculated only from tests that are chosen to be taken, and potentially those that are more likely to have COVID-19 symptoms may be more likely to look to confirm this with a test.

Pressure on the NHS (C5): This criterion gives an indication of the pressure that the cases are having on the NHS infrastructure for a given geographical area. This is an important consideration as if the pressure were to become so high that the NHS infrastructure becomes overwhelmed then, its ability to tackle the cases that result in hospitalisation would be severally hampered, which in turn could have highly negative impacts on its ability to prevent some cases from ultimately resulting in deaths. This criterion is based on numerical information regarding hospital admissions, so unlike other criterion, such as those taking COVID-19 tests, is not based on voluntary participation. However, it only considers situations requiring health service interventions.

Before we investigate the usefulness of DRSA in investigating fairness, let us discuss an illustrative example to understand different steps involved in the DRSA technique. This will help us contextualise our approach and its use within our COVID-19 data analysis case study.

3.2. Illustrative example of DRSA for COVID-19 related data

Table 1 provides synthetic data with ten observations, where each observation has three input readings relating to the number of cases, rate of change (in the number of cases from one day to the next), and the positivity rate. The fourth value can be considered as an output showing the Tier level allocated to that observation. In the DRSA literature, each observation is sometimes referred to as object; and the inputs are sometimes referred to as criteria attributes, and the output referred to as decision attribute.

Within rough set theory, the observation data is usually structured in an information table $S = \langle X, Q, V, f \rangle$, where X is a non-empty finite set of records and Q is a non-empty finite set of attributes observed for each record, such that $q : X \to V_q$ for every $q \epsilon Q$. The V_q is the domain of attribute q, $V = \cap_{q\epsilon Q} V_q$, and $f : X \times Q \to V$ is the information function defined such that $f(x, q) \epsilon V_q$ for each attribute q and record $x \epsilon X$. The set Q is often derived into a subset of criteria attributes $C \neq \emptyset$ and a subset of decision attributes $D \neq \emptyset$ such that $C \cup D = Q$ and $C \cap D = \emptyset$. In our illustrative example, **S** is the decision table shown in Table 1. In this table, X consists of 10 records pertaining to COVID-19 observations, C represents the set of criteria attributes of "Number of Cases", "Rate of change" and "Positivity Rate", while D is a singleton set consisting of only one decision attribute of "Tier". The union of these two sets C and D represent the overall set of attributes Q. Each cell in Table 1 represents the value $f(x, q)$ for its respective record x and attribute q.

The decision attribute $E \epsilon D$ makes a partition of X into a finite number of preference-ordered decision classes $\mathbf{Cl} = \{Cl_t, t \epsilon \{1, \dots, n\}\}$, such that each $x \epsilon X$ belongs to one, and only one, class Cl_t. For example, in Table 1, x_1, x_4, and x_5 all have the decision class of Tier 3, x_2, x_3, x_6 and x_8 all have the decision class of Tier 2, whilst x_7, x_9, and x_{10} all have the decision class of Tier 1. In this way, we obtain three subsets of preference-ordered decision classes.

The dominance relation Δ_p associated with $P \subseteq C$ is defined as: $x \Delta_p y \Leftrightarrow f(x, q) \geq f(y, q)$, $\forall q \epsilon P$, for each pair of observations $x \epsilon X$ and $y \epsilon X$. Here, the preference relation $\geq$ should be replaced with $\leq$ for criteria which are ordered according to decreasing preferences.

To each object $x \epsilon X$, we associate two sets: (1) *the P-dominating set* $\Delta_p^+\{x\} = \{y \epsilon X : y \Delta_p x\}$ containing the objects that dominate x, and (2) *the P-dominated set* $\Delta_p^-\{x\} = \{y \epsilon X : x \Delta_p y\}$ containing the objects dominated by x. For example, in Table 1 the second observation (x_2) shows that there were 92 (positive) cases, while the rate of change was 2.45, and the positivity rate was 7.89%. When comparing this observation with the first observation (x_1), we can see that all three input values of x_2 are lower than their respective values for x_1. Therefore, we can say that x_2 dominates x_1. In this way, we can compare the inputs of each observation with every other observation to determine its dominance. Table 2 summarises this dominance relationships for all observations. For example, the fifth observation (x_5) is clearly dominated by four other observations, those of x_4, x_8, x_9, x_{10}. Note that the table also includes x_5 itself as, by definition, the dominance relationship also includes equality.

After processing all the dominance relationships in Table 2, we also need to process the outputs, which are essentially the Tier values assigned to each observation. As we know that tiers are ordered categorical values, we can group these tiers into unions by enumerating all possible combinations. This is summarised in Table 3 where we have four unions of "at most T1", "at most T2", "at least T2" and "at least T3". The unions involving "at most" are also known as *downward* unions, and those involving "at least" are known as *upward* unions.[2] Formally, this information is a collection of upward unions $Cl_t^{\geq}$ and downward unions $Cl_t^{\leq}$ of classes as: $Cl_t^{\geq} = \bigcup_{s \geq t} Cl_s$ and $Cl_t^{\leq} = \bigcup_{s \leq t} Cl_s$ respectively. The assertion "$x \epsilon Cl_t^{\geq}$" means that "x belongs to at least class Cl_t" while the assertion "$x \epsilon Cl_t^{\leq}$" means that "x belongs to at most class Cl_t".

Now that the inputs and outputs are processed separately, the next step would be to link the two tables together. This is usually done by inducing certain, possible, and approximate rules using rough-set based algorithms like DOMLEM [34] or VC-DOMLEM [35]. The general structure of induced certain decision rules is as follows:

IF *antecedent*, THEN *At Most* Cl_t

IF *antecedent*, THEN *At Least* Cl_t

in which the *antecedent* specifies the conditions on one or more criteria, and the decision part (the *consequent*) specifies an assignment to one or more decision classes. The set of rules induced using DOMLEM for this example are shown in Table 4, where each rule has an antecedent and a consequent. The antecedent is the set of conditions that are required for this rule, and the consequent is the resulting union. Out of the nine rules extracted for this example, one can see that Rules 2, 4, 6 and 8 (see grey shaded rules) have only one condition to satisfy, and hence we can say that the rule length is 1. Consequently, the remaining

Table 1
Illustrative example to demonstrate the use of DRSA for COVID-related data.

Observation	Number of cases	Rate of change	Positivity rate	Tier
x_1	195	2.48	8.05	3
x_2	92	2.45	7.89	2
x_3	237	−2.74	8.94	2
x_4	515	2.82	1.43	3
x_5	528	7.54	5.3	3
x_6	434	1.65	5.41	2
x_7	143	−3.15	8.01	1
x_8	75	3.2	5.25	2
x_9	269	2.33	1.71	1
x_{10}	131	3.28	1.03	1

Table 2
Summary of dominating and dominated observations.

Observation	Dominating	Dominated
x_1	1	1,2,7
x_2	1, 2	2
x_3	3	3, 7
x_4	4, 5	4
x_5	5	4, 5, 8, 9, 10
x_6	6	6
x_7	1, 3, 7	7
x_8	5, 8	8
x_9	5, 9	9
x_{10}	5, 10	10

Table 3
The unions of classes for the illustrative example.

Union	Observations that are part of the union	Total
At most T1	7, 9, 10	3
At most T2	2, 3, 6, 7, 8, 9, 10	7
At least T2	1, 2, 3, 4, 5, 6, 8	7
At least T3	1, 4, 5	3

five rules are of rule length 2 as all of them need two conditions to be satisfied.

The first four rules in Table 4 have *consequents* involving "at most" while the remaining six rules involve "at least" in their *consequents*. The rules related to "at most" are termed as downward rules, whilst the rules related to "at least" are termed as upward rules. The last two columns in Table 4 show the support and strength for each rule induced from the data. The support refers to the number of observations that adhere to this rule, while the strength refers to the ratio of observations that met the rule against the observations that only met the first part of the rule (i.e. *antecedent*). For example, for the rule in row 1 of 4, there are three observations under the *antecedent*, of which 2 adhere to the rule (hence a support of 2). The strength for the rule is therefore, 2/3 = 66.67%.

Now that we have shown how DRSA can be used to extract If-Then rules about COVID-19 restriction data, the next section demonstrates how fairness, within domains such as the tired restrictions system, can be investigated.

3.3. Were the tiered restrictions fair?

The government claimed that the tiered system would result in greater transparency, through more fine-grained and area specific data driven restriction decisions. However, people questioned the integrity and fairness of the system from the start. For example, Manchester was forced to enter into Tier-3 without sufficient evidence to appease local policy makers and the local community, with The Mayor of Manchester responding to the city being put into Tier 3 when London remained in Tier 2 by saying "*That is clearly unfair. It gives the impression that jobs outside of London are not worth the same*" [36]. Similarly, when the area of Yorkshire was ordered to remain in Tier 3, local politicians raised objectives that, to them, the decision-making process was unfair [37].

To investigate the issue of fairness in assigning tiered restrictions, it is important to first understand the concept of fairness. When talking of regional fairness and equality, most research has focused on the distribution of resources. For example, fair distribution of food [38,39], assigning public health resources [40,41], or public transport [42] between different regions. In other cases, this concept is discussed for allocating budgets (finances) in a fair and equitable manner [43]. One can argue that budget is also a type of resource distribution where the resource is a limited monetary value. Both issues involve a direct (or indirect) competition as different regions are competing against each other for obtaining (or maximising) their share of finite resources. However, the idea of assigning different levels of restrictions does not fit into these two types, as the assignment of tier to any region is independent of how other regions are performing (or which state they are in). Thus the concept of fairness in a tiered allocation system is potentially different from the fairness in distribution of resources, where allocating more resources to one region implies the other regions will get less. Therefore, assigning heavy restrictions in one region has no dependence on the restrictions assigned to other regions.

Although assigning restrictions to one region is independent of other regions, people have a natural tendency to compare the restrictions imposed on them with those in other regions [44]. Any preferential treatment to one region over others might not be welcomed, and might end up in some form of protest. This protest can take a form of doing (or not doing) actions that lead to breaking the rules [45,46]. For example, in the case of allocating tiers, people in higher tiers still have the ability to break the rules, and act as if they were in a lower tier. Therefore, when they consider it unfair, it is more likely to see people breaking the rules, which makes these decisions critical for regional peace and stability.

3.4. The issue of fairness and the North–South divide in England

In England, the North–South divide is a term that refers to the socioeconomic differences between Southern and Northern parts of England [47]. A recent report by the Institute for Public Policy Research [48] states that the South of England consists of one-third of the UK population yet accounts for 45% of its economy and 42% of its wealth. Considering these numbers, it is important to investigate the possibility of this North–South divide penetrating in the COVID-19 related policies as well.

Before discussing this further, it is important to first define the boundaries separating the North and South. In terms of creating boundaries between North and South, a widely accepted belief is that the North consists of the regions of the North East, North West, Yorkshire and The Humber, East Midlands, and West Midlands[3]. On the other

[2] Please note that the unions of "at least T1" do not make sense as all observations will be at least T1.

[3] However, in some studies, the Midlands have also been considered a separate geographical entity from the North and the South [49].

Table 4
Illustrative example to demonstrate the use of DRSA for COVID-related data.

No.	Antecedent	Consequent	Support	Strength
1	(Number of cases <= 269) and (Positivity rate <= 1.71)	At most T1	2	66.67
2	(Rate of change <= 2.45)	At most T2	5	71.43
3	(Number of cases <= 434) and (Positivity rate <= 5.41)	At most T2	4	57.14
4	(Number of cases <= 131)	At most T2	3	42.86
5	(Number of cases >= 195) and (Rate of change >= 2.48)	At least T3	3	100.00
6	(Number of cases >= 515)	At least T3	2	66.67
7	(Rate of change >= 2.82) and (Positivity rate >= 1.43)	At least T2	3	42.86
8	(Number of cases >= 434)	At least T2	3	42.86
9	(Rate of change >= 1.65) and (Positivity rate >= 5.41)	At least T2	3	42.86

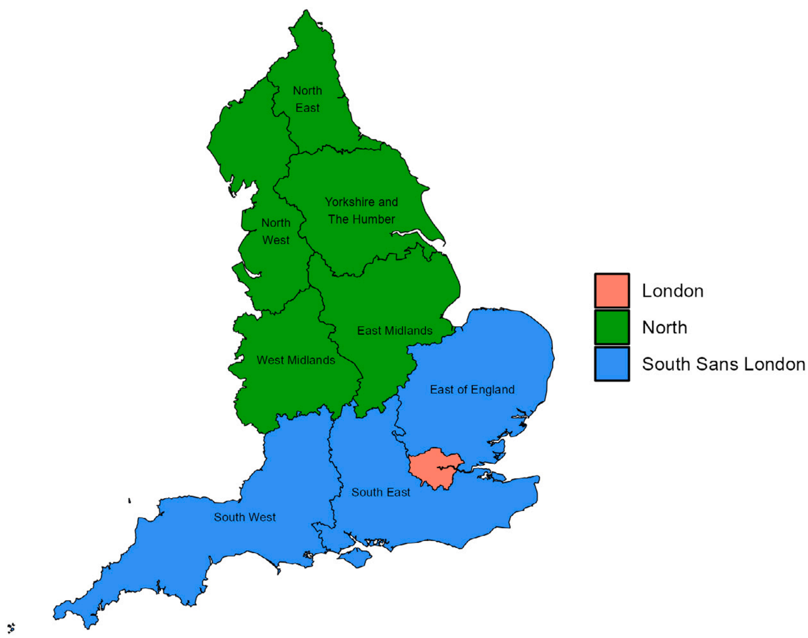

Fig. 1. England regions.

hand, the South consists of the East of England, London, South East, and South West. Although London is considered part of South, we are interested in exploring how it has been argued that London has its own unique dynamics due to its different socioeconomic demographics. Therefore, for research purpose, we will also analyse London separately due to its unique data properties (explained in Section 5.1).

For the purpose of our study, we will therefore consider three geographical area categories of (i) North, (ii) South Sans London, and (iii) London. These three areas can be seen geographically in Fig. 1 where each of the nine regions have been coloured green, red or blue. The green regions constitute North while the blue regions constitute South without London, and London is shown in red colour.

A recent review of the use of data analytics in COVID-19 concluded that it has been successfully employed in the healthcare sector [50]. We argue that the use of data analytical techniques can also help identify any anomalies or inconsistencies in the assignments of the tiered restrictions. Although the assignment of tiers is not a zero-sum game – where one person's gain is another person's loss – any inconsistency in the assignments can have serious implications as people would compare their level of freedom (or level of safety) to the other regions, and these comparisons might dissuade some of them from sticking to the rules [51].

In this context, investigating fairness amongst these regions is an important step towards addressing the concerns from various stakeholders. The next section outlines the methodology of our approach to investigate fairness.

4. Methodology

The UK government's tiered allocation assignments were based on data pertaining to the prevalence and risk posed by COVID-19 in different parts of England. In the illustrative example earlier, a set of rules was derived from a set of observations. Comparisons between different geographical areas can be performed by performing this process of rules derivation for different areas of England, utilising only the observations specific for each area, to generate a separate set of rules for each area. From the extracted sets of rules for each geographical area under comparison, we have the opportunity to assess the overall fairness of the tiered system. To explore the fairness of the tiered approach in relation to the North–South divide, we can look to generate three separate rule-sets from the observations from (i) the North, (ii) the South sans London, and (iii) London, as defined earlier and shown in Fig. 1.

The process pipeline of collecting and analysing data is shown in Fig. 2; which includes data acquisition, pre-processing, and the

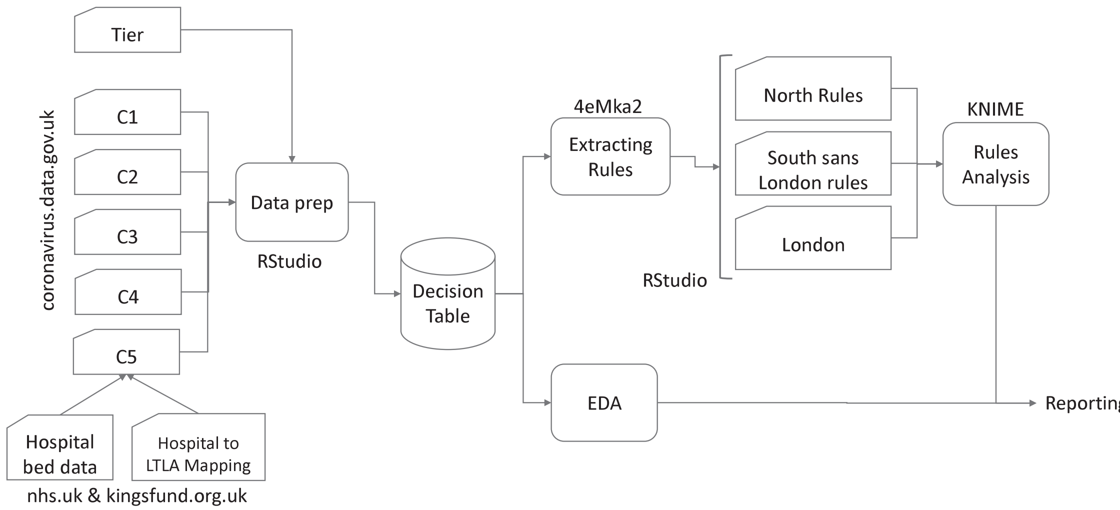

Fig. 2. The data extraction, pre-processing and analysis pipeline of our approach.

extraction and analysis of DRSA rules. Before detailing and discussing each of these pipeline stages in detail, we first briefly contextualise our approach within the approach space. Our approach focuses on investigating fairness through comparisons of separate sets of rules, derived from disjoint data slices — as opposed to focusing on predictions for future unseen data and thus classification prowess. Therefore, the proposed approach is not looking to develop, evaluate or compare with classification techniques. Moreover, bireduct analysis, with goals distinct from our approach, may be utilised in combination with our approach to explore, for exmaple, if desired, reduction before slicing data into disjoint slices. Bireducts within our approach are discussed in Section 4.5.

4.1. Data acquisition

To explore the UK's tired system requires the collection of observation data relating to the set of government criteria over the duration of the tired allocation system. For our approach data was collected from various different disparate sources.

For the different areas of England, data was obtained pertaining to the UK government's set of five criteria discussed earlier [33], along with the allocated tier values. Table 5 provides the list of data sources used to collect data on these five criteria.

England is geographically broken up into 9 Regions which are further broken down into a set of Lower Tier Local Authority (LTLA) areas [52]. Data relating to COVID-19 at the LTLA level of granularity can be collected. Therefore, we curate a data-set consisting of Tier value and criteria calculations, for each LTLA, for each day that the Tiered system was in place. We briefly outline the data acquisition process for each criterion below.

4.1.1. C1: Number of daily cases

The website www.coronavirus.data.gov.uk is a UK government website that provides official data relating to coronavirus (COVID-19). The data denoting the number of positive cases which were reported in each LTLA region every day can be obtained from the site.[4] Although the data is provided by age demographics, the overall number of cases were derived by totalling the cases across the set of age ranges.

As a common practice [53], we calculated a 7-day rolling average to alleviate discrepancies in the reporting velocity at different days of the week. The effects of this process are shown in Fig. 3, where, the figure on the left clearly shows that the number of reported cases around weekends is invariably lower than the cases on other days. The figure on the right shows the rolling average where this issue has been addressed.

[4] From the site's Supplementary downloads.

Table 5
Data sources for government-defined criteria set.

Criterion	Description	Data source(s)
C1	Number of daily cases (all ages)	coronavirus.data.gov.uk
C2	Number of daily cases in over 60 s	coronavirus.data.gov.uk
C3	Rate of change of daily cases	Derived from C1 data
C4	Positivity rate	api.coronavirus.data.gov.uk
C5	Pressure on the NHS	(1) api.coronavirus.data.gov.uk (2) www.kingsfund.org.uk (3) github.com/epiforecasts
Decision variable	Tier value	Parliament Legislation documentation

4.1.2. C2: Number of daily cases in those aged 60 plus

Recall that the number of cases collected for C1[5] are broken down by age demographic ranges, therefore, we derived C2 values as the sum of age demographic ranges 60 and over (i.e. for each LTLA for each day). As with C1, we then calculated a 7-day rolling average for C2 as well.

4.1.3. C3: Rate of change in number of daily cases

The rates of change in the number of cases (from one day to the next) can be determined from the data obtained for C1. A 7-day rolling average was calculated for these rate of change values for each LTLA region for each day.

4.1.4. C4: Positivity rate

The positivity rate of cases denotes the percentage of tests taken over a given time period that are returning positive results. The positivity rate for each LTLA region for each day can be sought. Such data can be retrieved via the UK government's COVID-19 API.[6] Through utilising

[5] From the coronavirus.data.gov.uk site.

[6] api.coronavirus.data.gov.uk - providing API calls for a range of COVID-19 related metrics and levels of geographical granularity.

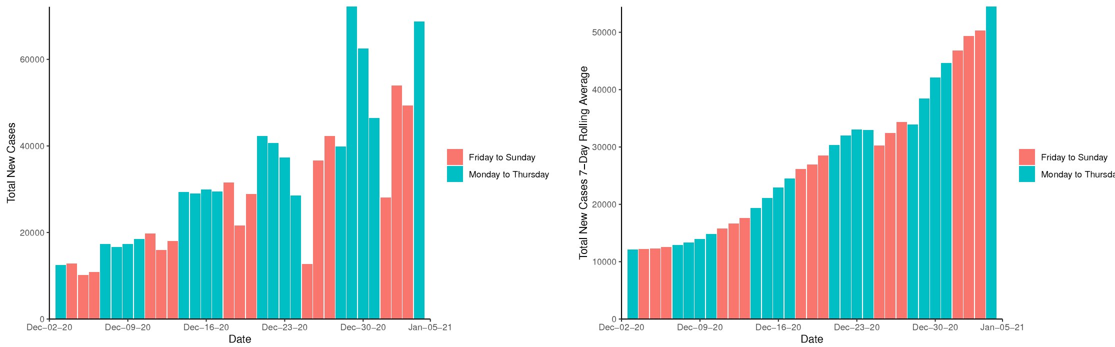

(a) Raw Data (b) 7-Day rolling average

Fig. 3. Raw data and 7-day rolling average data.

the API making a set of call requests, for each LTLA region, we obtained the positivity rate, for each LTLA region for each day, which again we processed into 7-day rolling average values.

4.1.5. C5: Pressure on the National Health Service (NHS)

The criterion of pressure on the NHS looks to measure the amount of pressure being placed on the local medical services, in terms of their ability to be able to cope with the COVID-19 cases prevalent within a given area. The UK government utilised this criterion in their tiers allocation system, however, the pressure on NHS cannot be quantified directly at LTLA level. To estimate the pressure, we relied on the number of beds occupied by COVID-19 patients at NHS Trusts level. The NHS service is broken geographically across England into a number of Trusts, 223 in total[7] (which each operate as a somewhat automated unit).

Due to lack of government transparency in quantifying the pressure on NHS, the Trust level values were collected and mapped to LTLA level to quantify the pressure on the NHS, as there was a need to align the Trust level data with the LTLA level data. This mapping (performed in R) is explained below.

Using the UK government's COVID-19 API provides the number of beds occupied by COVID-19 cases (for each trust for a given day). NHS trusts across England vary significantly in size (and thus levels of capacity) therefore, comparisons between such values on their own may be misleading. For example, one trust may have double the number of beds occupied with COVID-19 patients than another but still be under less pressure due to being three times the size. The government documentation regarding the tiers system remains opaque in terms of how such information was exactly utilised to help make decisions. However, we can utilise various other publicly available data to derive measures of the relative pressure on the NHS.

Each trust has multiple hospital sites for which data relating to the total number of available beds for each trust can be sought.[8] We already know the actual number of beds occupied by COVID-19 patients so therefore, we can estimate a representative ratio of occupied beds in relation to the overall number of beds in each trust.

Recall these values are available for each trust, however, to perform these analysis, we need all values at the LTLA level. A trust may serve more than one LTLA region, and an LTLA region might be served by multiple trusts. Therefore, in this way, there exists a many to many mapping between trusts and LTLAs.

One way to resolve this issue is to calculate probabilistic estimates through mapping the information at trust level to LTLA level. To achieve this we utilised the Trust to local authority mapping provided by Epiforecasts.[9] This package provides probabilistic mapping estimates that can be used to estimate COVID-19 hospital admissions at an LTLA level. From utilising this mapping on our data we obtained a ratio based pressure value for each LTLA region for every day, and thus at the same level of granularity that we have for the other four criteria. Whereas C1 to C4 are derived from deterministic distinct values, C5 values are estimated due to the fact that the areas covered by hospitals do not align with the areas covered by local regions. Therefore, although based on various official government data sources, C5 values are estimated through probabilistic calculations [54]. The implication of this, therefore, is that there is an inevitable element of approximation within these calculations.

4.1.6. Decision variable: Tiers data

Date pertaining to tier values (1, 2, 3 or 4) that different parts of England were placed under, and subsequent changes, were announced in the houses of Parliament and published in official legislation documentation by the government [55]. From this we were able to define the Tier that each LTLA region was in for every day that the tired system was in place.

4.2. Data processing and integration

From the acquisition and format alignment processing of data from these various sources making up the set of 5 criteria we then integrated the data to create an overall data-set. In the data-set each row (observation) denotes the data for the set of 5 criteria and the tier pertaining to an LTLA and Date pair, with additional information added to each observation regarding its membership of either Northern England, Southern England sans London, or London. The data-set consists of 10827 observations.[10]

The data-set was then utilised to gain understanding about the allocation of tiers, first with an exploratory data analysis, and then using the DRSA rules analysis approach. These two approaches are described in the next subsections below.

4.3. Exploratory data analysis

For better data understanding, we performed some statistical analysis of the data-set. This included analysing the spread of decision variable values, across the range of dates, and between the different geographical areas. We then analysed the relationships between and within the criteria set and tier values allocated. The results for exploratory data analysis are discussed in Section 5.1.

[7] www.kingsfund.org.uk/audio-video/key-facts-figures-nhs

[8] www.kingsfund.org.uk

[9] Available as a developer package on GitHub at - github.com/epiforecasts/covid19.nhs.data.

[10] The enriched data-set for this study is publicly available on github here: https://github.com/prioritization/DRSA_Covid19.

4.4. DRSA ruleset preparation, extraction, and comparisons

DRSA is usually tasked with taking a data-set of information and determining patterns within the data through the extraction of a single set of rules, as shown in Fig. 4(a). In such a scenario, the set of extracted rules can then be utilised to gleam insights about the data, or be utilised for making predictions on future unseen data. However, for our fairness investigation, we take a novel approach to utilise DRSA to extract multiple separate rule-sets, and then perform comparison analysis between the rule-sets. These stages are shown in Fig. 4(b). To investigate the tiers allocation system, first we segment our data set into three separate subsets representing the North, the South sans London, and London respectively. Before and/or after segmenting out data into separate slices, and performing separate rule extraction, if desired, bireduct analysis can be performed, to identify if reductive subsets exist within the data.

4.5. Bireducts analysis

Bireduct analysis to identify any bireducts within our data could be both informative, in terms of helping to identify the most important criteria (and undiscriminating criteria), and provide potential improvement in rule extraction processing time and robustness. Within out approach, bireduct analysis can be performed upon our whole dataset before sub dividing it into separate slices. Then any bireduct found, if desired and selected for use, could then be propagated when we subsequently subdivide our dataset into slices. For our COVID-19 5 attribute dataset bireduct analysis (utilising the DRSA tool 4eMka)[11] yielded no bireducts. Given each data slice contains the full set off attributes and a set of objects that are non-overlapping with any of the other data slices' objects. Therefore, it may be the case that separate bireducts exist at slice level. So, within our approach, if desired, bireduct analysis can be performed after sub dividing a dataset into separate slices, to then perform bireduct analysis separately on each data slice. However, from the perspective of looking to compare rule sets extracted from separate data slices, through identifying similar rules present in multiple data slices, bireducts at slice level may result in less wide-ranging comparisons being possible. For our COVID-19 dataset, sliced into three separate subsets representing the North, the South sans London, and London respectively, bireduct analysis yielded no bireducts in any of the three subsets. Although, for our COVID-19 data, with its relatively speaking small set of attributes, the analysis revealed no potential bireducts, we wish to emphasis that bireduct analysis can be a complementary part of our approach's operational stages.

4.5.1. Multiple rule-sets extraction

With out dataset segmented into three separate subsets representing the North, the South sans London, and London respectively.

Then, we perform separate independent rule extraction for each data slice. The data slices are processed using R scripts to generate information system files (.isf) readable by the DRSA tool 4eMka2.[12] These .isf files can then be processed by 4emka2 to generate three separate sets of extracted rules, providing us with corresponding outputs in the form of rules files (.rls). This process of separate data slices and rule-sets is shown in Fig. 4(b). The output .rls files are then parsed using R scripts for the rule-sets Comparisons.

As discussed in Section 3.2, the extracted rules may have different strength and support. These metrics can be utilised to define minimum threshold values for filtering rules. Such thresholds can be defined before the extraction process, or alternatively, they can be refined after extracting all rules. Taking the former approach will make the process computationally efficient whilst the latter provides more flexibility, to dynamically alter the thresholds without having to re-run the extraction process.

When comparing rules extracted from different data slices, the same threshold values should be applied across these slices. However, the support metric is size-dependent and the data slices from different regions are not necessarily exactly equal in size. Therefore, there is a need to transform the support metric into relative support, that is, the ratio of observations that support the rule relative to the total number of observations in the data slice. In this way, rules from different rule-sets could be filtered in a consistent manner.

4.5.2. Multiple rule-sets comparisons

The tiered restrictions applied to the North, the South sans London, and London can be compared using the three separate rule-sets. Rationally speaking, the data driven allocation of tiers should be indifferent to the region it represents, and should depend on the available data on concerned criteria. If this is not the case, we can consider it an inconsistency in the tiers allocation system. Therefore, identification of inconsistencies amongst different rule-sets may suggest unfairness in terms of restriction decisions. To compare different rule-sets, we identified where these rule-sets shared a rule. Shared rules are ones that relate to the same set of criteria in the antecedents, and the same tier value in the consequent.

For example, given a rule present in two different rule-sets, as shown in Table 6, we can compare the criteria values to identify any inconsistency between the regions. In Table 6, we observe the thresholds for both criteria that result in at least Tier 2, are lower for the North than for London. This suggests that London would be able to endure higher number of cases, and higher positivity rates, before it was treated the same as the North.

Such analysis allows us to compare separate rules, extracted from separate geographical areas, which in turn indicates whether the regions are being treated equally. Our approach collates all comparable rules which can then be presented to the user in tabular and graphical format for appraising equal treatment.

So far, we assumed that for each data slice, only one rule is possible to be extracted for given antecedent criteria and consequent. However, in practice, rules are found with the same set of criteria in the antecedent and the same consequent, but with different thresholds for the criteria at a different strength value. In such cases, we can have different criteria thresholds for different strength values. Therefore, comparing such rules across the data slices is not a straightforward task. In order to make a one-to-one comparison of such rules, we first need a representative rule from each data slice, as discussed next.

4.5.3. Obtaining representative rules for comparison

As mentioned in the previous section, we may obtain multiple rules within a data slice with the same set of criteria in the antecedent and the same consequent, and to make a one-to-one comparison, we need a representative rule. This can be achieved by aggregating such rules, for example, using weighted averaging where the weights are proportional to the rule strength. To perform such aggregation, we calculated the weighted average criteria values for the rules in such a group, weighted with respect to each rule's strength. In this way, each rule strength proportionally influences the amount of impact each rule in the group has upon the aggregate rule. The weights for aggregating these rules can be obtained in a number of ways, for example, Ordered weighted average operators [56], or other prominent aggregation operators such as outlined in [57] could alternatively be utilised in our appraoch.

Table 7 shows five rules from the rule-set for London. Here, all five rules share the same set of antecedent criteria and the same consequent, but have varying rule strength values. These can be aggregated into a single aggregated rule, becoming the single representative rule for

[11] 4eMka2 is a tool developed by the Laboratory of Intelligent Decision that implements the DRSA method Support System of the Institute of Computing science, Poznan University of Technology, idss.cs.put.poznan.pl.

[12] See footnote 11.

(a) Extracting a single set of rules from all data (b) Extracting separate rules from separate slices of data

Fig. 4. Flow diagrams to differentiate our proposed approach from conventional utilisation of DRSA.

Table 6
Example aggregated rule.

Region	Antecedent	Consequent
London	(Number of cases => 177.40) and (Positivity rate => 25.13)	At least T2
North	(Number of cases => 146.04) and (Positivity rate => 19.09)	At least T2

Table 7
Example LONDON Rules sharing matching criteria in the Antecedent and matching Consequent.

Antecedent	Consequent	Strength
(Number of cases >= 160) and (Positivity rate >= 13.8)	At least T2	27.76
(Number of cases >= 166) and (Positivity rate >= 13.7)	At least T2	33.5
(Number of cases >= 180) and (Positivity rate >= 29.6)	At least T2	62.67
(Number of cases >= 180) and (Positivity rate >= 28.2)	At least T2	73.02
(Number of cases >= 184) and (Positivity rate >= 27.6)	At least T2	77.76

Table 8
Example The NORTH Rules sharing matching criteria in the Antecedent and matching Consequent.

Antecedent	Consequent	Strength
(Number of cases >= 142) and (Positivity rate >= 12.1)	At least T2	28.5
(Number of cases >= 140) and (Positivity rate >= 13.1)	At least T2	34.5
(Number of cases >= 151) and (Positivity rate >= 25.3)	At least T2	65.4

London as shown earlier in Table 6. Similarly, given the three rules from the North rule-set, as show in Table 8, they can be aggregated into a single rule for the North, as shown earlier in Table 6.

After performing such aggregation, we obtain a single representative rule for each such group of rules, that can then be compared with like-for-like representative rules obtained from other slices.

5. Analysis and results

In this section, after some initial exploratory data analysis, we then discuss our approach for investigating fairness in tiers allocation by comparing the representative DRSA rules from each region.

5.1. Exploratory data analysis

A set of exploratory data analyses were carried out to gain data understanding before investigating further. These are described and discussed below.

5.1.1. Spread of tiers values

We explored the distribution within the data of Tier values across the time period. Fig. 5(a) shows these distributions for the whole data-set, where it can be seen that a majority of areas were in Tier-2 in early November but moved into Tier-4 by mid December. Tier-1, however, remains infrequent in the whole data-set, with only 0.5% of observations, (the other tier values percentages of observations in the whole data-set are Tier-2, 37.9%, Tier-3, 37.8%, and Tier-4, 25.9%).

Fig. 5(b) shows that the majority of the Northern regions started and stayed in Tier-3 while the Southern regions (sans London) started in Tier-2 (see 5(c)) and stayed in Tier-2 until towards the end of December. London followed a similar but even more pronounced pattern than the South, as visible in 5(d).

The analysis of Fig. 5 highlights the imbalance of the distribution of the Tier values across different regions. This imbalance might impact the coverage of DRSA rules that can be generated, for example, rules about Tier-1 might not be easy to derive due to the sparsity of observation data with Tier-1.

Observation 1. *There is an imbalance of the distribution of the Tier values.*

5.1.2. Analysis of relationships between Tier values and criteria values

Next, we analyse potential correlations between the criteria values and Tier values, across the different geographical areas. These correlations are shown in a set of tables provided in Table 9. Looking at Table 9a - showing pairwise correlations among criteria and Tiers within the whole data-set - one can see that there is a strong positive correlation between C1 and C2 (i.e. 0.94). Although this is true for all four tables, the correlation in the data from North (see Table 9b) is slightly stronger than in the data obtained from The South Sans London and London (see Tables 9c and 9d). However, regarding the other criteria the North has significantly lower correlations amongst the set of five criteria (and Tier values) than The South Sans London and London.

Observation 2. *The number of cases in age group 60+ (C2) is highly correlated with the overall number of cases (C1)*

Observation 3. *Data for the North has significantly lower correlations amongst the set of five criteria (and Tier values)*

A further insight from the correlation data is that C3 (Total New Cases Rate of Change) has a weak correlation with the Tier value (as well as to the other criteria). Conversely, C4 (Positively Rate) has the

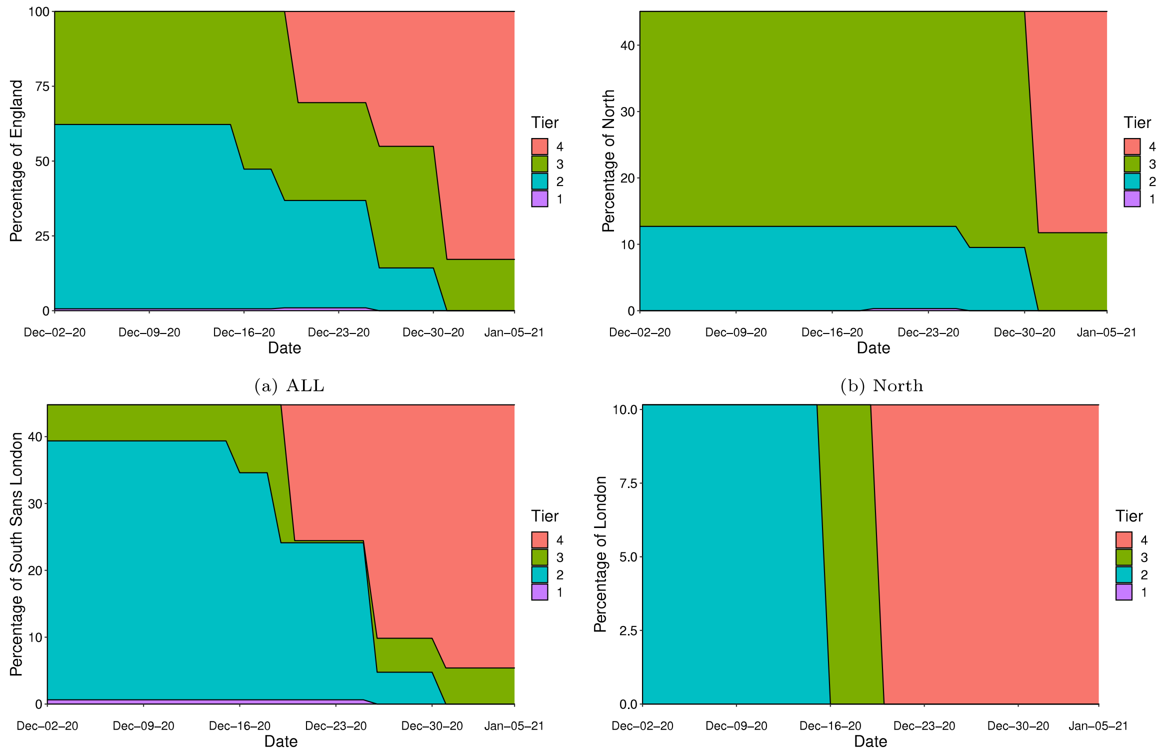

(a) ALL (b) North

(c) South sans London (d) London

Fig. 5. Density plots of the distribution of Tiers across the data range for different geographic areas.

Table 9
Correlation table.

	C1	C2	C3	C4	C5	Tier
C1		0.94	0.01	0.74	0.61	0.49
C2			−0.02	0.74	0.54	0.53
C3				−0.06	−0.06	−0.10
C4					0.56	0.71
C5						0.36
Tier						

(a) Correlations for overall data

	C1	C2	C3	C4	C5	Tier
C1		0.96	0.12	0.52	0.26	0.32
C2			0.11	0.53	0.25	0.34
C3				0.11	−0.03	−0.00
C4					0.25	0.65
C5						0.28
Tier						

(b) Correlations for North region

	C1	C2	C3	C4	C5	Tier
C1		0.93	−0.09	0.80	0.68	0.58
C2			−0.13	0.80	0.62	0.62
C3				−0.18	−0.10	−0.16
C4					0.61	0.75
C5						0.40
Tier						

(c) Correlations for South region Sans London

	C1	C2	C3	C4	C5	Tier
C1		0.93	−0.27	0.78	0.41	0.62
C2			−0.10	0.76	0.44	0.64
C3				−0.16	−0.11	−0.12
C4					0.52	0.76
C5						0.45
Tier						

(d) Correlations for Greater London

strongest correlation with the Tier value for the overall data (as well as for each of the three geographical breakdowns). This suggests that the 'Total New Cases Rate of Change' has little influence upon Tier values, whilst the 'Positivity Rate' does.

These relationships are shown visually through Box-plots in Fig. 6. Here, Box-plots are shown for C1, C3, C4 and C5. C2 is omitted due to the strong positive correlation previously identified between C1 and C2, which results in very similar box-plots for C1 and C2. In each of these plots, the criteria values are grouped according to the Tier values, and a statistical summary is shown for North, South Sans London, and London. For C1, C4 and C5, the Box Plots show that London was invariably assigned to be in a lower Tier for values that would have resulted in higher Tier assignment for the North and the South Sans London. For example, for C1, the median value for London being assigned Tier-3 is 239.43, where as conversely, the median value for North and South being assigned to Tier-4 is 77.14, and 115.50000 respectively. This suggests London was given different treatment, perhaps showing that is was being favoured for economic activities at the expense of public health; suggesting implicit additional criteria exist within the Tiers system.

Observation 4. *For C1, C4 and C5, London was allocated to be in Tier-2 or Tier-3 for the values that put other geographical regions in Tier-4.*

From these box-plots, for C3 (Total New Cases Rate of Change) shown in Fig. 6(b), there is apparently no relation between the criterion's value and the assigned Tier. Whereas, for C4 (Positivity Rate), a clear positive relationship is visible with the Tier value (albeit at different rates for the different geographical regions). This reiterates what the correlation analysis highlighted, namely that C3 appears to have a weak or no relationship with an assigned Tier value where as C4 does.

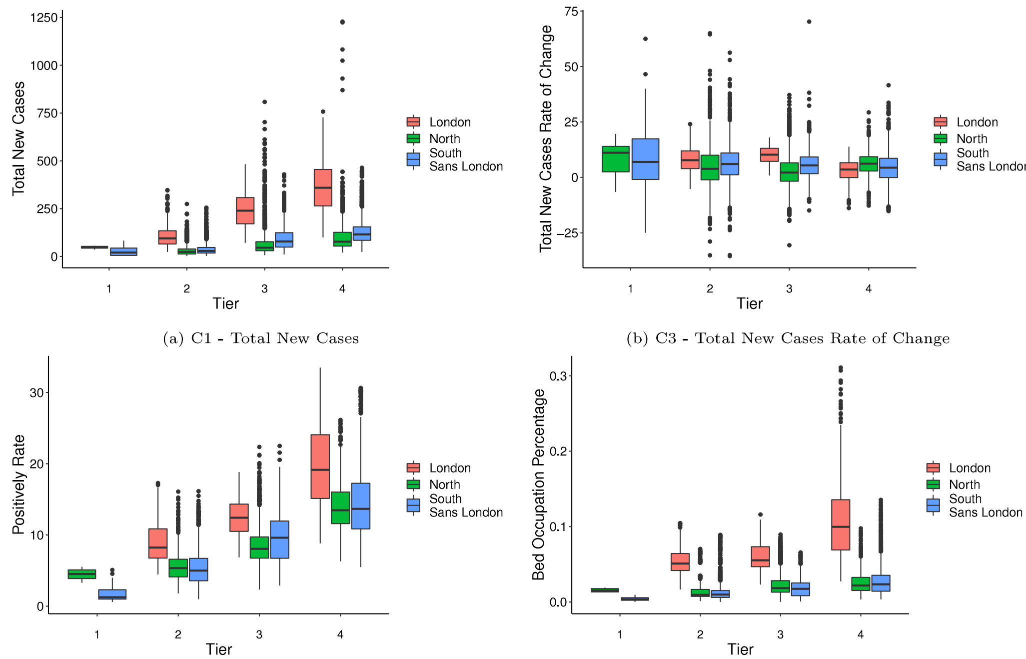

(a) C1 - Total New Cases
(b) C3 - Total New Cases Rate of Change
(c) C4 - Positivity Rate
(d) C5 - Pressure on the NHS

Fig. 6. Box plots.

Observation 5. *The Rate of change (C3) has a weak or no relation with Tiers allocation, whilst C4 (Positivity Rate) appears to have a strong relation with Tiers.*

So far, each criterion has been investigated individually, without analysing possible interrelationships amongst them and the assigned Tiers. Also, to investigate the fairness, it is more appropriate to investigate the combinations of tiers, for example, investigating geographical area in "at least" Tier-3 or "at most" Tier-2. Such analyses can be carried out using DRSA, as discussed in the next subsection. The DRSA models generated from such analysis can also be used to help us sort unseen data, and in turn, facilitate predictive analysis for evidence-based decision making.

5.2. DRSA rule analysis and comparison

Although the DRSA has been widely used in practical applications, these applications have focused on a single set of rules extracted from a data-set. On the contrary, in this research, we focus on generating multiple sets of rules created from different data-sets, and then comparing these rule-sets for gaining further insights.

The rules were extracted using the 4emka software tool using the rule strength as a filter to include/exclude rules for comparison. Keeping the rules with lower strength may reduce the quality of information for comparing different geographical regions. However, on the other hand, filtering out these rules might result in having sparse information which could be insufficient for comparative analysis. This can be seen as a quantity-quality trade-off problem and finding a balance between these two objectives is an important issue to address. Recall the lack of coverage of our data-set shown in Fig. 5(a). This already constrains our analysis due to a limited number of rules having higher confidence values.

Fig. 7 shows the results of an experimental trade-off analysis carried out by changing the threshold level of confidence which can help determine a trade-off between confidence level (i.e. quality) and the number of comparable rules (i.e. quantity). As shown in this figure, the number of comparable rules decrease significantly as we increase the threshold of acceptance for confidence. Therefore, in this case study, we used the minimum rule strength of 25% to keep sufficient information required to compare different regions. However, we believe the selection of this threshold is both (1) context-dependent, where other input data sets might result in alternative thresholds to be chosen, and (2) User dependent, subjective to different user dispositions regarding the trade-offs involved.

Instead of listing all the rules in a traditional if-then format, we display these rules in a table in order to better compare the three regions (with respect to different outcomes). The extracted upward rules are summarised in Table 10, while the downward rules are summarised in Table 11. The extracted upward rules are visualized in Fig. 8, and the extracted downward rules are visualized in Fig. 9 One may argue that the upward rules are about restricting regions from entering a lower Tier and so could be considered more relevant from health policy perspective, where as downward rules restrict the regions from entering a higher tier and so are more concerned with the economic considerations.

Taking the first criterion C1 (number of new cases) with consequent "tier at least 2" in the upward rules table (Table 10), the extracted rules suggest that London is at least in Tier 2 for values greater than 297.14, whereas the North is at least in Tier 2 for values greater than 62.04. This implies that the thresholds for London are relatively more relaxed than the North. For example, if the number of new cases is over 62.04 and below 297.14, the extracted rule suggests that North should not be in Tier-1, whilst London might still be in Tier-1. The ratio between these values is 1:4, suggesting London's cases can be almost five fold more before it is treated the same as the North. Note that here there was no such rule extracted for the South to compare with the North and London.

Observation 6. *The thresholds for upward rules give evidence of relaxed rules for London for C1*

Focusing on the second criterion of C2 (number of cases in those aged 60+), DRSA has extracted rules for London and the South, however, no rule was extracted for the North. As observed for C1, the threshold for London (i.e. 33.71) are also more relaxed than the threshold for the South (i.e. 14.19). The threshold for London is more than

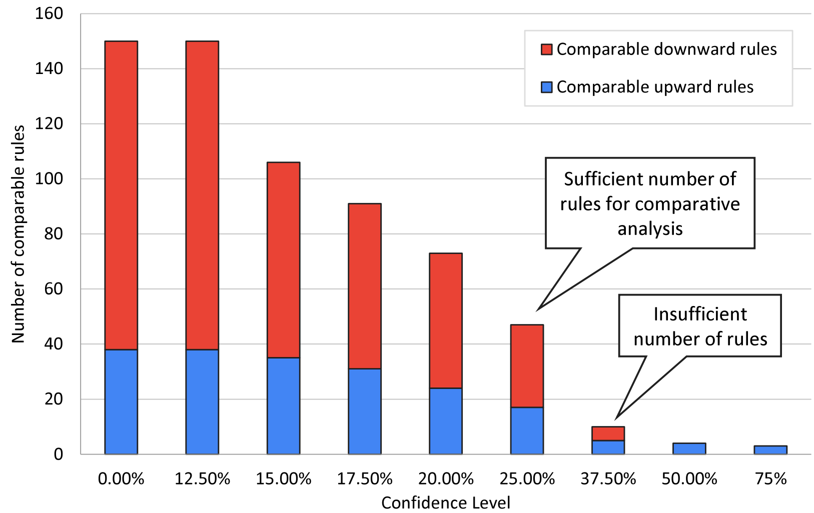

Fig. 7. Increasing the confidence threshold reduces the number of possible comparisons.

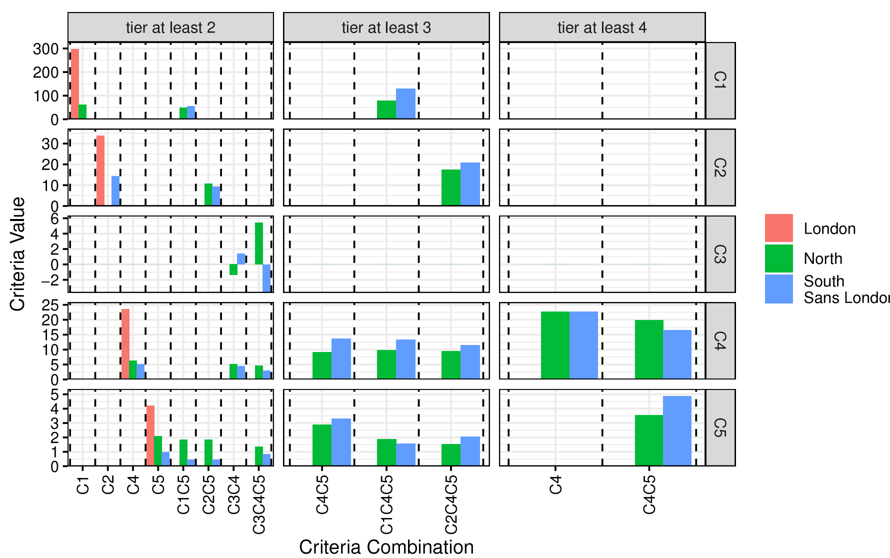

Fig. 8. Upward rules plot.

Table 10
Upwards rules table.

		tier at least 2								tier at least 3			tier at least 4	
		C1	C2	C4	C5	C1C5	C2C5	C3C4	C3C4C5	C4C5	C1C4C5	C2C4C5	C4	C4C5
C1	**London**	297.14				-					-			
	North	62.04				49					80.35			
	South	-				55.43					130.39			
C2	**London**		33.71				-					-		
	North		-				10.71					17.44		
	South		14.19				9.18					20.8		
C3	**London**							-	-					
	North							-1.37	5.44					
	South							1.42	-3.68					
C4	**London**			23.5				-	-	-	-	-	-	-
	North			6.36				5.14	4.64	9.09	9.84	9.43	22.67	19.81
	South			5.12				4.5	2.86	13.64	13.25	11.55	22.63	16.49
C5	**London**				4.22	-	-		-	-	-	-		-
	North				2.09	1.86	1.86		1.35	2.91	1.87	1.54		3.55
	South				1	0.46	0.46		0.83	3.3	1.56	2.07		4.88

twice the threshold for the South. In this case, there was no rule extracted for the North.

For the fourth criterion of C4 (Positivity rate), DRSA has extracted rules for all three of the geographical regions, facilitating a comparison between all three. Here again, the threshold for London (23.5) is many folds higher than others, whilst there is a negligible difference between the North (6.36) and the South (5.12).

Observation 7. *The thresholds for upward rules give evidence of relaxed rules for London for C4 as London was in Tier-2 when others were in Tier-3*

Table 11
Downward rules table.

		tier at most 3													
		C1	**C4**	**C1C2**	**C1C4**	**C2C4**	**C3C4**	**C4C5**	**C1C4C5**	**C2C3C4**	**C2C4C5**	**C3C4C5**	**C1C3C4C5**	**C2C3C4C5**	**C1C2C3C4C5**
C1	**London**	-		-	-				-				-		-
	North	19.5		21.14	22.96				31.55				56.93		33.51
	South	21.09		31.24	42.24				38.29				37.71		42.6
C2	**London**			-		-				-	-			-	-
	North			4.29		6.51				8.84	8.41			7.63	6.54
	South			3.91		5.18				4.86	6.23			7.08	6.52
C3	**London**						-			-		-	-	-	-
	North						-0.08			5.19		6.58	12.64	11.44	11.17
	South						10.44			4.42		8.82	8.26	8.61	6.78
C4	**London**		-		-	-	-	-	-	-	-	-	-	-	-
	North		6		8.14	7.78	8.36	6.75	8.24	8.2	7.19	9.13	7.94	8.81	8.43
	South		5.32		5.57	7.28	5.55	6.22	6.18	7.31	6.65	5.72	7.09	6.76	6.95
C5	**London**							-	-		-	-	-	-	-
	North							1.3	1.77		1.67	1.04	1.48	1.81	2.28
	South							1	1.27		1.27	1.6	1.15	1.28	1.63

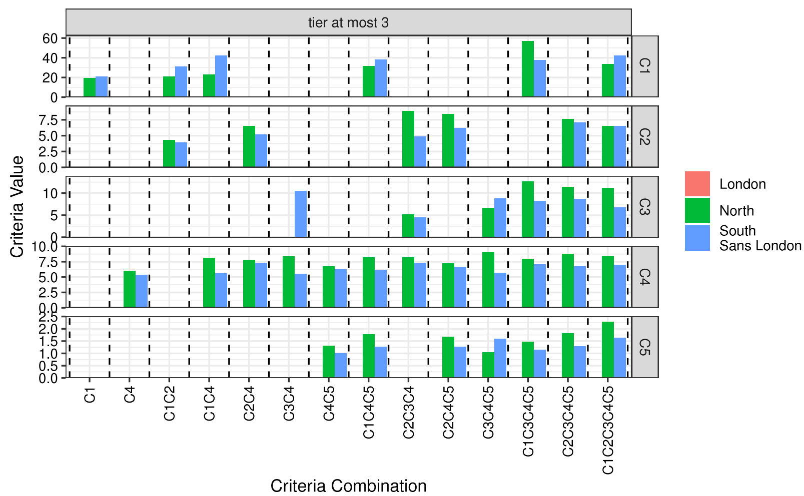

Fig. 9. Downward rules plot.

Considering the fifth criterion (Pressure on NHS), again the threshold for London remains very high (4.22), followed by the North (2.09) and then the South (1.00). To summarise, the extracted rules suggest that the thresholds for London are considerably higher than other parts of England.

Observation 8. *C5 in London remained significantly higher than other regions.*

Focusing on the consequent of tier at least 4, the rules with a single criterion in its antecedent were found only for C4 (see the second last column of Table 10). Here, the rules for the North and the South are showing almost identical threshold values (i.e. 22.67 and 22.63).

So far, we have discussed the rules involving only a single criterion in its antecedent. There also exist other rules involving multiple criteria in their antecedents. For example, the consequent of tier at least 3 shows that there are rules extracted with the criteria C4 and C5 together. The threshold values for these two criteria are 9.09 & 2.91 for the North, and 13.64 & 3.3 for the South, respectively. Please note that as these thresholds are paired together, the values cannot be compared individually. In this case, we can see that the pair of values for the North are both considerably smaller than the pair of values extracted for the South.

The downward rules are summarised in Table 11, and are visualised in Fig. 9 for inspection. Looking at the first row showing thresholds for C1, there are five instances where South sans London has higher threshold values than the North, while there is one instance where North has higher threshold value. On the contrary, looking at the second row, showing thresholds for C2, the North has higher threshold values in all cases. There are no rules extracted for London, however there are several rules extracted for North and South sans London, where, generally there is little to separate the threshold values for these two geographic regions.

Overall, given that the two geographical areas of the North and the South sans London are quite similar, and that London has more relaxed threshold values, it can be argued that although there is a North–South divide in terms of tiers allocation, it is one that is driven by London and not the other regions of the South. This results in London appearing to be given preferential treatment, with respect to the set of Criteria under consideration at least.

With London representing such a large percentage of the economic activity in England, this suggests that there may be additional implicit input into tier allocation reflecting economic considerations as additional criteria,

Observation 9. *There is a north south divide, in terms of tiers allocation, which is driven mostly by London.*

5.3. Discussions

In this section, we have shown how our proposed approach can be applied to conduct fairness analysis on the UK government's tiered restrictions system. The results have identified differences in the rules extracted from the North and the South of England, driven in large part by the region of London. Such inconsistencies suggest unfairness treatment based on geographical location. Our approach looks to extract a separate set of rules from each data slice, where the resulting rules are extracted via (trained via) the observations within the data slice. In a machine learning context, one may argue that it may result in over-fitting, this might be due to the presence of noise, the limited size of training set, and/or model complexity [58]. Rule-set induction methods, like other machine learning models, can be susceptible to over-fitting, especially when the induced rules become too specific [59]. The impact of over fitting is that a model may subsequently fail to fit additional data or predict future observations reliably. Within our approach, the focus in on inducing rules from specific slices of data, so that the underlining patterns are unearthed. Ergo, the focus of our approach is on fairness analysis, not on prediction, we utilise historical data to identify specific patterns and relationships. Therefore, the generalisability of a model is less relevant in conducting investigative analysis i.e. it may be better to have a "reluctance to simplify" in order to identify and address the causes of unfairness [60]. However, we could explore curtailing over-fitting within the approach, if required. For the Covid-19 case study, strategies such as data-expansion [58] would not be viable due to the finite nature of the available data observations (and the set of attributes for this specific problem). However, this might be viable for the application of the approach within other domains. In this context, regularisation may help avoid creating overly complex models, which may over-fit the training data. For example, one way of doing this is through reducing the number of features used to build a model. Within a rule inducing method, having a smaller set of attributes may result in less complex rules, in terms of the number of attributes in the antecedent. Our proposed approach has the ability to incorporate bireduct analysis (see Section 4.5) which facilitates potential regularisation, in terms of reducing the number of attributes to induce rules from. In this way, less complex, and thus less specific to the dataset, rules may be induced. Subsequently, regarding over-fitting, it is worth pondering the trade-off between strict descriptive models that are prone to over-fitting against overly general models that provide little useful information for investigative purposes [59].

From an application perspective, although we illustrated the approach within the COVID-19 pandemic, the approach could provide valuable future insights through both post hoc analysis and through its application to future pandemic data. The post hoc analysis may provide policy makers greater perspicacity and knowledge to bring to bear at the start of a future pandemic. The application to future pandemic data can facilitate greater transparency after decision announcements, for example, by showing that there are no inconsistencies within the decision outcomes. Such issues are pertinent to the government's post hoc analysis of its COVID-19 pandemic handling, as can be seen within the "Transparency and accountability during COVID-19 report" [61]. Here, the committee's conclusions included that communication regarding local lockdown were not transparent enough, leading to confusion and mistrust, and that data underpinning decisions to put some areas under greater restrictions were not clear enough.

For the COVID-19 application of the approach, a number of attributes were utilised to make tiered restrictions of movement decisions. The attribute of geographical location is an inappropriate attribute within the tiered system. Through slicing the data with respect to this attribute, we were able to identify inconsistencies in the assignment of tiered restrictions. The approach could be utilised to explore fairness analysis within other domains with an underlying dataset containing a number of attributes being utilised to make decision outcomes, along with additional attribute(s) which, for the decisions being made, should not be utilised to impact decision outcomes. Such data sets could be sliced with respect to such an additional attribute that should not be impacting decisions, for example, attributes such as age within recruitment decisions, or gender within exam grade decisions. Then fairness analysis can be explored through comparing rules extracted from the separate slices.

6. Conclusion

In this paper, we have proposed a novel approach to investigate fairness in data-driven decisions with the help of DRSA, and explored its application to analyse inconsistencies in allocating tiered restrictions of movement during COVID-19. Our approach is able to identify patterns in the data, pertaining to the UK government criteria set for their tiered allocations, which were translated into "if-then" type rules. We demonstrate a novel way of investigating fairness by extracting separate sets of rules for separate segments, and then comparing these rule-sets to investigate fairness discrepancies in these segments. We found differences in the rules extracted from the North and the South of England, driven in large part by the region of London.

We found that there is a high level of correlation between C1 and C2, from which one could question how useful was it to have both in deciding the tiered restrictions. Intuitively, it makes sense to consider separate criteria to capture age disparities, given the consensus that age has a direct impact upon the severity of illness. However, our analysis suggests that the use of C2 as a separate criterion was statistically redundant. Our analysis also suggests that C3 was also not very useful, in a sense that it had almost no explanatory power in allocating tiers. Moreover, there is no information made public regarding the relative importance assigned to each of these criteria. As the problem involved multiple criteria, the focus of released information remained on public health which ignored the use of other important information related to, for example, economy, society and technology. Essentially, the allocation of tiers is a problem involving conflicting objectives where minimising the risk conflicts with the maximising of economic prosperity. Therefore, this is an area of research involving trade-off analysis and other multi-criteria decision making techniques.

We argue that, despite being lauded as transparent, the proposed systematic approach was still not transparent enough. The UK government did collect the data on the set of criteria to assign tiers, but the use of this (and any other implicate) information was obfuscated. One of the implications of a perceived lack of transparency of such a system, combined with an apparent disparity in treatment for different segments, is that it could lead to a breakdown of trust. This in turn could lead to non-conformity to the rules, thus defeating the purpose of the whole tiered-allocation system.

This study can be further researched in a number of ways both theoretically and in terms of its applications. From a theoretical perspective, as discussed earlier, topics of over-fitting and regularisation can be further investigated. From a practical perspective, future work can also explore fairness analysis within other domains for which there exists attribute(s) which, with respect to the decisions being made, should not be utilised to affect decision outcomes (such as age or gender within recruitment or exam grade decisions). Within the COVID-19 restrictions domain, further work can delve into the inconsistencies the approach has revealed, in order to explore hypothesises regarding the underlying causes. For example, hypothesising about potential additional criteria which are obfuscated within the system, and experimenting with how including such criteria might impact inconsistencies, could provide suggestions regarding obfuscated criteria at play within the decision making. Moreover, the creation of an interactive software tool is also an area of further work that can allow users to utilise the proposed approach on datasets in their application domains. Such a tool could enable users to load their raw data set, define the method for segmenting data into separate slices, and then perform fairness analysis with respect to these segments.

CRediT authorship contribution statement

Edward Abel: Conceptualization, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing. **Sajid Siraj:** Conceptualization, Investigation, Methodology, Resources, Software, Validation, Visualization, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

[1] D. Yu, Z. Xu, W. Pedrycz, Bibliometric analysis of rough sets research, Appl. Soft Comput. 94 (2020) 106467.

[2] K.-Y. Shen, G.-H. Tzeng, Combined soft computing model for value stock selection based on fundamental analysis, Appl. Soft Comput. 37 (2015) 142–155, http://dx.doi.org/10.1016/j.asoc.2015.07.030, URL https://www.sciencedirect.com/science/article/pii/S1568494615004706.

[3] J.J.H. Liou, M.-H. Chang, H.-W. Lo, M.-H. Hsu, Application of an MCDM model with data mining techniques for green supplier evaluation and selection, Appl. Soft Comput. 109 (2021) 107534, http://dx.doi.org/10.1016/j.asoc.2021.107534, URL https://www.sciencedirect.com/science/article/pii/S1568494621004579.

[4] S. Barocas, M. Hardt, A. Narayanan, Fairness and Machine Learning: Limitations and Opportunities, 2019, fairmlbook.org. URL https://fairmlbook.org/.

[5] A. Narayanan, 21 Fairness Definitions and Their Politics, Tech. rep., 2018, URL https://fairmlbook.org/tutorial2.html.

[6] R. Berk, H. Heidari, S. Jabbari, M. Kearns, A. Roth, Fairness in criminal justice risk assessments: The state of the art, Sociol. Methods Res. 50 (1) (2018) 3–44, http://dx.doi.org/10.1177/0049124118782533, URL https://journals.sagepub.com/doi/full/10.1177/0049124118782533.

[7] A.N. Carey, X. Wu, The statistical fairness field guide: perspectives from social and formal sciences, AI Ethics 3 (1) (2022) 1–23, http://dx.doi.org/10.1007/S43681-022-00183-3.

[8] Department of Health and Social Care, GOV.UK, Tiering decisions to be made through national command structure, in: Department of Health and Social Care, 2020, URL https://www.gov.uk/government/news/tiering-decisions-to-be-made-through-national-command-structure.

[9] S. Greco, B. Matarazzo, R. Slowinski, Rough sets theory for multicriteria decision analysis, European J. Oper. Res. 129 (1) (2001) 1–47.

[10] J.J. Liou, C.H. Tang, W.C. Yeh, C.Y. Tsai, A decision rules approach for improvement of airport service quality, Expert Syst. Appl. 38 (11) (2011) 13723–13730, http://dx.doi.org/10.1016/J.ESWA.2011.04.168.

[11] N. Rehman, A. Ali, K. Hila, Soft dominance based multigranulation decision theoretic rough sets and their applications in conflict problems, Artif. Intell. Rev. 53 (8) (2020) 6079–6110, http://dx.doi.org/10.1007/S10462-020-09843-4/METRICS, URL https://link.springer.com/article/10.1007/s10462-020-09843-4.

[12] M.G. Augeri, S. Greco, V. Nicolosi, Planning urban pavement maintenance by a new interactive multiobjective optimization approach, Eur. Transp. Res. Rev. 11 (1) (2019) 1–14, http://dx.doi.org/10.1186/S12544-019-0353-9/FIGURES/3, https://link.springer.com/article/10.1186/s12544-019-0353-9.

[13] C.J.J.d. Figueiredo, C.M.d.M. Mota, A.G.F. Rosa, A.P.G.d. Souza, S.M.d.S. Lima, Vulnerability to COVID-19 in Pernambuco, Brazil: A geospatial evaluation supported by multiple-criteria decision aid methodology, Geospatial Health 17 (s1) (2022) http://dx.doi.org/10.4081/gh.2022.1000, URL https://www.geospatialhealth.net/index.php/gh/article/view/1000.

[14] P. Ferretti, M.B. Zolin, G. Ferraro, Relationships among sustainability dimensions: evidence from an Alpine area case study using Dominance-based Rough Set Approach, Land Use Policy 92 (2020) 104457, http://dx.doi.org/10.1016/J.LANDUSEPOL.2019.104457.

[15] M. Karolak, Å. Pałkowski, B. Kubiak, J. Błaszczyński, R. Łunio, W. Sawicki, R. Słowiński, J. Krysiński, Application of dominance-based rough set approach for optimization of pellets tableting process, Pharmaceutics 12 (11) (2020) 1–14, http://dx.doi.org/10.3390/PHARMACEUTICS12111024, /pmc/articles/PMC7692369/ /pmc/articles/PMC7692369/?report=abstract https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7692369/.

[16] J. Błaszczyński, A.T. de Almeida Filho, A. Matuszyk, M. Szeląg, R. Słowiński, Auto loan fraud detection using dominance-based rough set approach versus machine learning methods, Expert Syst. Appl. 163 (2021) 113740, http://dx.doi.org/10.1016/J.ESWA.2020.113740.

[17] N. Dehouche, Predicting seminal quality with the dominance-based rough sets approach, ICIC Express Lett. 14 (7) (2020) 653–659, http://dx.doi.org/10.48550/arxiv.2012.13204, URL https://arxiv.org/abs/2012.13204v1.

[18] H.R. Bhapkar, P.N. Mahalle, G.R. Shinde, M. Mahmud, Rough sets in COVID-19 to predict symptomatic cases, in: Lecture Notes on Data Engineering and Communications Technologies, Vol. 60, Springer, Singapore, 2021, pp. 57–68, http://dx.doi.org/10.1007/978-981-15-9682-7{_}7, URL https://link.springer.com/chapter/10.1007/978-981-15-9682-7.

[19] Q. Hu, S. Chakhar, S. Siraj, A. Labib, Spare parts classification in industrial manufacturing using the dominance-based rough set approach, European J. Oper. Res. 262 (3) (2017) 1136–1163, http://dx.doi.org/10.1016/J.EJOR.2017.04.040.

[20] O.A. Egaji, S. Chakhar, D. Brown, An innovative decision rule approach to tyre pressure monitoring, Expert Syst. Appl. 124 (2019) 252–270, http://dx.doi.org/10.1016/J.ESWA.2019.01.051.

[21] I. Brigui, S. Chakhar, Dominance relation based ranking procedure for automated reverse auctions, in: Lecture Notes in Networks and Systems, vol. 441 LNNS, Springer, 2022, pp. 137–145, http://dx.doi.org/10.1007/978-3-030-98531-8{_}14.

[22] M. Palangetić, C. Cornelis, S. Greco, R. Słowiński, Fuzzy extensions of the dominance-based rough set approach, Internat. J. Approx. Reason. 129 (2021) 1–19, http://dx.doi.org/10.1016/J.IJAR.2020.10.004.

[23] K.Y. Shen, G.H. Tzeng, Fuzzy inference-enhanced VC-DRSA model for technical analysis: Investment decision aid, Int. J. Fuzzy Syst. 17 (3) (2015) 375–389, http://dx.doi.org/10.1007/S40815-015-0058-8/TABLES/10, URL https://link.springer.com/article/10.1007/s40815-015-0058-8.

[24] Y. Kusunoki, M. Inuiguchi, A unified approach to reducts in dominance-based rough set approach, Soft Comput. 14 (5) (2010) 507–515, http://dx.doi.org/10.1007/S00500-009-0450-0/TABLES/3, URL https://link.springer.com/article/10.1007/s00500-009-0450-0.

[25] M.J. Benítez-Caballero, J. Medina, E. Ramírez-Poussa, D. Ślézak, Bireducts with tolerance relations, Inform. Sci. 435 (2018) 26–39, http://dx.doi.org/10.1016/J.INS.2017.12.037.

[26] D. Ślęzak, S. Stawicki, The problem of finding the simplest classifier ensemble is NP-hard – a rough-set-inspired formulation based on decision bireducts, in: Rough Sets, Vol. 12179, Nature Publishing Group, 2020, p. 204, http://dx.doi.org/10.1007/978-3-030-52705-1{_}15, URL https://www.ncbi.nlm.nih.gov/pmc/articles/PMC7338167/.

[27] S. Stawicki, D. Ślęzak, A. Janusz, S. Widz, Decision bireducts and decision reducts – a comparison, Internat. J. Approx. Reason. 84 (2017) 75–109, http://dx.doi.org/10.1016/J.IJAR.2017.02.007.

[28] A. Janusz, D. Ślęzak, S. Stawicki, K. Stencel, A Practical Study of Methods for Deriving Insightful Attribute Importance Rankings Using Decision Bireducts, Tech. rep., 2023, http://dx.doi.org/10.2139/SSRN.4418685, URL https://papers.ssrn.com/abstract=4418685.

[29] F. Cabitza, A. Campagner, V. Basile, Toward a perspectivist turn in ground truthing for predictive computing, in: Proceedings of the AAAI Conference on Artificial Intelligence, Vol. 37, No. 6, Association for the Advancement of Artificial Intelligence (AAAI), 2023, pp. 6860–6868, http://dx.doi.org/10.1609/AAAI.V37I6.25840, URL https://ojs.aaai.org/index.php/AAAI/article/view/25840.

[30] U. Hébert-Johnson, M.P. Kim, O. Reingold, G.N. Rothblum, Multicalibration: Calibration for the (computationally-identifiable) masses, in: Conference on Machine Learning, PMLR, 2018, pp. 1939–1948, URL https://proceedings.mlr.press/v80/hebert-johnson18a.html.

[31] M.M. Bassiouni, R.K. Chakrabortty, O.K. Hussain, H.F. Rahman, Advanced deep learning approaches to predict supply chain risks under COVID-19 restrictions, Expert Syst. Appl. 211 (2023) 118604, http://dx.doi.org/10.1016/J.ESWA.2022.118604.

[32] GOV.UK, Analysis of the Health, Economic and Social Effects of COVID-19 and the Approach to Tiering Executive Summary, Tech. rep., HM Government, UK, 2020, pp. 1–48.

[33] DHSC, M. Hancock, Written statement to Parliament - Returning to a regional tiered approach, 2020, URL https://www.gov.uk/government/speeches/returning-to-a-regional-tiered-approach.

[34] S. Greco, B. Matarazzo, R. Slowinski, J. Stefanowski, An algorithm for induction of decision rules consistent with the dominance principle, in: Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics). 2005, Springer, Berlin, Heidelberg, 2000, pp. 304–313.

[35] J. Błaszczyński, R. Słowiński, M. Szelg, Sequential covering rule induction algorithm for variable consistency rough set approaches, Inf. Sci. Inf. Comput. Sci. Intell. Syst. Appl. An Int. J. 181 (5) (2011) 987–1002, http://dx.doi.org/10.1016/J.INS.2010.10.030, URL https://dl.acm.org/doi/abs/10.1016/j.ins.2010.10.030.

[36] J. Scheerhout, Andy Burnham slams ministers for the way they spared London from Tier 3 - Manchester Evening News, 2020, Manchester Evening News. URL https://www.manchestereveningnews.co.uk/news/greater-manchester-news/andy-burnham-slams-ministers-way-19376371.

[37] NewsMediaOrganization, Covid: Yorkshire tier 3 decision 'unfair', councillors say, 2020, BBC News. URL https://www.bbc.com/news/uk-england-55349559.

[38] F. Alkaabneh, A. Diabat, H.O. Gao, A unified framework for efficient, effective, and fair resource allocation by food banks using an Approximate Dynamic Programming approach, Omega 100 (2021) 102300, http://dx.doi.org/10.1016/j.omega.2020.102300.

[39] G. Busch, A. Spiller, Farmer share and fair distribution in food chains from a consumer's perspective, J. Econ. Psychol. 55 (2016) 149–158.

[40] M. Tavana, K. Govindan, A.K. Nasr, M.S. Heidary, H. Mina, A mathematical programming approach for equitable COVID-19 vaccine distribution in developing countries, Ann. Oper. Res. (2021).

[41] O.F. Norheim, Ethical priority setting for universal health coverage: challenges in deciding upon fair distribution of health services, BMC Med. 14 (1) (2016) 75, http://dx.doi.org/10.1186/s12916-016-0624-4, URL http://bmcmedicine.biomedcentral.com/articles/10.1186/s12916-016-0624-4.

[42] G.L. Gaile, Effiquity: A comparison of a measure of efficiency with an entropic measure of the equality of discrete spatial distributions, Econ. Geogr. 53 (3) (1977) 265–282, http://dx.doi.org/10.2307/143192.

[43] R. Shankar, A. Shah, Bridging the economic divide within countries: A scorecard on the performance of regional policies in reducing regional income disparities, World Dev. 31 (8) (2003) 1421–1441, http://dx.doi.org/10.1016/S0305-750X(03)00098-6.
[44] M.D. Alicke, E. Zell, Social comparison and envy, in: R.H. Smith (Ed.), Envy: Theory and Research, Oxford University Press, 2008, pp. 73–93, http://dx.doi.org/10.1093/acprof:oso/9780195327953.003.0005.
[45] S.N. Williams, C.J. Armitage, P. of Health Psychology, K. Dienes, L. in Clinical, H. Psychology, Public Perceptions of Non-Adherence to COVID-19 Measures by Self and Others in the United Kingdom, Cold Spring Harbor Laboratory Press, 2020, http://dx.doi.org/10.1101/2020.11.17.20233486, medRxiv. 2020.11.17.20233486. https://www.medrxiv.org/content/10.1101/2020.11.17.20233486v1.abstract.
[46] J. Carthey, S. Walker, V. Deelchand, C. Vincent, W.H. Griffiths, Breaking the rules: understanding non-compliance with policies and guidelines, BMJ 343 (2011) http://dx.doi.org/10.1136/bmj.d5283, URL https://www.bmj.com/content/343/bmj.d5283.
[47] The Economist, The great divide, 2012, The Economist. URL https://www.economist.com/britain/2012/09/15/the-great-divide.
[48] J. Webb, M. Johns, E. Roscoe, A. Giovannini, A. Qureshi, R. Baldini, State of the North 2021/22 Powering Northern Excellence the Progressive Policy Think Tank, Tech. rep., 2022, URL www.ippr.org/north.
[49] M. Smith, What regions make up the North and South of England? | YouGov, 2018, YouGov.UK. URL https://yougov.co.uk/topics/politics/articles-reports/2018/01/04/what-regions-make-north-and-south-england.
[50] M. Khan, M.T. Mehran, Z.U. Haq, Z. Ullah, S.R. Naqvi, M. Ihsan, H. Abbass, Applications of artificial intelligence in COVID-19 pandemic: A comprehensive review, Expert Syst. Appl. 185 (2021) 115695, http://dx.doi.org/10.1016/J.ESWA.2021.115695.
[51] R. Moss, Why People Aren\'t Going to Follow The Three-Tier Lockdown Rules, Tech. rep., HuffPost UK, 2020, URL https://bit.ly/3qWJ8xr.
[52] ONS, England. Detailed Information on the Administrative Structure Within England, Office for National Statistics, 2016, URL www.ons.gov.uk/methodology/geography/ukgeographies/administrativegeography/england.
[53] C. Katris, A time series-based statistical approach for outbreak spread forecasting: Application of COVID-19 in Greece, Expert Syst. Appl. 166 (2021) 114077, http://dx.doi.org/10.1016/J.ESWA.2020.114077.
[54] S. Meakin, S. Funk, Quantifying the Impact of Hospital Catchment Area Definitions on Hospital Admissions Forecasts: COVID-19 in England, September 2020 - April 2021, Cold Spring Harbor Laboratory Press, 2023, http://dx.doi.org/10.1101/2023.07.12.23292451, medRxiv. 2023.07.12.23292451. https://www.medrxiv.org/content/10.1101/2023.07.12.23292451v1.abstract.
[55] GOV.UK, The Health Protection (Coronavirus, Restrictions) (All Tiers) (England) Regulations 2020, Tech. rep., Queen's Printer of Acts of Parliament, Public Health England, 2020.
[56] M. Flores-Sosa, E. Avilés-Ochoa, J.M. Merigó, J. Kacprzyk, The OWA operator in multiple linear regression, Appl. Soft Comput. 124 (2022) 108985, http://dx.doi.org/10.1016/J.ASOC.2022.108985.
[57] F. Blanco-Mesa, E. León-Castro, J.M. Merigó, A bibliometric analysis of aggregation operators, Appl. Soft Comput. 81 (2019) 105488, http://dx.doi.org/10.1016/J.ASOC.2019.105488.
[58] X. Ying, An overview of overfitting and its solutions, J. Phys. Conf. Ser. 1168 (2) (2019) 022022, http://dx.doi.org/10.1088/1742-6596/1168/2/022022, https://iopscience.iop.org/article/10.1088/1742-6596/1168/2/022022/meta.
[59] E. Gil-Herrera, G. Aden-Buie, A. Yalcin, A. Tsalatsanis, L.E. Barnes, B. Djulbegovic, Rough set theory based prognostic classification models for hospice referral, BMC Med. Inf. Decis. Mak. 15 (1) (2015) 1–18, http://dx.doi.org/10.1186/S12911-015-0216-9/TABLES/7.
[60] K.E. Weick, K.M. Sutcliffe, Managing the Unexpected : Sustained Performance in a Complex World, Jossey-Bass, 2015, p. 224, URL https://www.wiley.com/en-gb/Managing+the+Unexpected%3A+Sustained+Performance+in+a+Complex+World%2C+3rd+Edition-p-9781118862414.
[61] P. Administration, C. Affairs Committee, Government Transparency and Accountability During COVID 19: the Data Underpinning Decisions, Tech. rep., 2021, URL www.parliament.uk/copyright..