Using Model Class Reliance to Measure Group Effects on Non-Adherence to Asthma Medication

Vanja Ljevar*, James Goulding*, Gavin Smith*, Alexa Spence†
*N/LAB, University of Nottingham, UK
{vanja.ljevar1, james.goulding, gavin.smith}@nottingham.ac.uk
†School of Psychology, University of Nottingham, UK
alexa.spence@nottingham.ac.uk

***Abstract*—Asthma affects an estimated 300 million people across the world. Despite being a highly treatable condition using preventative inhalers, mortality rates remain unacceptably high, with lack of adherence to medications cited as a major cause. While various drivers for non-adherence have been considered in isolation, interactions between demographic, behavioural and situational factors have never been modelled in concert mostly due to the limited ability of traditional methods to group such a large variety of features. This was addressed in this paper through a non-linear modelling approach, leveraging a novel dataset obtained via online surveying of asthma patients. Application of traditional variable importance methods to examine explanatory factors, however, is not possible. This is due to the presence of high multicollinearity in the data, a highly common occurrence in big data, or any datasets which include a large number of input features. This results in insights being obfuscated by extensive shared information and non-linear interactions occurring across variables. To mitigate this, we introduce the first Grouped Feature approach to Model Class Reliance (Group-MCR), that is able to quantify the importance of specific variable sets in underpinning explanations. Cross-validated models achieve 71% accuracy, with Group-MCR revealing the importance of perceptual factors. Out of all the perceptual factors *denial* proves to be most predictive of non-adherence to asthma medication, indicating that public health interventions should not only target the physical aspects of asthma, but additionally focus on patients' beliefs and perceptions as valuable parts of their treatment.**

I. Introduction

Despite being a highly treatable condition, asthma mortality rates remain unacceptably high. In the UK alone an average of 3 people die from an asthma attack each day (Asthma UK). Yet while symptoms are rapidly alleviated by preventative inhalers, and regular use of asthma medication is essential in preventing negative asthma outcomes [1], only 38% to 50% of patients actually adhere to their prescribed treatment regimes [2]. This is despite the use of inhalers decreasing hospital admissions by 30%, making patient cooperation in using medication prescribed key to successfully managing the condition. Many drivers for this non-adherence have been proposed, and there is little doubt that patients' use of medication is impacted by demographics and lifestyle choices, as well as contextual and psychological factors. Yet, there exists little consensus as to which of these groups of factors are most central to non-adherence behaviours [3]. Cross-sectional data is lacking, and while many covariates have been examined in an attempt to explain patients' reticence to medicate, studies have generally occurred in isolation, considering only one or few variables at a time.

Identifying the most relevant group of factors is valuable not only to practitioners, but to policy makers and those implementing health interventions. Better identification of the categories of features most linked to non-adherence promises to provide not only a better understanding of the condition itself, but holds potential to generate actionable early warning signals, identifying those who may be at risk. It is possible to identify individual feature importances and act on them. However, acting on specific *groups of factors* related to non-adherence, in concert, is highly desirable given that such groups often require different channels of communications, different types of interventions and share similar behavioural underpinnings [3], [4]. Therefore, knowing *which group*, rather than which individual feature, is of greatest importance can lead to different interventions structures and informed planning. For example, interventions may be most effective by reaching out to patients in rural areas, by increasing education in poorer communities, or by focusing more on simply creating behaviourally focused therapies [4]. Understanding which groups of factors are most predictive of non-adherence seeks to lead us to better patient outcomes and the lowering of both personal and stakeholders' costs [3].

This work therefore examines two research questions via an inductive data modelling approach: • *What are the main groups of risk factors that predict non-adherence to asthma medication?* and • *Which individual features within these groupings are particularly representative of non-adherence?*. Measuring the performance of predictive factors is a multi-dimensional problem. With indicators of non-adherence not necessarily being mutually exclusive [3] and, in reality highly likely to interact with one another, a distinct methodological challenge is presented.

A wide range of approaches have been applied in an attempt to better understand the risk factors behind negative health outcomes (see Section III). However, a key challenge, and a common characteristic of big datasets in general, is the large range of input features faced - an issue widely

Vanja Ljevar is supported by the Horizon Centre for Doctoral Training at the University of Nottingham (UKRI Grant No. EP/L015463/1).

Pre-print. 978-1-6654-3902-2/21/$31.00 ©2021 IEEE

recognised as restricting traditional regression methods. Additionally, traditional analytical approaches have struggled to cope with the non-linearity and high information sharing occurring between input features (the most simple version of which is multicollinearity). Methods in social sciences and pharmacology, in particular, have had problems unpicking group effects, with remedial technical solutions being severely prone to information loss (due to elimination of features that do not contribute to a 'compound scale', for example [5]). Yet there also remain deficiencies in methods derived from the Computer Sciences, with explanatory approaches such as *grouped permutation importance* providing insights only about the mechanism of a single model, rather than the generative processes of the underlying phenomena - again, multi-collinearity and interaction effects can often result in misleading conclusions [6], [7].

To meet these challenges, in this paper we not only take a non-linear modelling approach, leveraging a novel dataset drawn from surveying asthma patients, but additionally introduce a concept of *Grouped Model Class Reliance* (Group-MCR). In contrast to group permutation importance, this extension is able to take into account multiple models to uncover underlying relationships between groups of factors and indicate each group's utility in predicting adherence. The rest of the paper is structured as follows: Section II details related work with regard to risk factors, Section III highlights the challenges that exist with regard to applying recent/traditional approaches to achieve the aims of this work and Section IV details the dataset and the proposed methodology/experimental design, including the use of Group-MCR. Results are then provided in Section V, followed by a discussion and conclusion in Sections VI and VII respectively.

II. Related Work

Adherence is defined as the extent to which a patient's behaviour matches the recommendations from a medical professional in relation to their treatment [8]. While never previously assessed in concert, various potential drivers of non-adherence to asthma medication have been considered in isolation. A wide range of predictive factors include: the complexity of therapy; patients' understanding of the illness and its complications; social support of patients; drug availability; disappointment due to lack of improvement; and level of patient's education [9], [10]. Many such factors share a common thread, all speaking to patients' ability and motivation to obtain effective support or acquire knowledge [11], [12].

An extensive stream of work, specifically considering asthma non-adherence, has focused on socio-demographic factors. Previous research, however, has not yet identified clear and consistent relationships between medication usage and demographic variables, such as gender and age [3], [8]. There is some evidence indicating that patients with better adherence behaviour tend to be female [10], from older demographics [4] and are most frequently married [8]. Despite these links being made, studies show that socio-demographics and clinical factors can only account for a small amount of variance in adherence [8] - and that additional other factors are relevant. In particular, lifestyle factors, such as smoking and lack of exercise are seen as negatively associated to adherence to prescribed medication [13].

The fact that socio-demographics and lifestyle features have only managed to yield partial explanations for non-adherence has increasingly led to examination of the impact of factors such as personality traits. Individual differences such as personality traits are the focus of several studies on adherence, yet consensus is still to be achieved on what characteristics may relate to non-adherence. For example, *neuroticism* (a personality trait from the well-established five-factor psychological model [14]–[16]) was found to have a negative impact on adherence, while *conscientiousness* and *agreeableness* had positive effects [17], [18]; yet, contradictory results have also been published indicating that personality does not have significant impact after all [8], [19]. Other psychological studies have considered the impact of the emotional states of patients, or affectivity. Negative affect, a common symptom of depression, has been correlated with low adherence to asthma mediation [2], a result supported by recent work evidencing that emotions are also predictive of intentions to adhere [20]. Other studies have focused on specific emotional states. Evidence of fears of asthma consequences, for example, have been strongly associated with reduced willingness to use medication [21] - whereas 'hedonic capacity' of patients, has been positively correlated with effective asthma control [2].

Less well examined, however, are the relationships between *perceptual* factors and adherence to asthma medication. This is perhaps surprising for asthma, which has been associated with issues of both stigma and avoidance [22]. Perceptual influences are rooted in the two steps patients must take before they adhere to a medical regimen: 1. deciding whether they accept they have a condition; and 2. whether they consider the prescribed medication is suitable and worth any perceived risks [23], [24]. It has been reported that such decisions relate to patients' notion of their own personal identity, and that maintaining an existing quality of life dominates many individuals' experiences of asthma [25]. Perceptual factors relating to asthma, and the perceived 'weakness' some attach to the condition, have related effects on coping strategies [8]. Coping efforts vary widely across individuals - in some responses to diagnosis, patients can be 'active, problem-focused and adaptive'; while others are 'passive, emotion-focused and focus on condition negation' [9]. Where an individual falls across this spectrum is highly relevant to long-term behavioural patterns related to appropriate asthma management [26].

As described above, numerous potential factors have been found to impact on non-adherence - yet the number of available candidate variables emphasizes the key challenge faced - it is extremely difficult to isolate the impact of one factor individually. Underlying drivers of non-adherence tend to be grouped, deeply embedded in people's daily lives, ranging from behavioural to situational circumstances, psychological to demographic. This situation is exacerbated by the number of (potentially non-linear) interactions that exist across indicators.

An example is age and stigma: as individuals age, their risk of hospitalization due to asthma increases; in concert the stigma patients anticipate when carrying a preventative inhaler in public decreases, reducing risk; examining age and perceived stigma separately may inflate relationships observed. Such interactions, and the range of mediating and moderating factors at play across patients' lives, are the motivation behind this study, which takes a non-linear and inductive approach to the problem domain. Leveraging machine learning techniques, we model a new dataset of 80 factors relevant to asthma adherence, drawn from the empirical and theoretical literature, to interrogate explanations of non-adherence.

III. Identifying actionable feature groups

Due to the complexity of the non-parametric approaches common in machine learning, derivation of stable model explanations is often non-trivial. A goal of this work, therefore, is not only to construct an accurate predictive model of adherence to asthma medication, but to use outputs to elicit robust insights into specific *groups of risk factors* that might inform future medical interventions [27]. To achieve actionable variable groups of this nature requires feature importance techniques. Importantly, research into model interpretability can also be dichotomised into two approaches in assessing the predictive performance of feature groups:

- *a priori combination:* features are combined prior to modelling. This corresponds to the tradition of feature engineering in machine learning and includes, but is not limited to, compression techniques such as PCA; Multidimensional Scaling; Kernel PCA; Maximum Variance Unfolding; and Partial Least Squares. All such techniques identify components (or latent variables) that reduce dimensionality - yet which are notoriously hard to interpret, restricting explanatory insights about an application domain [28].
- *posteriori combination:* feature combinations are assessed after modelling has occurred. Such approaches are particularly appropriate when input features sets are associated with specific groupings from the outset - either based on the theoretical or domain knowledge, or due to processes of data collection. Increasingly, improvements in algorithmic efficiency have reduced the need for exhaustive feature compression, allowing group importance to be assessed after modelling in linear and non-linear models (e.g. via grouped permutation importance) [6].

The main focus of this research is to provide explanatory analysis of non-adherence to asthma medication in order to inform medical interventions to promote adherence. Use of *a priori* combinations are therefore less desirable, due to complications in interpretation. Moreover, the dependent variables used in health research are often categorized due to their conceptual nature (e.g. therapy-, condition-, patient-related), making a *posteriori* approach to group analyses more attractive.

A priori reduction of features into actionable sets is, actually, a common task in 'survey-based' research. Here, individual items are often combined into scales (constructed from variables that were designed to measure the same theorised underlying concept) with reliability assessed using measures such as Cronbach's alpha [29]. A high Cronbach's alpha supports the conjecture that features are interrelated with one another, and can be combined into a compound variable by taking a mean value, or sum [30]. While extremely common in social sciences research, this process has numerous drawbacks. First, features that are not linearly correlated cannot be allocated to the same compound 'group', even if they link thematically. Second, items of varying types, such as categorical and continuous features, cannot be incorporated. Finally, Cronbach's alpha is only valid for data that is linear and normally distributed [29] which means valid variable combinations can be overlooked. It is also common for variables that cannot be grouped to be eliminated from modelling altogether, a situation that preserves the integrity of analysis, but at the cost of information loss and predictive accuracy - and consequently weakening the case for the validity of explanations produced [5].

Increasingly, machine learning approaches are eschewing *a priori* feature construction of this form, using approaches such as grouped permutation importances, performed on models trained on all original features [6]. Groups of features are simultaneously permuted to measure joint effects as their information is disrupted. In a similar fashion, feature exclusion and retraining methods[1] [31] include or exclude groups of features and assess the significance of their inclusion in the model (although the use of methods to then determine such significance such as split counts and SHAP values [32] remains non-trivial with feature sets).

A major, and sometimes overlooked, weakness of these approaches is that they only tell us about the mechanism of a single model - and not the underlying mechanism of the phenomenon itself. It has, therefore, been noted that the explanations they output can be misleading regarding an underlying generative process [7], [31]. This is due to the fact that input features typically share predictive information (i.e. multicollinearity). Focusing on learning a single functional relationship, despite many possible and equally predictive relationships existing within the data set, can provide deceptive results [31]. To address these issues, Model Class Reliance (MCR) has recently been proposed [7], [31]. This method implicitly considers the full set of models with equal predictive performance to the learnt 'best' model, known as its *Rashomon Set*, reporting variable importance bounds for each variable. Due to the goal of finding actionable explanatory insights for non-adherence, we use MCR within this study, extending its capabilities to understand grouped features. The goal of this work is to propose a way to solve these methodological challenges in order to observe several groups of well-known non-adherence drivers and discover the most powerful one.

[1]This includes step-wise approaches more often used for variable selection

IV. METHOD

A. Participants

Data was collected via an online survey, distributed via Google Forms and advertised through Prolific. The survey was reviewed and approved by an independent ethics committee at the University of Nottingham. Each participant was rewarded with points by Prolific, equivalent to £2. No identifying personal information about participants was collected. Only people who live in the UK, were older than 18, and possessed a current asthma diagnosis were considered. Participants were also required to complete a consent form in order to participate. Survey responses were collected during the period 1st to 5th May in 2020. Responses obtained in less than 4 minutes were excluded as this was considered too short to have read the questions posed. A total of 511 valid individual responses were collected from the survey. 344 (67.3%) participants were female and 165 (32.3%) were male. 214 participants (41.87%) had high adherence, with 297 participants (58.12%) having low adherence.

B. Materials

The survey lasted approximately 20 minutes and had 80 individual questions, designed to investigate the relationship between asthma medication regime adherence and a range of dependent feature groups. Adherence was measured via the Morisky-8 item questionnaire [33], composed of eight questions with 'yes' or 'no' response options, but producing a compound adherence score ranging between [0,7] (with higher scores implying higher adherence). In accordance with previous literature, scores can be further categorized into a binary adherence score using a cut-off point of six, with a score <6 corresponding to low adherence and a score ≥ 6 to high adherence [33]. The survey's other questions can be categorized into seven groups of factors, all typically measured using (1-7) Likert Scales (Table I details exemplar items from each group and includes some key research references. The full question list is available[2]).

- *Asthma traits:* characteristics of asthma and its symptoms were measured via: condition type (mild/severe); duration; inhaler type; and severity of symptoms [34].
- *Demographics:* Socio-demographics questions: age; gender; education; employment; and income [10], [26]).
- *Perceptual Factors:* based upon qualitative examinations of adherence to asthma medication from a variety of studies, perceptions measured included perceived stigma of having asthma [35]; perceived sense of support [36]; negative impacts of media, and disparaging humour [37].
- *Coping mechanisms:* Coping mechanisms were examined using the adjusted Asthma Specific Coping Scale, which is composed of questions related to: information seeking; hiding; ignoring; worrying about asthma; restrictive lifestyles; and positive reappraisal [38].

[2]http://www.nlab.org.uk/wp-content/uploads/Asthma_Survey_Question_List.pdf

TABLE I
LIST OF FEATURES GROUPS, DESCRIPTIONS AND EXEMPLAR REFERENCES

Variable Group	Description	References
Adherence Score	Morisky-8 question adherence scale (0-7)	[33], [40]
Asthma Traits	Asthma impact (level of the individual): • asthma severity (mild/severe) • duration of condition (years) • inhaler type (categorical)	[34]
Demographics	Age, gender, education, employment and income	[10], [26]
Perceptual Factors	Asthma perceptions influenced by exogenous factors • adjusted Stigma Scale • perceived sense of support • perceived effects of media and disparaging humour	[35]–[37]
Coping mechanisms	Asthma Specific Coping Scale	[38]
Emotional Affect	Range of features based upon PANAS scoring system	[39]
Lifestyle Habits	Identification of Patient controllable behaviours: • smoking • exercise	[13], [25]
Personality traits	Big-5 trait measurement scale	[14], [16]

- *Emotional Affect:* Propensity for positive and negative emotions measured in the survey was based upon the PANAS scoring system [39]. Emotions in relation to asthma that were measured were: feeling sad, strong, scared, hostile, ashamed, nervous, determined, guilt.
- *Lifestyle Factors:* Lifestyle behaviour items were collected, relating to factors within patients' realm of control, identifying habits such as smoking and exercise.
- *Personality traits:* Personality traits were collected using the Ten Item Personality Inventory (TIPI) version of the well-established Big Five: neuroticism, openness, extroversion, agreeableness & conscientiousness [14]–[16].

C. Experimental Design

The analysis of asthma adherence was framed as a binary classification task with non-adherence (the target class) coded as zero. Missing values were handled using median imputation, and the issue of class imbalance was solved by downsampling [41], resulting in both low and high adherence classes containing 214 data points (N=428). Data was stratified into a training (75%) and test set (25%), with the performance of each model examined against the held-out test data. Classification accuracy was used as the standard metric against which results are assessed. Three classes of models (Logistic Regression, Support Vector Machines and Random Forests) were trained and evaluated. Meta-parameters for each model class were optimized via a grid search and 10-fold Cross-Validation[3], with training data being further split to obtain a 10% validation set for each fold. Once an optimal set of meta-parameters was identified, the corresponding model was

[3]The parameters searched over: LR: class weight: balanced, solver: liblinear, C: np.logspace(-4,3,20); SVC: kernel: (linear, rbf), C: (0.1,1,5,10,100), γ: (scale, auto, 0.1,0.01,0.001); RF: n_estimators: 2400, bootstrap: False, min_samples_split: (13,15,18,20,24,35), max_features: (sqrt, log2, 15).

then re-fit to the full training dataset and evaluated against the held out sample. While nested cross-validation could be employed, a split sample approach was preferred here due to the explanatory focus of the analysis and the subsequent goal of isolating an optimal reference model against which variable importance analysis could be performed. However, to confirm that the optimal reference model found was representative of generalized classification accuracy, the experiment was re-run in its entirety ten times to ensure the representativeness of the reference model's performance.

With an optimal model identified, we next investigated the factors underpinning non-adherence through the lens of feature importance analysis. Due to the grouped nature of the feature sets used, risk factors were assessed via permutation importance within their typed groups. In contrast to the proposal in [6], where group permutation importance scores are divided by the number of features in each variable set (thus 'normalizing' them), we instead report only total increase in error. This is motivated by the fact that our setting is not tasked with *minimal feature selection* (as per [6]), but the challenge of *maximal information collection*. That is, if practitioners must focus their data collection efforts on one pre-defined feature set, whether psychological, behavioural or physical, it is crucial to identify which ought be prioritised. This distinction is key, as the minimal feature selection problem favours inclusion of smaller groups of features per step - even if such groups contribute less predictive power than a larger group overall. Thus, for our problem setting, only total mean decrease in accuracy for permuted feature sets was assessed.

Permutation importance can lead to misleading results if practitioners are unaware that outputs only encode a single functional relationship between input feature and response feature - which is not necessarily representative of feature importance in the underlying phenomenon. MCR was introduced to help address this problem [31], requiring that analysis considers not one arbitrary solution, but all models that can provide optimal predictive performance. By computing MCR scores [7] for this Rashomon set, we are able to identify those features that contain indispensable information - and those which are readily substituted for another feature. As well as moving our understanding towards explanations of the underlying phenomenon being modelled, rather than the workings of an individual model instantiation, this allows us to mitigate against any feature selection bias that might occur during model building.

D. Grouped Model Class Reliance

To get closer to the underlying mechanism behind asthma non-adherence, we add functionality to existing MCR methods that allows analysis of groups of features within the framework (echoing the addition of grouping mechanisms to naive permutation methods developed in [6]). Group-MCR[4] is realized though a modification of the RF-MCR algorithm made available in [7]. Before developing our Group-MCR extension to MCR, we briefly introduce MCR below.

Model Class Reliance: Model Class Reliance (MCR) provides importance bounds for an input variable, highlighting the range of explanatory power each variable has across all possible equally predictive relationships consistent with the observational data (the phenomena) rather than just a single (typically arbitrary) relationship learnt by a machine learning algorithm. Conceptually the approach builds on permutation importance for a single model. Consider a set of data points Q. Each data point is defined as $q = (x_1, x_2, \ldots, x_v, y), q \in Q$ where each x_i denotes the value for the input feature, X_i, and y represents the value for the output feature, Y. The notion of Model Reliance (MR) of a specific feature (e.g. X_1) can then be defined as the mean decease in accuracy (MDA)[5], between a given model (functional relationship, f) when X_1 is available and when it is disrupted via permutation X_1^{ϕ}, with all other input variables $X_2, \ldots, X_n$ left unchanged:

$$MR_{X_1}(f) = \frac{1}{|Q|} \sum_{q \in Q} \mathbb{1}\left(f(x_1, x_2, \ldots, x_n) = y\right) - \frac{1}{|Q^{\phi}|} \sum_{q \in Q^{\phi}} \mathbb{1}\left(f(x_1^{\phi}, x_2, \ldots, x_n) = y\right) \quad (1)$$

Note that the use of the indicator function is simply recording if a model f is predicting the label of item q correctly. Intuitively the value of MR_{X_i} represents the damage done when the information held by X_i is rendered useless - reflecting its importance to the model. With this concept in hand, MCR can then be defined as the minimum ($MCR^{-}_{X_1}$) and maximum ($MCR^{+}_{X_1}$) MR observable across all of the equally predictive models producible by a given machine learning model class, and consistent with the observational data. The restriction to a model class enables all equally predictive relationships to be represented as the implicit set of models $F = (f_1, \ldots, f_m)$, resulting in the definition:

$$MCR(X_1) = \left[MCR^{-}_{X_1}, MCR^{+}_{X_1}\right] = \left[\min_{f \in F} MR_{X_1}(f), \max_{f \in F} MR_{X_1}(f)\right] \quad (2)$$

Algorithms and implementations for estimating MCR in a tractable fashion currently exist for Kernel Regression (polynomial run-time [31]) and Random Forests (classification and regression, logarithmic runtime [7]), with the latter being extended in this work.

Grouped Model Class Reliance: MCR computes the range of potential explanatory power for each individual variable - but tells us nothing about the importance of groups of variables in concert. Group-MCR is realized as an extension to the MCR for a Random Forests algorithm (RF-MCR) [7]. Noting that MCR is based on permutation importance, we follow a similar approach to the extension of permutation based methods for

[4]An open source implementation of Group-MCR is made available at the repository: https://github.com/gavin-s-smith/mcrforest

[5]For simplicity of explanation we focus on classification due to the problem type of this work. However, the approach holds more generally with the MR being defined as the change in any loss function, enabling easy generalization to regression and different performance measures (See [7], [31]).

single models from [6], proposing a modification to the MCR for Random Forest algorithm. While this technical contribution is an extension to an existing algorithm, its utility enables the novel analysis presented herein, and leads us to insights that would have otherwise been unobtainable.

Again let the (arbitrary) variable of interest be X_1. RF-MCR estimates $MCR^-_{X_1}$ ($MCR^+_{X_1}$) by defining a conceptual set of models with equal performance to a reference model assumed to represent an optimal model. A two staged search algorithm then searches across this conceptual set to locate the forest, $\hat{f}_{X_1}$, that utilizes X_1 the least (most). The MR for $\hat{f}_{X_1}$ is then the $MCR^-_{X_1}$ ($MCR^+_{X_1}$). At the level of an individual tree in the forest, Stage 1 assumes that all surrogate splits[6] for each decision node have been identified, and this is assumed for all trees. Specifically for each tree in the forest ($\forall T \in \hat{f}$), all decision nodes ($\forall D \in T$) in each tree have been annotated not only with a primary split variable (p) but also with all possible surrogate splits [42] - other input features that would split the data equally well ($S = \{s_1, s_2, \ldots\}$) with $D = (p, S)$. While all split variables (primary or surrogate) in practice have associated values, these are omitted in the notation for convenience. Based on this information, on a per tree basis, the tree which utilizes the variable the lowest (highest) is substituted to create the Stage 1 forest for $MCR^-_{X_1}$ ($MCR^+_{X_1}$). Let $a \in_R B$ denote the random selection of element a from set B and let p_D denote the primary split variable for a given decision node D in the tree. Formally for $MCR^-_{X_1}$, $\forall T \in \hat{f}, \forall D \in T$:

$$p_D = \begin{cases} p_D, & \text{if } x_1 \neq p_D \\ x \in_R (S_D - \{x_1\}), & \text{if } (S_D - \{x_i\}) \neq \emptyset \\ x_1, & otherwise \end{cases} \quad (3)$$

and for MCR^+, $\forall T \in \hat{f}, \forall D \in T$:

$$p_D = \begin{cases} x_1, & \text{if } x_1 \in p_D \cup S_D \\ p_D, & otherwise \end{cases} \quad (4)$$

For each substituted tree T', it is guaranteed that $MR_{X_1}(T') \leq MR_{X_1}(T)$ ($MR_{X_1}(T') \geq MR_{X_1}(T)$). This is because the trees remain structurally equivalent yet the permuted variable within $MR(\cdot)$ is utilized less (more) [7].

Once the Stage 1 forest has been built the final forest is constructed via, for each tree, searching across trees in the Stage 1 forest and replacing the tree, if a prediction equivalent tree that relies on X_1 less (more) than the current tree can be found. As shown in [7], combined, these transforms result in a final forest that represents the forest which relies on X_1 the least (most) and for which the MR equals the $MCR^-_{X_1}$ ($MCR^+_{X_1}$) under mild conditions.

To achieve Group-MCR we reconsider the definition of Model Reliance to be with respect to a group of variables, $G = X_1, \ldots, X_g$ where $g = |G|$ rather than a single variable X_1. Specifically:

$$MR_G(f) = \frac{1}{|Q|} \sum_{q \in Q} \mathbb{1}\left(f(x_1, \ldots, x_g, x_{g+1}, \ldots, x_n) = y\right) - \frac{1}{|Q^\phi|} \sum_{q \in Q^\phi} \mathbb{1}\left(f(x_1^\phi, \ldots, x_g^\phi, x_{g+1}, \ldots, x_n) = y\right) \quad (5)$$

This subsequently generalizes the GROUP-MCR definition to:

$$MCR(G) = \left[MCR^-_G, MCR^+_G\right] = \left[\min_{f \in F} MR_G(f), \max_{f \in F} MR_G(f)\right] \quad (6)$$

To compute GROUP-MCR we then follow the algorithm from [7] but modifying Equation 3 to, $\forall T \in \hat{f}, \forall D \in T$:

$$p_D = \begin{cases} p_D, & \text{if } p_D \notin G \\ x \in_R (S_D - G), & \text{if } (S_D - G) \neq \emptyset \\ p_D, & otherwise \end{cases} \quad (7)$$

and Equation 4 to, $\forall T \in \hat{f}, \forall D \in T$:

$$p_D = \begin{cases} p_D, & \text{if } p_D \in G \\ x \in_R (G \cap (\{p_D\} \cup S_D)), & \text{if } (G \cap (\{p_D\} \cup S_D)) \neq \emptyset \\ p_D, & otherwise \end{cases} \quad (8)$$

Note that under such modification, we retain the requirement that for each substituted tree T', that $MR_G(T') \leq MR_G(T)$ ($MR_G(T') \geq MR_G(T)$). Again, this is due to the fact the trees remain structurally equivalent and the permuted variables within $MR(G)$ are utilized more. Having made these changes, Stage 2 of the algorithm and its proofs are unaffected, allowing the final computation of MCR^-_G (MCR^+_G).

V. Results

Following parameterization via 10-fold cross validation of Logistic Regression, Support Vector Machine, and Random Forest model classes, and evaluation against the held-out test set, Random Forest models produced best predictive performance, with classification accuracy of 71% (precision = 0.73; recall = 0.69). SVM produced the second best result with the accuracy of 64% (precision = 0.67; recall = 0.56). Logistic Regression Classifier produced an accuracy of 63% (precision = 0.67; recall = 0.52). As the optimal classifier, the Random Forest model was first analysed using unconditional permutation importance [43], examining the importance of various feature groups to its predictive mechanism. Results of permutation importance are illustrated in Figure 1a (left). Perceptual factors provided the strongest influence on successful predictions of non-adherence, followed by emotions. Lifestyle choices and coping behaviours were of moderate importance to the model, with other features of only minor relevance to model outputs.

Next, Group-MCR analysis was applied, with Figure 1b illustrating the arbitrary nature of raw permutation importance

[6]Variables and their decision values that split the data in the same way as the variable chosen as the primary split [42]

analysis. Group-MCR analysis takes into account all equally predictive models from the model class, with importance boundaries being notably at odds to results in Figure 1a. First we note that, as expected (and by definition), the results of analysis of the reference model lie within the Group-MCR bounds. However, the Minimum Model Class Reliance (MCR-) and Minimum Model Class Reliance (MCR+) scores imply that there are clear differences between competing 'optimal' models in terms of the variable groups they leverage. This indicates that information sharing is occurring between the groups - in some instances feature groups are reflective of the same information, and could be use interchangeably. The only exception to this is the group of Perceptual factors, which not only contributes the most to any optimal model, but contains information (as indicated by a non-zero MCR- score) that is indispensable to prediction of adherence.

The perceptual factors, which showed the strongest importance to the model class, were then considered at an individual feature level via RF-MCR analysis. Individual features here correspond to individual questions from the survey, and have been ranked in Figure 2 with respect to their importance ranges across all models. Immediately, we recognize that much information is shared across questions, as one would expect given they all relate to patient perceptions of their condition. However, statements reflecting *Denial* are noticeable in achieving the highest available MCR+ levels: *'Even though I am diagnosed, I think I may not have asthma'* and *'My asthma is not as serious as my doctor and my diagnosis say it is'*. Another notable perception is *Perceived Discrimination*, as statements that illustrate this perceived stigma had a high MCR- and MCR+ score (e.g.*'I have been discriminated against at work because of my asthma'*).

In order to provide an indication of the relationship that exists between adherence scores and individual perceptions [44], Shapley additive explanation (SHAP) values were also computed based on the reference model. Figure 2b illustrates SHAP values for 28 individual features within the perceptions variable set, each item reflecting an individual question from the survey. Interpretation of these SHAP values indicates that predominantly, negative perceptions, indicative of discrimination and denial, correspond to low adherence scores (note that in the figure negative perceptual factors are predominantly indicated in red - but this is dependent on exact wording of individual questions).

VI. Discussion

Results have shown that Perceptual factors are the most important category of predictors indicative of non-adherence of asthma patients to prescribed medication. This conclusion is supported by both permutation importance and Group-MCR methods, with the groups' high MCR- score indicating that the information contained by perception features cannot be replaced by any other factors. Even in the predictive model where they are needed the least, perceptions still play a key role. Based on this it can be argued that perceptions patients create about their asthma are inseparable from the way they manage their condition - perceptions are essential and should be taken into account in all models.

The result that perceptions are primary predictors of adherence to asthma medication could be considered intuitive since perceptions relate to how we see the world. However, perceptions about one's condition would traditionally not be in the focus like other groups of factors such as demographics or lifestyle behaviours. This is why this provides valuable implications for future intervention strategies. The relationship of perceived stigma and non adherence indicates the likely benefit from a greater focus on actively addressing the reputation of asthma, which continues to be caricatured as a sign of 'weakness' in the media [35]. Identifying perceptions as the most relevant is also important because perceptions, such as stigma, often do not only originate from patients but are constantly reshaped by public opinions. Therefore, unlike for some other groups of factors that focus on patient-characteristics, such as demographics, interventions aimed at targeting perceptions should not only focus on communications to patients regarding stigma, but also on the wider public beyond asthma patients as well [35].

Group-MCR analysis also highlighted the existence of models in which some feature groups have an MCR- score of 0 (e.g. lifestyle choices). This insight emphasizes the dangers of only considering results of permutation importance of a single model. Permutation importance (Figure 1) implied 'lifestyle choices' could be considered as the third most relevant group. Yet, Group-MCR demonstrates this is not necessarily representative of the phenomena; some models simply do not need to use this group and still achieve the same prediction accuracy. The difference in terms of MCR- and MCR+ scores indicates that groups have intertwined relationships and likely share a significant proportion of information between each other. Instead of focusing on changing patients lifestyle, it is likely to be more impactful to focus on changing perceptions and emotions. The differences between graphs in Figure 1 also emphasize that when features share information, ranking of features can be highly arbitrary - demographics, for example, were unimportant in the reference model, but achieved the third highest MCR+ score, indicating that in some models demographics are highly valuable in predicting non-adherence.

In the breakdown of individual perception importances (see Figure 2), RF-MCR finds strongest predictive power to be reflected in the question: *'Even though I am diagnosed, I think I may not have asthma'*, illustrating *denial* of the condition [45]. SHAP values associated with statements about denial show the negative association denial has with adherence. This confirms previous qualitative analysis, which has identified denial as one of the most 'dangerous' aspects of stigma; whether the result of mental health challenges, or simply inability of a patient to accept the identity of being an 'asthmatic', denial can be a significant barrier to getting appropriate medical attention, care and support [4].

Many individual perceptual factors have a very low score for MCR- indicating they share a lot of information with other features (again, see Figure 2a). This is to be expected, as many

Fig. 1. Comparison of unconditional permutation importance (left) and Group-MCR (right).

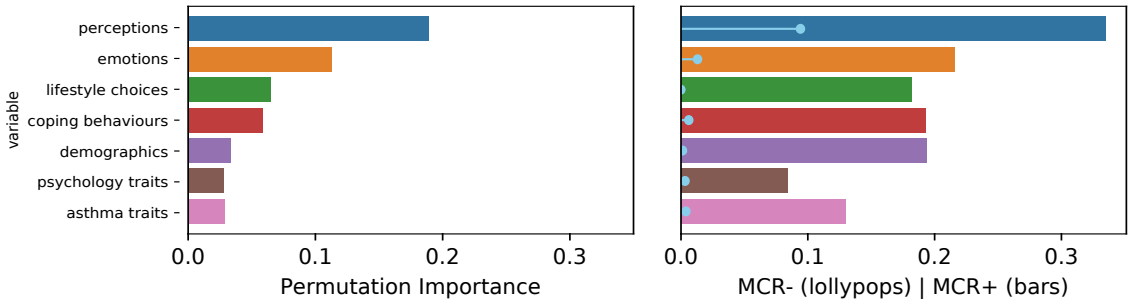

Fig. 2. Perception questions (features) only: MCR Plot (left) and a SHAP Summary Plot for a single, arbitrary, model from the Rashomon Set (right). Higher adherence is indicated by positive SHAP values, and non-adherence by negative SHAP values.

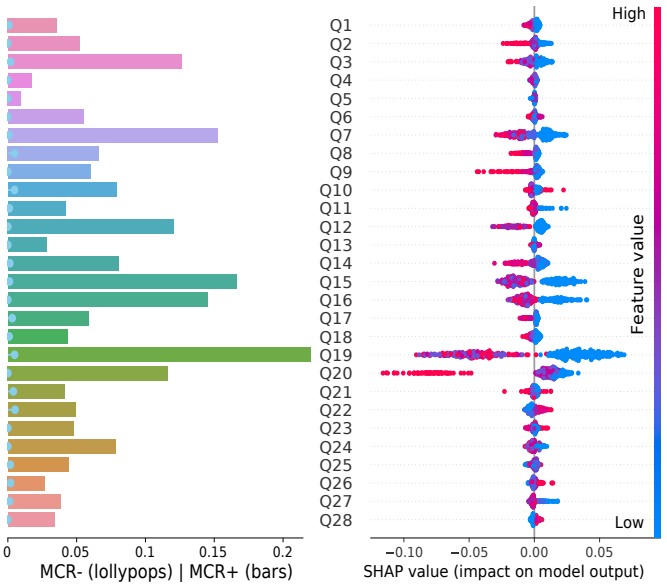

Q1. People sometimes joke about asthma to put me down.
Q2. People would make jokes about asthma in front of me.
Q3. I do not use my inhaler in public because people might make fun of me.
Q4. Media generally portrays asthma in a negative light.
Q5. Society perceives people with asthma as weak.
Q6. People would treat me differently if they knew I had asthma.
Q7. I have never been discriminated against at work because of asthma.
Q8. Sometimes I feel that I am being talked down to because of asthma.
Q9. I worry about telling people I use an inhaler.
Q10. Some people with asthma are weak.
Q11. I do not think people understand what asthma really is.
Q12. I think people with asthma are not as reliable co-workers as anybody else.
Q13. I worry about how people might react if they found out about my asthma.
Q14. I would have had better chances in life if I had not had asthma.
Q15. I do not tell people at my workplace that I have asthma.
Q16. I prefer if people did not see me using my inhaler.
Q17. I am angry with the way some people have reacted about my asthma.
Q18. I had trouble with people because of my asthma.
Q19. Even though I am diagnosed I think I may not have asthma.
Q20. My asthma is not as serious as my doctor and my diagnosis say it is.
Q21. Using an inhaler means you are not coping well with your asthma.
Q22. I would not ask others to change their behaviour...it is my own problem.
Q23. I get valuable information on how to cope with asthma from people online.
Q24. I feel more understood about asthma problems by people online...
Q25. I don't think that support groups for asthmatics are of any use to me.
Q26. I am concerned that I might get incorrect information...online.
Q27. Realizing that my experience is not unique helped me cope better.
Q28. I would rather suffer from smoke than explain that I have asthma.

questions in the survey are related to one another. This view is further supported by observed SHAP values (Figure 2b) which implies features mostly negatively impact adherence. Many perceptions examined in our models are related to stigmatization (including denial and discrimination) and these effects of stigma on non-adherence, and consequently on mortality rates, deserve further examination. The potential for perceptions of asthma patients, such as denial, to negatively impact adherence has an implication for intervention design, which currently remains focused on the physical components of asthma and related conditions [46]. That said, it is important to acknowledge the constraints on generalization of our findings. While the list of features (and groups) included in our model is based on hypothesized predictors from the literature, this is not an exhaustive list; increased dataset sizes would better serve our analysis; and a range of models, for which RF-MCR is not available, are yet to be considered.

VII. Conclusion

This work presented a study of a new asthma survey dataset, allowing investigation of combined groups of features, which have been independently proposed as related to asthma non-adherence. Considering intricate interactions between different features we considered their group's utility in predicting non-adherence using a novel Grouped Variable approach to Model Class Reliance (Group-MCR). This method overcomes challenges such as masking issues due to multicollinearity and the non-linear nature of relationships occurring between variable groups. Group-MCR Analysis also highlighted the risk of assuming that feature importances derived from a single 'optimal' model of non-adherence are fully representative.

Results indicated that perceptual factors were the strongest predictor of non-adherence. Additionally, investigation of individual features revealed that denial and perceived discrimination play a crucial role in non-adherence. This means that asthma is more than a pharmacological challenge. These insights can be used to develop better markers for non-adherence and open potential routes to tailored communications and services. The indispensability of perceptual factors produces a clear recommendation - policy should attempt to reduce both stigma and discrimination surrounding the condition, focusing on allaying denial and patient fears.

REFERENCES

[1] L. Stern, J. Berman, W. Lumry, L. Katz, L. Wang, L. Rosenblatt, and J. J. Doyle, "Medication compliance and disease exacerbation in patients with asthma," *Annals of Allergy, Asthma & Immunology*, vol. 97, no. 3, pp. 402–408, 2006.

[2] M. Axelsson, M. Emilsson, E. Brink, J. Lundgren, K. Torén, and J. Lötvall, "Personality, adherence, asthma control and health-related quality of life in young adult asthmatics," *Respiratory medicine*, vol. 103, no. 7, pp. 1033–1040, 2009.

[3] J. T. Lindsay and L. G. Heaney, "Nonadherence in difficult asthma–facts, myths, and a time to act," *Patient preference and adherence*, vol. 7, p. 329, 2013.

[4] B. Bender, H. Milgrom, and A. Apter, "Adherence intervention research: what have we learned and what do we do next?" *J. of Allergy & Clinical Immunology*, 2003.

[5] P. K. Kopalle and D. R. Lehmann, "Alpha inflation? the impact of eliminating scale items on cronbach's alpha," *Org. Behavior & Human Decision Proc.*, vol. 70, no. 3, 1997.

[6] B. Gregorutti, B. Michel, and P. Saint-Pierre, "Grouped variable importance with random forests and application to multiple functional data analysis," *Computational Statistics & Data Analysis*, vol. 90, pp. 15–35, 2015.

[7] G. Smith, R. Mansilla, and J. Goulding, "Model class reliance for random forests," in *Advances in Neural Information Processing Systems*, vol. 33, 2020.

[8] R. Horne, "Compliance, adherence, and concordance: implications for asthma treatment," *Chest*, vol. 130, no. 1, pp. 65S–72S, 2006.

[9] M. Olszanecka-Glinianowicz, J. Chudek, and A. Almgren-Rachtan, "Therapeutic preferences and factors determining the use of inhaled corticosteroids with long-acting β2-agonists in patients with asthma and chronic obstructive pulmonary disease," *Adv. in Dermatology & Allergology/Postepy Dermatologii i Alergologii*, 2020.

[10] C. Bolman, T. G. Arwert, and T. Völlink, "Adherence to prophylactic asthma medication: Habit strength and cognitions," *Heart & Lung*, vol. 40, no. 1, pp. 63–75, 2011.

[11] A. J. Apter, R. C. Boston, M. George, A. L. Norfleet, T. Tenhave, J. C. Coyne, K. Birck, S. T. Reisine, A. J. Cucchiara, and H. I. Feldman, "Modifiable barriers to adherence to inhaled steroids among adults with asthma: it's not just black and white," *Journal of Allergy and Clinical Immunology*, vol. 111, no. 6, pp. 1219–1226, 2003.

[12] G. Roberts, "Predicting the long-term outcome of preschool wheeze: are we there yet?" *Journal of Allergy and Clinical Immunology*, vol. 124, no. 5, pp. 911–912, 2009.

[13] B. Granger, I. Ekman, C. Granger, J. Ostergren, B. Olofsson, E. Michelson, J. McMurray, S. Yusuf, M. Pfeffer, and K. Swedberg, "Adherence to medication according to sex and age in the charm programme," *Euro. J. of Heart Failure*, 2009.

[14] P. T. Costa and R. R. McCrae, *The NEO personality inventory*. Psychological Assessment Resources Odessa, FL, 1985.

[15] L. R. Goldberg, "The development of markers for the big-five factor structure." *Psychological assessment*, vol. 4, no. 1, p. 26, 1992.

[16] R. R. McCrae and O. P. John, "An introduction to the five-factor model and its applications," *J. Personality*, vol. 60, no. 2, pp. 175–215, 1992.

[17] M. Axelsson, E. Brink, and J. Lötvall, "A personality and gender perspective on adherence and health-related quality of life in people with asthma and/or allergic rhinitis," *J. of the American Association of Nurse Practitioners*, vol. 26, no. 1, pp. 32–39, 2014.

[18] C. S. Stilley, S. Sereika, M. F. Muldoon, C. M. Ryan, and J. Dunbar-Jacob, "Psychological and cognitive function: predictors of adherence with cholesterol lowering treatment," *Annals of Behavioral Medicine*, vol. 27, no. 2, pp. 117–124, 2004.

[19] C. M. Brooks, J. M. Richards, C. L. Kohler, S.-J. Soong, B. Martin, R. A. Windsor, and W. C. Bailey, "Assessing adherence to asthma medication and inhaler regimens: a psychometric analysis of adult self-report scales," *Medical care*, pp. 298–307, 1994.

[20] Y. Kim, J. P. Dillard, and R. A. Smith, "Communicating antibiotic stewardship: emotional responses and their impact on adherence," *Health communication*, 2019.

[21] L.-P. Boulet, "Perception of the role and potential side effects of inhaled corticosteroids among asthmatic patients," *Chest*, vol. 113, no. 3, pp. 587–592, 1998.

[22] M. Axelsson, E. Brink, J. Lundgren, and J. Lötvall, "The influence of personality traits on reported adherence to medication in individuals with chronic disease: an epidemiological study in west sweden," *PloS one*, vol. 6, no. 3, p. e18241, 2011.

[23] C. A. Mancuso, W. Sayles, and J. P. Allegrante, "Knowledge, attitude, and self-efficacy in asthma self-management and quality of life," *Journal of Asthma*, vol. 47, no. 8, pp. 883–888, 2010.

[24] C. S. Ulrik, V. Backer, U. Søes-Petersen, P. Lange, H. Harving, and P. P. Plaschke, "The patient's perspective: adherence or non-adherence to asthma controller therapy?" *Journal of Asthma*, vol. 43, no. 9, pp. 701–704, 2006.

[25] M. R. DiMatteo, "Variations in patients' adherence to medical recommendations: a quantitative review of 50 years of research," *Medical care*, pp. 200–209, 2004.

[26] S. MÄkinen, T. Suominen, and S. Lauri, "Self-care in adults with asthma: how they cope," *Journal of Clinical Nursing*, vol. 9, no. 4, pp. 557–565, 2000.

[27] M. Olszanecka-Glinianowicz and A. Almgren-Rachtan, "The adherence and illness perception of patients diagnosed with asthma or chronic obstructive pulmonary disease treated with polytherapy using new generation cyclohaler," *Advances in Dermatology and Allergology/Postepy Dermatologii i Alergologii*, vol. 31, no. 4, p. 235, 2014.

[28] L. Van Der Maaten, E. Postma, and J. Van den Herik, "Dimensionality reduction: A comparative review," *J Mach Learn Res*, vol. 10, 2009.

[29] R. R. Wilcox, "Robust generalizations of classical test reliability and cronbach's alpha," *British J. of Mathematical and Statistical Psychology*, vol. 45, no. 2, pp. 239–254, 1992.

[30] P. E. Green and V. R. Rao, "Rating scales and information recovery—how many scales and response categories to use?" *Journal of Marketing*, vol. 34, no. 3, pp. 33–39, 1970.

[31] A. Fisher, C. Rudin, and F. Dominici, "All models are wrong, but many are useful: Learning a variable's importance by studying an entire class of prediction models simultaneously," *J. of Machine Learning Research*, vol. 20, no. 177, pp. 1–81, 2019.

[32] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems*, 2017.

[33] D. E. Morisky, "Predictive validity of a medication adherence measure for hypertension control," *Journal of clinical hypertension*, vol. 10, pp. 348–354, 2008.

[34] M. Thomas, K. Gruffydd-Jones, C. Stonham, S. Ward, and T. Macfarlane, "Assessing asthma control in routine clinical practice: use of the royal college of physicians '3 questions'," *Primary Care Respiratory Journal*, vol. 18, no. 2, pp. 83–88, 2009.

[35] K. L. Andrews, S. C. Jones, and J. Mullan, "Stigma: Still an important issue for adults with asthma," *Journal of Asthma & Allergy Educators*, vol. 4, no. 4, pp. 165–171, 2013.

[36] G. S. O'Keeffe, K. Clarke-Pearson *et al.*, "The impact of social media on children, adolescents, and families," *Pediatrics*, vol. 127, no. 4, pp. 800–804, 2011.

[37] T. E. Ford and M. A. Ferguson, "Social consequences of disparagement humor: A prejudiced norm theory," *Personality and social psychology review*, vol. 8, no. 1, pp. 79–94, 2004.

[38] A.-M. Aalto, K. Härkäpää, A. R. Aro, and P. Rissanen, "Ways of coping with asthma in everyday life: validation of the asthma specific coping scale," *J. Psychosomatic Research*, vol. 53, no. 6, pp. 1061–1069, 2002.

[39] D. Watson, L. A. Clark, and A. Tellegen, "Development and validation of brief measures of positive and negative affect: the panas scales." *Journal of personality and social psychology*, vol. 54, no. 6, p. 1063, 1988.

[40] H. Al-Qazaz, M. Hassali, A. Shafie, S. Sulaiman, S. Sundram, and D. Morisky, "The eight-item morisky medication adherence scale mmas: translation and validation of the malaysian version," *Diabetes Research & clinical practice*, vol. 90, no. 2, pp. 216–221, 2010.

[41] S. Liu, M. L. Ong, K. K. Mun, J. Yao, and M. Motani, "Early prediction of sepsis via smote upsampling and mutual information based downsampling," in *2019 Computing in Cardiology (CinC)*. IEEE, 2019.

[42] L. Breiman, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.

[43] A. Altmann, L. Toloşi, O. Sander, and T. Lengauer, "Permutation importance: a corrected feature importance measure," *Bioinformatics*, vol. 26, no. 10, pp. 1340–1347, 2010.

[44] K. Futagami, Y. Fukazawa, N. Kapoor, and T. Kito, "Pairwise acquisition prediction with shap value interpretation," *The Journal of Finance and Data Science*, 2021.

[45] L. M. Osman, "Psychological factors in asthma control and attack risk," *Thorax*, 2002.
[46] G. J. Johnson and P. J. Ambrose, "Neo-tribes: The power and potential of online communities in health care," *Communications of the ACM*, vol. 49, no. 1, pp. 107–113, 2006.