



groupShapley: Efficient prediction explanation with Shapley values for feature groups

Note no
Authors

SAMBA/20/21
Martin Jullum
Annabelle Redelmeier
Kjersti Aas

Date

Feb 20th, 2021

Norwegian Computing Center

Norsk Regnesentral (Norwegian Computing Center, NR) is a private, independent, non-profit foundation established in 1952. NR carries out contract research and development projects in information and communication technology and applied statistical-mathematical modelling. The clients include a broad range of industrial, commercial and public service organisations in the national as well as the international market. Our scientific and technical capabilities are further developed in co-operation with The Research Council of Norway and key customers. The results of our projects may take the form of reports, software, prototypes, and short courses. A proof of the confidence and appreciation our clients have in us is given by the fact that most of our new contracts are signed with previous customers.

Title **groupShapley: Efficient prediction explanation with Shapley values for feature groups**

Authors **Martin Jullum** <jullum@nr.no>
Annabelle Redelmeier <anr@nr.no>
Kjersti Aas <kjersti@nr.no>

Date Feb 20th, 2021

Publication number SAMBA/20/21

Abstract

Shapley values has established itself as one of the most appropriate and theoretically sound frameworks for explaining predictions from complex machine learning models. The popularity of Shapley values in the explanation setting is probably due to its unique theoretical properties. The main drawback with Shapley values, however, is that its computational complexity grows exponentially in the number of input features, making it unfeasible in many real world situations where there could be hundreds or thousands of features. Furthermore, with many (dependent) features, presenting/visualizing and interpreting the computed Shapley values also becomes challenging. The present paper introduces *groupShapley*: a conceptually simple approach for dealing with the aforementioned bottlenecks. The idea is to group the features, for example by type or dependence, and then compute and present Shapley values for these groups instead of for all individual features. Reducing hundreds or thousands of features to half a dozen or so, makes precise computations practically feasible and the presentation and knowledge extraction greatly simplified. We prove that under certain conditions, *groupShapley* is equivalent to summing the feature-wise Shapley values within each feature group. Moreover, we provide a simulation study exemplifying the differences when these conditions are not met. We illustrate the usability of the approach in a real world car insurance example, where *groupShapley* is used to provide simple and intuitive explanations.

Keywords Shapley values, grouping, prediction explanation

Target group ML/XAI researchers

Availability Open

Project Big Insight Explanations

Project number 220865

Research field XAI

Number of pages 23

© Copyright Norwegian Computing Center

Contents

1	Introduction	5
1.1	Challenges with the Shapley value framework	5
1.2	The present paper	6
2	groupShapley	7
2.1	Defining feature groups	8
2.2	Other types of grouped Shapley values	9
2.3	Comparing groupShapley with post-grouped Shapley	10
3	Simulations	11
3.1	Experiment 1: linear regression model	12
3.2	Experiment 2: generalized additive model	13
3.3	Experiment 3: different dependence within and between feature groups	13
4	Real data example	15
5	Concluding remarks	16
	References	17
A	Proofs	20
A.1	Proof of Theorem 2.1	21
A.2	Proof of Theorem 2.2	23

1 Introduction

Consider a predictive modelling/machine learning setting with a given model $f(\cdot)$ that takes an M dimensional feature vector $\mathbf{x} = \{x_1, \dots, x_M\}$ as input and provides a prediction $f(\mathbf{x})$ of an unknown response y . Suppose that for a specific feature vector, $\mathbf{x} = \mathbf{x}^*$, we want to understand how the different features (or types of features) contribute to the specific prediction outcome $f(\mathbf{x}^*)$. This task is called prediction explanation and is a type of *local* model explanation, as opposed to a *global* model explanation, which attempts to explain the full model at once, through concepts such as global feature importance (Molnar, 2020, Ch. 2).

Shapley values (Shapley, 1953) has established itself as one of the leading frameworks for prediction explanation. The methodology has, in particular, received increased interest following the seminal paper of Lundberg and Lee (2017). The feature-wise Shapley values ϕ_1, \dots, ϕ_M for a predictive model $f(\mathbf{x})$ at $\mathbf{x} = \mathbf{x}^*$ are given by

$$\phi_j = \sum_{\mathcal{S} \subseteq \mathcal{M} \setminus \{j\}} \frac{|\mathcal{S}|!(M - |\mathcal{S}| - 1)!}{M!} (v(\mathcal{S} \cup \{j\}) - v(\mathcal{S})), \quad (1)$$

for $j = 1, \dots, M$, where M is the total number of features, \mathcal{M} is the set of all features, and $v(\cdot)$ is the characteristic function, also referred to as the *contribution function* in the prediction explanation setting. Roughly speaking, the information obtained by observing feature j modifies the prediction by the amount of its Shapley value, ϕ_j . Although there are different choices available for specifying the contribution function, the conditional expectation

$$v(\mathcal{S}) = \mathbb{E}[f(\mathbf{x}) | \mathbf{x}_{\mathcal{S}} = \mathbf{x}_{\mathcal{S}}^*], \quad (2)$$

where $\mathbf{x}_{\mathcal{S}}$ denotes the subvector of \mathbf{x} corresponding to feature subset \mathcal{S} , has the nice property of providing explanations which properly account for the dependence between the features (Chen et al., 2020). This specific contribution function was originally the idea of Lundberg and Lee (2017) and has since been used by several others (Aas et al., 2019, 2021; Frye et al., 2020; Redelmeier et al., 2020).

1.1 Challenges with the Shapley value framework

There are, in principle, two challenges when it comes to the practical use of Shapley values for prediction explanation. Since the true (conditional) distribution of \mathbf{x} is rarely known, the first challenge is estimating (2) precisely and efficiently. Poor estimates may lead to severely inaccurate Shapley values and, therefore, also explanations and conclusions (Aas et al., 2019). We will not be dealing with this issue in the present paper and, instead, refer to Aas et al. (2019); Frye et al. (2020) for work along these lines.

Instead, this paper concerns itself with the second challenge of handling the fundamental computational complexity of the Shapley formula, which contains 2^{M-1} terms and therefore grows exponentially in the number of features M . Hence, even with limited time spent on appropriately estimating (2), the computational burden of the arithmetic

operations in (1) becomes intractable when there are more than one or two dozen features (depending on the computational resources and time available).

There exists some model-agnostic approaches for approximating the Shapley value formula. One alternative is to, instead, sum over a smaller number of sampled feature subsets \mathcal{S} (Covert and Lee, 2020; Lundberg and Lee, 2017). However, as the number of features grows, the number of sampled subsets also needs to grow (exponentially) in order to retain an acceptable accuracy for the approximated Shapley values. Thus, this approach only minimally lifts the roof on the number of features that can be handled. In any case, when there are tens, hundreds or even thousands of features, which is not uncommon in the modern machine learning world, these approaches simply do not do. Chen et al. (2019) and Li et al. (2019) also propose approaches for reducing the complexity of the Shapley value computation. These approaches are very different from the direct sampling approach above in that they assume that the contribution of features to the prediction respects the structure of an underlying graph, and are therefore not generally applicable. There are, however, some conceptual similarities with the method proposed in the present paper. We will return to this in Section 2.3.

When the number of features increases, it also becomes harder to visualize and extract information from Shapley values. This is especially the case when some of the features are (highly) dependent, which in most applied fields is the rule rather than the exception. When features are highly dependent, the joint contribution of these features is distributed among the features, resulting in many small Shapley values. Thus, interpreting the Shapley values and transferring this to practical knowledge requires understanding of the feature dependence structure and appropriate modifications of the interpretation. This clearly becomes an overwhelming task when there are more than a few features.

Note that some Shapley based prediction explanation approaches (Janzing et al., 2020; Sundararajan and Najmi, 2020) replace the contribution function $v(\mathcal{S})$ in (2) by a so-called interventional conditional expectation (Chen et al., 2020) which is easier to compute and, therefore, reduces the computational burden. However, the computational complexity still grows exponentially in the number of features; it just kicks in a bit later. Furthermore, these methods have the drawback of ignoring the feature dependence, which may lead to inconsistent explanations in the presence of feature dependence. There also exists algorithms for specific model classes, like TreeSHAP (Lundberg et al., 2020) for tree-based models, which uses the feature dependence learned by the tree(s) to estimate the conditional expectations and efficiently compute Shapley values. However, as seen in Aas et al. (2019), the resulting Shapley values may be highly inaccurate. Since the method is model-specific, it is also difficult to use in a development/testing stage where models of different types are compared.

1.2 The present paper

In this paper we propose a novel approach, *groupShapley*, to bypass the aforementioned computational issue of explaining predictions using the Shapley value framework. The idea is to move away from computing Shapley values for single features and rather compute and present Shapley values for *groups of features*. Our definition of grouped Shapley

values simply replaces the individual features in (1) by feature groups, creating perfectly well-defined Shapley values with all of the usual Shapley properties. The computational complexity is kept small as long as the number of groups is small. In addition, both the presentation/visualization and the interpretation and knowledge extraction in the presence of high feature dependence, are simplified. Consider e.g. the case when a few base features/data sources are used to construct a myriad of engineered features. Then, *groupShapley* can help improve the interpretability of this model by only presenting Shapley values for the original base features/data sources.

The rest of the paper is organized as follows: Section 2 provides the mathematical definition of *groupShapley* and discusses how groups can be constructed in different practical settings. We discuss and compare *groupShapley* with other ways of establishing a Shapley based measure for the contribution of a group of features. Moreover, we present a theorem saying that under certain conditions, *groupShapley* is equivalent to summing the feature-wise Shapley values within each feature group. Under the same conditions we also derive a simplified formula for the feature-wise Shapley values. In Section 3, we provide a simulation study investigating how *groupShapley* differs from the alternative method under different violations of the conditions where they are equal. In Section 4, we give a real world example using a car insurance data set. Finally, we provide some general discussion and concluding remarks in Section 5. The supplementary material accompanying the present paper contains proofs of the theoretical results from Section 2.

2 groupShapley

Recall the predictive modelling/machine learning setting from Section 1 with the predictive model $f(\cdot)$ producing predictions based on an M -dimensional feature vector \mathbf{x} that one wishes to explain. Let us now define a partition¹ $\mathcal{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_G\}$ of the feature set \mathcal{M} , consisting of G groups. Then, the Shapley value for the i -th group \mathcal{G}_i explaining the prediction $f(\mathbf{x}^*)$ is given by

$$\phi_{\mathcal{G}_i} = \sum_{\mathcal{T} \subseteq \mathcal{G} \setminus \mathcal{G}_i} \frac{|\mathcal{T}|_g!(G - |\mathcal{T}|_g - 1)!}{G!} (v(\mathcal{T} \cup \mathcal{G}_i) - v(\mathcal{T})), \quad (3)$$

where the summation index \mathcal{T} runs over the groups (not the individual features) in the set of groups $\mathcal{G} \setminus \mathcal{G}_i$ and $|\mathcal{T}|_g$ refers to the number of groups (not the individual elements) in \mathcal{T} . *groupShapley* also adopts the contribution function (2), meaning that

$$v(\mathcal{T}) = \mathbb{E}[f(\mathbf{x}) | \mathbf{x}_{\mathcal{T}} = \mathbf{x}_{\mathcal{T}}^*], \quad (4)$$

where $\mathbf{x}_{\mathcal{T}}$ denotes the subvector of \mathbf{x} corresponding to all features contained in all groups in \mathcal{T} . Thus, any procedure applicable for estimating (2), can also be directly applied to estimate (4). As *groupShapley* is simply the game theoretic Shapley value framework applied directly to groups of features instead of individual features, the *groupShapley* values

1. A partition is a grouping of non-empty subsets of the elements of a set where each element is represented in exactly one group.

possess all the regular Shapley value properties. In particular, the *efficiency property* states that $\sum_{i=1}^G \phi_{G_i} = f(\mathbf{x}^*) - \mathbb{E}[f(\mathbf{x})]$, the *symmetry property* roughly states that groups which contribute equally (regardless of the influence of other feature groups) have identical Shapley values, and the *null player property* states that feature groups with no contribution to the prediction (either directly or through other features which it correlates with) have a Shapley value of zero.

By directly comparing the formula for the feature-wise Shapley values in (1) with the *groupShapley* formula in (3), we see that the computational complexity of the sum reduces from 2^{M-1} to 2^{G-1} , which gives a relative computational cost reduction of 2^{M-G} . With, for example, $M = 50$ features and $G = 5$ feature groups, the relative cost reduction is $> 10^{13}$. Naturally, in some settings, reducing feature-based explanation to group-based explanation, may not provide a sufficient detail or disguise single feature contributions of interest. However, for large number of features, the benefits of *groupShapley* far outweigh the risk of over-simplification.

2.1 Defining feature groups

With the general *groupShapley* procedure established in (3), the next question is how to define these groups. We propose two alternative strategies: Grouping based on 1) feature type, and 2) feature dependence. These two grouping strategies can be viewed as grouping based on the *practical* properties or the *theoretical* properties of the features, respectively. Note that there is no gold standard for grouping features and it may be informative to compute several explanations of the same predictions using different groupings.

The first strategy groups together features with similar practical meaning, origin or definition. Consider, for instance, a housing price prediction model with a wide selection of features. Then, groups made according to feature *type* could be *house attributes* (housing type, square footage, number of rooms, building year), *luxury amenities* (has hot tub/swimming pool/wine cellar/home theatre), *locality* (distance to nearest grocery store/public transport/gym), *environmental/surrounding factors* (area crime rate/air pollution/traffic noise, balcony sunshine hours), and *historical turnover* (previous sold price, details of any refurbishing). As a second example, consider a model that predicts whether a customer purchases a specific item, e.g. as part of a recommendation system. Then, groups could be specified as follows: *customer details* (gender, age, occupation), *customer purchasing history* (previous purchases, last time item bought), *date/time* (time of day purchases made, day of the week, any upcoming holiday), and *pricing* (item price compared to other brands/stores, current sale, and quantity discount). In these cases, using *groupShapley* may inform that the *locality* of a very central apartment has increased its predicted price, or that it was the *day/time* that reduced the probability of making a specific purchase on a Friday night. In Section 4, we provide a real data example within car insurance where feature type grouping is used.

The second strategy groups together features based on mutual dependence. The point of this approach is to simplify the interpretability of the computed Shapley values. As mentioned in Section 1, the joint contribution of highly dependent features is spread out on the feature-wise Shapley values making it more difficult to detect the significance of

these features. If such highly dependent features are grouped together, it is easier to spot these contributions and, therefore, infer a more ‘correct’ explanation. Such a grouping can be created through a clustering method that uses a dependence/correlation based similarity measure, see e.g. [Kaufman and Rousseeuw \(2009\)](#). See also the treatment in [Aas et al. \(2019, Sec. 5\)](#).

2.2 Other types of grouped Shapley values

In the general game theoretic setting, using Shapley values to determine the influence of groups of players is not entirely new. [Marichal et al. \(2007\)](#) defines the concept of a generalized Shapley value for a set of players and studies its theoretical properties. [Al-gaba et al. \(2019, Ch. 13\)](#) refers to this as the Shapley group value, studies the behavior of the players when acting as a group rather than individually, and applies it to social network ‘games’. An important distinction between this methodology and *groupShapley* is that they study how a single group of players acts relative to the other *individual* players, while we compare groups with other groups. This makes their method suitable for studying concepts such as which players gain from joining a cooperative group and which are stronger on their own, rather than comparing the relative performance of disjoint groups. Furthermore, their formulation still has a computational complexity that grows exponentially in the number of players/features. We will, therefore, not devote more space to this method.

An alternative way to define the contribution of a group of features based on the Shapley value framework is to utilize the additive structure of Shapley values and simply take the sum of the feature-wise Shapley values in (1) for the features within each group. That is, for each group $\mathcal{G}_i, i = 1, \dots, G$, we define

$$\phi_{\text{post-}\mathcal{G}_i} = \sum_{j \in \mathcal{G}_i} \phi_j, \quad (5)$$

where the summation index j runs over the features in group \mathcal{G}_i and ϕ_j is the Shapley value for feature j computed using (1). We call this *post-grouped Shapley* since we explicitly calculate the Shapley values for *individual features* first and then group these Shapley values afterwards. In Section 2.3, we show that under certain conditions, *groupShapley* and *post-grouped Shapley* are identical. However, this is generally not the case, as exemplified through a series of simulation experiments in Section 3. Since *groupShapley* provides proper Shapley values when using (4) and Shapley values are unique, *post-grouped Shapley* does *not* in general correspond to proper Shapley values using (4). Moreover, (5) is inherently impractical to compute when there are many features, as it requires the computation of all feature-wise Shapley values – exactly the bottleneck we are trying to avoid. That being said, (5) is a natural approach to calculating group contributions that has been used before ([Aas et al., 2019](#); [Redelmeier et al., 2020](#)). Therefore, we will use *post-grouped Shapley* as a reference method when studying the properties of *groupShapley*. Note, however, that there is no reason to prefer *post-grouped Shapley* since it does not provide proper Shapley values and takes longer to compute.

2.3 Comparing groupShapley with post-grouped Shapley

In this section we provide two conditions and then show two different theoretical results that hold under the conditions: 1) The feature-wise Shapley values take a simplified form and 2) *groupShapley* and *post-grouped Shapley* are equivalent, i.e. $\phi_{\mathcal{G}_i} = \phi_{\text{post-}\mathcal{G}_i}$, for $i = 1, \dots, G$. We first state the two conditions and then present the results through two theorems. The proofs of the theorems are given in the supplementary material.

Condition 2.1 (Partially additively separable). *The predictive function $f(\mathbf{x})$ is partially additively separable with respect to the partition $\mathcal{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_G\}$, i.e. we may write*

$$f(\mathbf{x}) = \sum_{i=1}^G f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i}),$$

where for all $\mathcal{G}_i \in \mathcal{G}$, $\mathbf{x}_{\mathcal{G}_i}$ is the subset of \mathbf{x} corresponding to the features in subset \mathcal{G}_i and $f_{\mathcal{G}_i}$ is a real function involving only $\mathbf{x}_{\mathcal{G}_i}$.

Condition 2.2 (Group independence). *All feature groups in partition \mathcal{G} are independent ($\mathbf{x}_{\mathcal{G}_i} \perp\!\!\!\perp \mathbf{x}_{\mathcal{G}_j}$ for all $i \neq j$), meaning that all features in \mathcal{G}_i are mutually independent of all the features in \mathcal{G}_j , for all $i \neq j$.*

Before turning to the theorems, let us introduce the contribution function corresponding to the sub-functions from Condition 2.1: $v_{\mathcal{G}_i}(\mathcal{S}) = \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i}) | \mathbf{x}_{\mathcal{S}} = \mathbf{x}_{\mathcal{S}}^*]$ for any $\mathcal{S} \subseteq \mathcal{G}_i$.

Theorem 2.1 (Simplified feature-wise Shapley formula). *Assume Conditions 2.1 and 2.2 hold. Then for any $j \in \mathcal{G}_i$, $i = 1 \dots G$, the Shapley formula in (1) simplifies to*

$$\phi_j = \sum_{\mathcal{S} \subseteq \mathcal{G}_i \setminus \{j\}} \frac{|\mathcal{S}|!(|\mathcal{G}_i| - |\mathcal{S}| - 1)!}{|\mathcal{G}_i|!} (v_{\mathcal{G}_i}(\mathcal{S} \cup \{j\}) - v_{\mathcal{G}_i}(\mathcal{S})), \quad (6)$$

i.e. the $\phi_j, j \in \mathcal{G}_i$ are identical to the Shapley values for explaining the predictive model $f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})$.

Theorem 2.1 does not only provide a simplified formula for the feature-wise Shapley values which reduces the computation complexity from 2^{M-1} to $2^{|\mathcal{G}_i|-1}$ under the stated conditions. It also says that under the stated conditions, the Shapley values for $f(\mathbf{x}^*)$ can actually be obtained by computing the Shapley values for $f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})$ based only on the $|\mathcal{G}_i|$ features relevant for that function, for $i = 1, \dots, G$. An important special case of Theorem 2.1 is when a single feature j is independent of all the other features and joins the predictive model formula $f(\mathbf{x})$ in a purely additive way. Then, $\phi_j = f_j(x_j^*) - \mathbb{E}[f_j(x_j)]$, which generalizes a well known simplification of the Shapley value formula for linear models with independent features, see e.g. Aas et al. (2019, Section 2.2). The formula in (6) of Theorem 2.1 is conceptually similar to a simplified Shapley value formula in Li et al. (2019). They do, however, work with a different contribution function and rely on an assumed underlying graph structure.

Returning to feature groups, the theorem below states the special case equivalence between *groupShapley* and *post-grouped Shapley*.

Theorem 2.2 (groupShapley equivalence). *Assume Conditions 2.1 and 2.2 hold. Then,*

$$\phi_{\text{post-}\mathcal{G}_i} = \phi_{\mathcal{G}_i} = v_{\mathcal{G}_i}(\mathcal{G}_i) - v_{\mathcal{G}_i}(\emptyset), \quad i = 1, \dots, G,$$

i.e. *post-grouping in terms of (5) is equivalent to groupShapley.*

Theorem 2.2 does not only say that *groupShapley* and *post-grouped Shapley* are equivalent under the stated conditions, but also that the Shapley values are very easy to compute. Conditions 2.1 and 2.2 are, of course, rarely met in practical situations. It is important to note that this does not, in any way, invalidate *groupShapley* or *post-grouped Shapley*; it just means that the two approaches do not necessarily always yield identical explanations.

3 Simulations

We perform a simulation study with the goal of uncovering differences between group values calculated under *groupShapley* and *post-grouped Shapley*. The idea is to start with a model satisfying Conditions 2.1 and 2.2 and then show the differences between *groupShapley* and *post-grouped Shapley* as we violate these conditions. We defy the *partially additively separable* condition by including pair-wise linear and non-linear interactions in our models. We defy the *group independence* condition by adding dependence between features that belong to different groups. Note that Conditions 2.1 and 2.2 are *sufficient*. Thus, if the conditions do not hold, this does *not* necessarily mean that *groupShapley* and *post-grouped Shapley* will yield different results. We explore this in more detail below.

There are three versions of the simulation study: Experiment 1 uses a linear regression model; Experiment 2 uses a generalized additive model; and Experiment 3 uses both of these models and a slightly different feature covariance matrix. The full parameters of the simulation study are:

- *Underlying feature distribution model*: We simulate 10 features with a multivariate Gaussian distribution $p(\mathbf{x}) = N_{10}(\mathbf{0}, \Sigma)$. This makes it easy to insert dependence between features.
 - Σ : The covariance matrix of the joint Gaussian distribution when simulating the features \mathbf{x} . For Experiments 1 and 2, we set the variance to 1 and the correlation ρ which takes values in $\{0, 0.1, 0.3, 0.7, 0.9\}$. For Experiment 3, we use a slightly different Σ discussed in that section.
- n_{test} : The number of testing observations. Set to $n_{test} = 100$.
- *Groups*: We put the 10 features into two distinct groupings/partitions. \mathcal{G}_A has $G = 3$ groups: $\mathcal{G}_1 = \{1, 2, 3, 4\}$, $\mathcal{G}_2 = \{5, 6, 7, 8\}$, and $\mathcal{G}_3 = \{9, 10\}$. \mathcal{G}_B has $G = 5$ groups: $\mathcal{G}_1 = \{1, 2\}$, $\mathcal{G}_2 = \{3, 4\}$, $\mathcal{G}_3 = \{5, 6\}$, $\mathcal{G}_4 = \{7, 8\}$, and $\mathcal{G}_5 = \{9, 10\}$.
- *Predictive model*: We use two different types of predictive models: a linear regression model (lm) and a generalized additive model (GAM). Both lm and GAM are used to fit three different models: one model without interactions, one model with interactions between features of the *same group*, and one model with interactions between features of *different groups*. We discuss each response function in more detail in Sections 3.1 and 3.2.

Once the features are simulated, we compute the Shapley values of each group using both (3) and (5). We compare the results between (3) and (5) using the *Mean Absolute Deviation* (MAD). Assuming G groups, the MAD for individual i is defined as:

$$\text{MAD}_i = \frac{1}{G} \sum_{j=1}^G |\phi_{\text{pre-}\mathcal{G}_j}(\mathbf{x}_i) - \phi_{\text{post-}\mathcal{G}_j}(\mathbf{x}_i)|, \quad (7)$$

where $\phi_{\text{pre-}\mathcal{G}_j}(\mathbf{x}_i)$ and $\phi_{\text{post-}\mathcal{G}_j}(\mathbf{x}_i)$ denote, respectively, the value estimated with *group-Shapley* and the value estimated with *post-grouped Shapley*. When calculating the *group-Shapley* and *post-grouped Shapley* values, we use the true Gaussian distribution to compute (4) and (2), since we are not concerned with the estimation of those. We use Monte Carlo integration with 1000 samples from the true Gaussian distribution to compute the conditional expectations. This introduces some randomness in our empirical results, and is the reason that the MAD is not exactly zero under Conditions 2.1 and 2.2, see Figures 1-3.

To be able to compare the MAD between different experiments and parameters, the predictive functions specified in the next subsections are standardized such that all Shapley value computations are performed on predictive functions with standard deviation 1 (the mean values cancel each other in the MAD formula).

3.1 Experiment 1: linear regression model

In this experiment, we explore how *groupShapley* and *post-grouped Shapley* vary when using a linear regression model. We fit three linear models:

- lm_1 uses a simple linear combination of features **without interactions**: $f_{lm,1}(\mathbf{x}) = \beta_0 + \sum_{i=1}^{10} \beta_i x_i$, where $\beta = \{-0.6, 0.2, -0.8, 1.6, 0.3, -0.8, 0.5, 0.7, 0.6, -0.3, 1.5\}$.
- lm_2 introduces **interactions between some features of the same group**: $f_{lm,2}(\mathbf{x}) = f_{lm,1}(\mathbf{x}) + \gamma_1 \cdot x_1 \cdot x_2 + \gamma_2 \cdot x_3 \cdot x_4 + \gamma_3 \cdot x_5 \cdot x_6 + \gamma_4 \cdot x_7 \cdot x_8 + \gamma_5 \cdot x_9 \cdot x_{10}$, where $\gamma = \{0.4, -0.6, -2.2, 1.1, 0.0\}$.
- lm_3 introduces **interactions between features in different groups**: $f_{lm,3}(\mathbf{x}) = f_{lm,1}(\mathbf{x}) + \gamma_1 \cdot x_1 \cdot x_5 + \gamma_2 \cdot x_1 \cdot x_7 + \gamma_3 \cdot x_1 \cdot x_9 + \gamma_4 \cdot x_3 \cdot x_5 + \gamma_5 \cdot x_3 \cdot x_7 + \delta_1 \cdot x_3 \cdot x_9 + \delta_2 \cdot x_5 \cdot x_9$, where $\delta = \{0.1, 0.9\}$ and γ are as above.

Since models lm_1 and lm_2 don't have interactions between features of different groups, we can observe strictly what happens when we violate the group independence condition. Note that when we add dependence, for example setting $\rho = 0.3$, this experiment does not distinguish between features of the same group and of different groups. This is extended in Section 3.3.

The results of Experiment 1 are in Figure 1. Notice that for a simple linear model, interactions between features of different groups are *not enough* to increase the MAD. However, adding stronger dependence between all features (but specifically features of different groups – see Section 3.3), increases the MAD. Although this trend is seen for both groupings \mathcal{G}_A and \mathcal{G}_B , the mean, median, and spread is larger for grouping \mathcal{G}_A (with only three groups). The reason may be that more features *per group* adds more variance to the *post-grouped Shapley* values.

3.2 Experiment 2: generalized additive model

In this experiment, we explore how the difference between *groupShapley* and *post-grouped Shapley* varies when using a GAM. This experiment is analogous to Experiment 1 but with a non-linear base model and (pairwise) non-linear interactions. The non-linear interactions takes the form $h(a, b) = a \cdot b + a \cdot b^2 + b \cdot a^2$. The three models are as follows:

- GAM_1 uses a simple linear combination of cosines **without interactions**: $f_{GAM,1}(\mathbf{x}) = \beta_0 + \sum_{i=1}^{10} \cos(x_i)$ and $\beta_0 = -0.6$.
- GAM_2 introduces **interactions between some features of the same group**. The entire model is $f_{GAM,2}(\mathbf{x}) = f_{GAM,1}(\mathbf{x}) + h(x_1, x_2) + h(x_3, x_4) + h(x_5, x_6) + h(x_7, x_8) + h(x_9, x_{10})$.
- GAM_3 introduces **interactions between features in different groups**. The entire model is $f_{GAM,3}(\mathbf{x}) = f_{GAM,1}(\mathbf{x}) + h(x_1, x_5) + h(x_1, x_7) + h(x_1, x_9) + h(x_3, x_5) + h(x_3, x_7) + h(x_3, x_9) + h(x_5, x_9)$.

Notice that these models are similar to those in Experiment 1, except that we use cosines when adding single features to the model and a bivariate function $h(a, b)$ when adding interactions to the model. The results are given in Figure 2. For the models without interactions (GAM_1) and with within-group interactions (GAM_2), the MAD increases with increasing correlation and otherwise behave fairly similar to their linear counterparts in Experiment 1. On the other hand, the model with between-group interactions (GAM_3) behaves completely different compared to its linear counterpart. It has a large MAD already for independent features, which does not seem to increase further with increasing correlation.

3.3 Experiment 3: different dependence within and between feature groups

In this experiment, we fix the dependence between features of the *same* group and vary the dependence between features of *different* groups. The covariance matrix Σ continues to have 1 on the diagonal. However, the correlation for features belonging to the same group is now fixed to 0.87, while the correlations between features in different groups takes values in $\{0, 0.1, 0.3, 0.7, 0.9\}$. We focus on models lm_2 and GAM_2 defined above since we are only interested in violating the group independence condition. The point of this experiment is to investigate the effect of increased between-group correlation, when holding the within-group correlation fixed.

The results for both lm_2 and GAM_2 are given in Figure 3. We notice that even if there is strong dependence between features within the same group, small between-group correlation gives a small MAD. For grouping \mathcal{G}_A , any increase in the between-group correlation increases the MAD. For grouping \mathcal{G}_B , this is not the case and only a strong between-group correlation results in an increase in MAD. Like Experiment 1, the MAD of grouping \mathcal{G}_A seems to have a larger mean, median, and spread than grouping \mathcal{G}_B across all correlations.

The simulation study showed that dependence between features of *different* groups and non-additive models often yield different *groupShapley* and *post-grouped Shapley* values. When using a simple linear model, the MAD increased with the dependence between

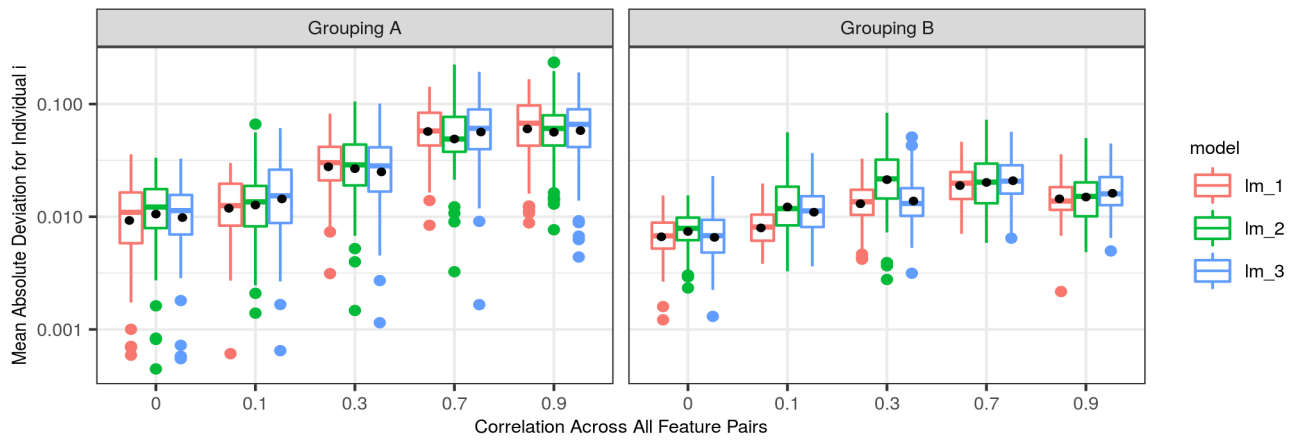


Figure 1. Boxplots of the mean absolute difference on log scale for all three lm models fit in Experiment 1. Black dot indicates the mean.

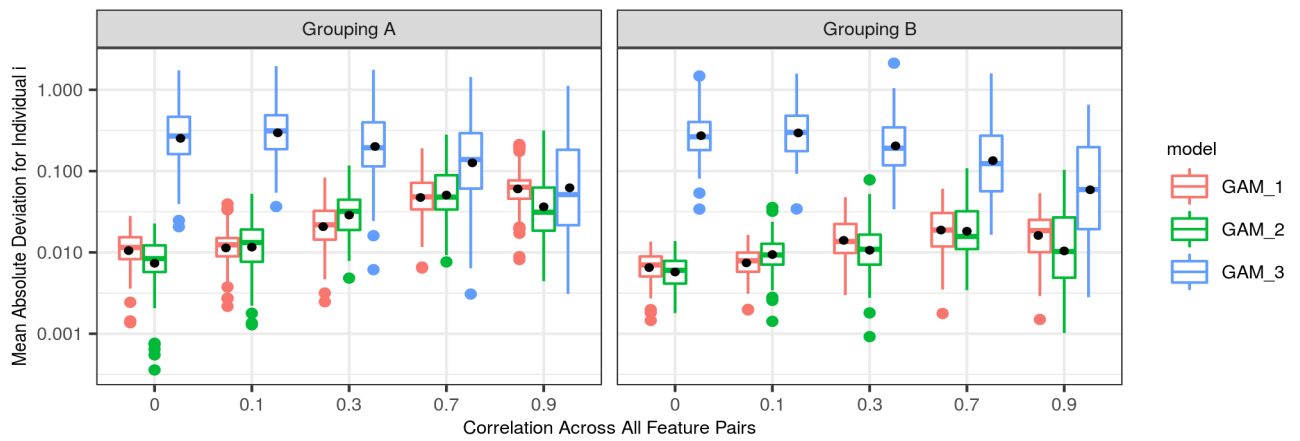


Figure 2. Boxplots of the mean absolute difference on log scale for all three GAMs fit in Experiment 2. Black dot indicates the mean.

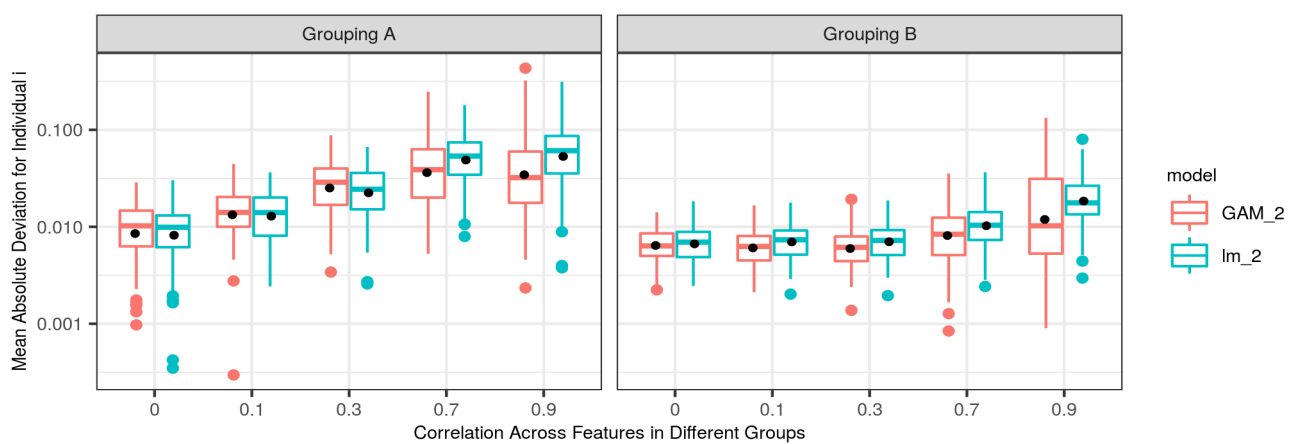


Figure 3. Boxplots of the mean absolute difference on log scale for all lm_2 and GAM_2 fit in Experiment 3. Black dot indicates the mean. The correlation between features of the same group is fixed at 0.87.

the features and did not seem to be affected by the linear interactions. For the GAMs, the MAD was large for the model with non-linear interactions between features of *different* groups, but did not increase further with increased dependence. Finally, we saw that using a fixed and non-zero between-group correlation did not significantly change the trends seen in Experiments 1 and 2.

4 Real data example

In this section, we demonstrate how *groupShapley* can be applied to a real data set. We use a car insurance data set found on the Kaggle website (<https://www.kaggle.com/xiaomengsun/car-insurance-claim-data>). The data contains two different response variables, 23 features, and 10,302 observations. The response we use is the binary variable *customer had a claim*. The features can be naturally partitioned into the following groups:

- *Personal Information*: age of driver, highest education level, number of children living at home, value of home, income, job category, number of driving children, marital status, single parent, gender, distance to work, whether driver lives in an urban city, how many years driver has had job.
- *Track Record*: number of claims in the past five years, motor vehicle record points, licence revoked in past seven years, amount of time as customer.
- *Car Information*: value of car, age of car, type of car, whether car is red.

Five of the variables have missing data so we use predictive mean matching to impute these. To model the probability of a claim, we train a random forest model with 500 trees using the *ranger* R-package (Wright and Ziegler, 2017) on the binary response and all 23 features. We use 10 fold cross-validation and get an average out-of-fold AUC of 0.815. The average predicted probability of a claim is 0.273. We then use *groupShapley* to calculate the group Shapley values of the *Personal Information*, *Car Information*, and *Track Record* groups for four different individuals. Since there is a mix of continuous, discrete and categorical features, the conditional inference tree approach of Redelmeier et al. (2020) is used to estimate (4). We plot the three grouped Shapley values for four different individuals in Figure 4.

The first individual is a single mother of four (where two children drive). She drives an SUV and drives 27 miles to work. She has had one claim in the last five years and has three motor vehicle record points. *Personal Information* gives the largest increase in the predicted probability, which is not surprising given her travel distance and two young drivers. The second individual is a 37-year-old father of two (where one child drives). He has had one claim in the last five years, his licence revoked in the last seven years, and ten motor vehicle points. His *Track Record* significantly increases his predicted probability, which is natural given his misdemeanors.

The third individual is a 60-year-old married male with no kids at home. He drives a red sports car and has had three claims in the last five years. He has a PhD and currently

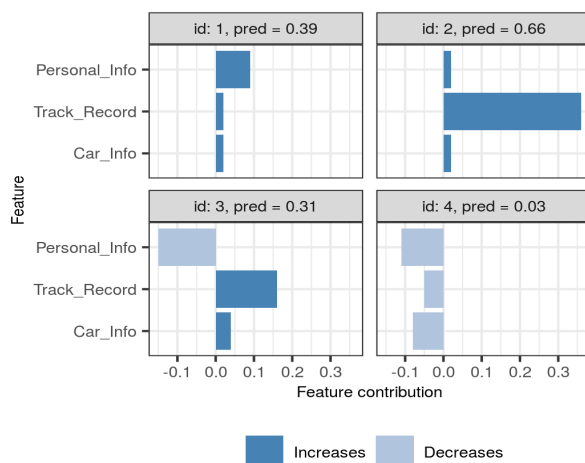


Figure 4. Estimated *groupShapley* values for four individuals with the groups defined in Section 4.

works as a doctor. His *Personal Information* naturally reduces his predicted probability, while his poor *Track Record*, and to some extent his luxurious *Car Information* increases his predicted probability. The fourth individual is a 50-year-old female with no kids at home. She drives a minivan and has no previous claims or revoked licences. She also has a PhD and drives 42 miles to work. She appears to be on the safer side of things, which is reflected in all negative *groupShapley* values and a smaller predicted probability.

This real world example and Figure 4 exemplifies that *groupShapley* is a simple and efficient way to produce intuitive prediction explanations.

5 Concluding remarks

In addition to explaining predictions based on feature groups, *groupShapley* can be utilized to compute the relevant feature-wise scores in the case where there are categorical features which are one-hot-encoded when passed to the model $f(\cdot)$. In this case, one is not really interested in the Shapley values for the individual one-hot-encoded features, but for the original features. One can then group the one-hot-encoded features into one group and apply *groupShapley*. This might be viewed as an alternative to the approach in Redelmeier et al. (2020) for computing feature-wise Shapley values in the presence of categorical features.

Another possible application of *groupShapley* is to time series classification settings (Kvamme et al., 2018; Ordóñez and Roggen, 2016). Suppose the model takes one or more time series as input and predicts class probabilities. It may not be interesting or even relevant to compute Shapley values per time point, but by splitting the time series into different chunks, *groupShapley* can be used to explain how different parts of the time series contribute.

It is worth noting that *groupShapley* is a highly general, and computationally efficient solution that can be used in a range of other Shapley value settings. In addition to the

prediction explanation setting, it can be reused to bypass the computational burden for *global* explanation methods (Grömping, 2007; Owen and Prieur, 2017), or even outside the fields of explainable artificial intelligence and interpretable machine learning (Moretti and Patrone, 2008).

Finally, the groupShapley methodology is implemented in the shapr R-package (Sellereite and Jullum, 2020)².

Acknowledgements

We thank Anders Løland for useful discussions in the early-stages with this paper and Jens Christian Wahl for discussions related to the simulation study. This work was supported by the Norwegian Research Council grant 237718 (Big Insight).

References

- Aas, K., Jullum, M., and Løland, A. (2019). Explaining individual predictions when features are dependent: More accurate approximations to shapley values. *arXiv preprint arXiv:1903.10464*.
- Aas, K., Nagler, T., Jullum, M., and Løland, A. (2021). Explaining predictive models using shapley values and non-parametric vine copulas. *arXiv preprint arXiv:2102.06416*.
- Algaba, E., Fragnelli, V., and Sánchez-Soriano, J. (2019). *Handbook of the Shapley value*. CRC Press.
- Chen, H., Janizek, J. D., Lundberg, S., and Lee, S.-I. (2020). True to the model or true to the data? *arXiv preprint arXiv:2006.16234*.
- Chen, J., Song, L., Wainwright, M. J., and Jordan, M. I. (2019). L-shapley and c-shapley: Efficient model interpretation for structured data. In *International Conference on Learning Representations*.
- Covert, I. and Lee, S.-I. (2020). Improving kernelshap: Practical shapley value estimation via linear regression. *arXiv preprint arXiv:2012.01536*.
- Frye, C., de Mijolla, D., Cowton, L., Stanley, M., and Feige, I. (2020). Shapley-based explainability on the data manifold. *arXiv preprint arXiv:2006.01272*.
- Grömping, U. (2007). Estimators of relative importance in linear regression based on variance decomposition. *The American Statistician*, 61(2):139–147.

2. As of February 2021, the groupShapley methodology is only available in the development version of the package on GitHub: <https://github.com/NorskRegnesentral/shapr>.

- Janzing, D., Minorics, L., and Blöbaum, P. (2020). Feature relevance quantification in explainable ai: A causal problem. In *International Conference on Artificial Intelligence and Statistics*, pages 2907–2916.
- Kaufman, L. and Rousseeuw, P. J. (2009). *Finding groups in data: an introduction to cluster analysis*, volume 344. John Wiley & Sons.
- Kvamme, H., Sellereite, N., Aas, K., and Sjursen, S. (2018). Predicting mortgage default using convolutional neural networks. *Expert Systems with Applications*, 102:207–217.
- Li, X., Dvornek, N. C., Zhou, Y., Zhuang, J., Ventola, P., and Duncan, J. S. (2019). Efficient interpretation of deep learning models using graph structure and cooperative game theory: Application to asd biomarker discovery. In *International Conference on Information Processing in Medical Imaging*, pages 718–730. Springer.
- Lundberg, S. M., Erion, G., Chen, H., DeGrave, A., Prutkin, J. M., Nair, B., Katz, R., Himmelfarb, J., Bansal, N., and Lee, S.-I. (2020). From local explanations to global understanding with explainable ai for trees. *Nature Machine Intelligence*, 2(1):2522–5839.
- Lundberg, S. M. and Lee, S.-I. (2017). A Unified Approach to Interpreting Model Predictions. In *Advances in Neural Information Processing Systems*, pages 4768–4777. Curram Associates Inc.
- Marichal, J.-L., Kojadinovic, I., and Fujimoto, K. (2007). Axiomatic characterizations of generalized values. *Discrete Applied Mathematics*, 155(1):26–43.
- Molnar, C. (2020). *Interpretable machine learning: A guide for making black box models explainable*. Christoph Molnar, Leanpub.
- Moretti, S. and Patrone, F. (2008). Transversality of the shapley value. *Top*, 16(1):1–41.
- Ordóñez, F. J. and Roggen, D. (2016). Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition. *Sensors*, 16(1):115.
- Owen, A. B. and Priour, C. (2017). On Shapley value for measuring importance of dependent inputs. *SIAM/ASA Journal on Uncertainty Quantification*, 5:986–1002.
- Redelmeier, A., Jullum, M., and Aas, K. (2020). Explaining predictive models with mixed features using shapley values and conditional inference trees. In *International Cross-Domain Conference for Machine Learning and Knowledge Extraction*, pages 117–137. Springer.
- Sellereite, N. and Jullum, M. (2020). shapr: An r-package for explaining machine learning models with dependence-aware shapley values. *Journal of Open Source Software*, 5(46):2027.
- Shapley, L. S. (1953). A Value for N-Person Games. *Contributions to the Theory of Games*, 2:307–317.

Sundararajan, M. and Najmi, A. (2020). The many shapley values for model explanation. In *International Conference on Machine Learning*, pages 9269–9278. PMLR.

Wright, M. N. and Ziegler, A. (2017). ranger: A fast implementation of random forests for high dimensional data in C++ and R. *Journal of Statistical Software*, 77(1):1–17.

A Proofs

This appendix contains proofs of Theorem 2.1 and 2.2. Before heading to those proofs, we will present and prove two lemmas that will be useful in those proofs.

Lemma A.1 (Contribution function identity). *Assume that Conditions 2.1 and 2.2 holds. Further, let \mathcal{T}_0 and \mathcal{T}_{AB} be any two disjoint subsets of the partition $\mathcal{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_G\}$ and $\mathcal{S}_0, \mathcal{S}_A, \mathcal{S}_B$ be any three feature-subsets where $\mathcal{S}_0 \subseteq \mathcal{T}_0$, $\mathcal{S}_A \subseteq \mathcal{T}_{AB}$, $\mathcal{S}_B \subseteq \mathcal{T}_{AB}$. Then, we have that*

$$v(\mathcal{S}_0 \cup \mathcal{S}_A) - v(\mathcal{S}_0 \cup \mathcal{S}_B) = v(\mathcal{S}_A) - v(\mathcal{S}_B). \quad (\text{A.1})$$

Proof. Let us write $\mathbf{x}_{\mathcal{G}_i}$ for the subvector of \mathbf{x} corresponding to the features in group \mathcal{G}_i . Let $\mathcal{G}_{\mathcal{S}}$ be the set of group indexes with at least one feature in the subset \mathcal{S} , i.e. $\mathcal{G}_{\mathcal{S}} = \{i : \mathcal{G}_i \cap \mathcal{S} \neq \emptyset\}$, and let $\bar{\mathcal{G}}_{\mathcal{S}}$ be the group indexes not part of $\mathcal{G}_{\mathcal{S}}$. To show the equality in (A.1), we will first split $v(\mathcal{S})$ into simpler terms by utilizing the additive structure $f(\mathbf{x}) = \sum_{i=1}^G f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})$ and the group independence $\mathbf{x}_{\mathcal{G}_i} \perp\!\!\!\perp \mathbf{x}_{\mathcal{G}_j}$ for all $i \neq j$.

We may then utilize Conditions 2.1 and 2.2 to decompose $v(\mathcal{S})$ into a sum of conditional expectations involving the groups associated with \mathcal{S} and unconditional expectations for the remaining groups:

$$\begin{aligned} v(\mathcal{S}) &= \mathbb{E}[f(\mathbf{x})|\mathbf{x}_{\mathcal{S}}] = \mathbb{E}\left[\sum_{i=1}^G f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i}) \middle| \mathbf{x}_{\mathcal{S}}\right] = \sum_{i=1}^G \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{S}}] \\ &= \sum_{i \in \mathcal{G}_{\mathcal{S}}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{S}}] + \sum_{i \in \bar{\mathcal{G}}_{\mathcal{S}}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{S}}] \\ &= \sum_{i \in \mathcal{G}_{\mathcal{S}}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{G}_i \cap \mathcal{S}}] + \sum_{i \in \bar{\mathcal{G}}_{\mathcal{S}}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})]. \end{aligned} \quad (\text{A.2})$$

Now, proving the equality in (A.1) is just about utilizing the separability of the formula in (A.2) and writing out the four terms:

$$\begin{aligned} v(\mathcal{S}_0 \cup \mathcal{S}_A) &= \sum_{i \in \mathcal{G}_{\mathcal{S}_0}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{G}_i \cap \mathcal{S}_0}] + \sum_{i \in \mathcal{G}_{\mathcal{S}_A}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{G}_i \cap \mathcal{S}_A}] + \sum_{i \in \bar{\mathcal{G}}_{\mathcal{S}_0 \cup \mathcal{S}_A}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})], \\ v(\mathcal{S}_0 \cup \mathcal{S}_B) &= \sum_{i \in \mathcal{G}_{\mathcal{S}_0}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{G}_i \cap \mathcal{S}_0}] + \sum_{i \in \mathcal{G}_{\mathcal{S}_B}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{G}_i \cap \mathcal{S}_B}] + \sum_{i \in \bar{\mathcal{G}}_{\mathcal{S}_0 \cup \mathcal{S}_B}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})], \\ v(\mathcal{S}_A) &= \sum_{i \in \mathcal{G}_{\mathcal{S}_A}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{G}_i \cap \mathcal{S}_A}] + \sum_{i \in \bar{\mathcal{G}}_{\mathcal{S}_0 \cup \mathcal{S}_A}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})], \\ v(\mathcal{S}_B) &= \sum_{i \in \mathcal{G}_{\mathcal{S}_B}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{G}_i \cap \mathcal{S}_B}] + \sum_{i \in \bar{\mathcal{G}}_{\mathcal{S}_0 \cup \mathcal{S}_B}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})]. \end{aligned}$$

From the above expressions we easily see that both sides of (A.1) equal

$$\begin{aligned} &\sum_{i \in \mathcal{G}_{\mathcal{S}_A}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{G}_i \cap \mathcal{S}_A}] - \sum_{i \in \mathcal{G}_{\mathcal{S}_B}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{G}_i \cap \mathcal{S}_B}] \\ &+ \sum_{i \in \bar{\mathcal{G}}_{\mathcal{S}_0 \cup \mathcal{S}_A}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})] - \sum_{i \in \bar{\mathcal{G}}_{\mathcal{S}_0 \cup \mathcal{S}_B}} \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})], \end{aligned}$$

and the proof is complete. ■

Lemma A.2 (Simplified contribution function identity). *Assume that Conditions 2.1 and 2.2 holds and that $j \in \mathcal{G}_i$. Further, let \mathcal{S} be a feature subset of \mathcal{G}_i not containing j . Define also $v_{\mathcal{G}_i}(\mathcal{S}) = \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{S}}]$. Then we have that*

$$v(\mathcal{S} + \{j\}) - v(\mathcal{S}) = v_{\mathcal{G}_i}(\mathcal{S} + \{j\}) - v_{\mathcal{G}_i}(\mathcal{S}).$$

Proof. The result follows almost directly from writing out the expressions for v in terms of the $f_{\mathcal{G}_k}$ -functions. Since both \mathcal{S} and j only contain features in \mathcal{G}_i , we may, similarly to (A.2), write

$$\begin{aligned} v(\mathcal{S}) &= \sum_{k=1}^G \mathbb{E}[f_{\mathcal{G}_k}(\mathbf{x}_{\mathcal{G}_k})|\mathbf{x}_{\mathcal{S}}] \\ &= \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{S}}] + \sum_{k \neq i} \mathbb{E}[f_{\mathcal{G}_k}(\mathbf{x}_{\mathcal{G}_k})], \\ v(\mathcal{S} + \{j\}) &= \sum_{k=1}^G \mathbb{E}[f_{\mathcal{G}_k}(\mathbf{x}_{\mathcal{G}_k})|\mathbf{x}_{\mathcal{S} + \{j\}}] \\ &= \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{S} + \{j\}}] + \sum_{k \neq i} \mathbb{E}[f_{\mathcal{G}_k}(\mathbf{x}_{\mathcal{G}_k})]. \end{aligned}$$

As the sums in the equations cancel each other out, we get

$$\begin{aligned} v(\mathcal{S} + \{j\}) - v(\mathcal{S}) &= \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{S} + \{j\}}] - \mathbb{E}[f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})|\mathbf{x}_{\mathcal{S}}] \\ &= v_{\mathcal{G}_i}(\mathcal{S} + \{j\}) - v_{\mathcal{G}_i}(\mathcal{S}), \end{aligned}$$

which completes the proof. ■

A.1 Proof of Theorem 2.1

We are now going to prove that the two conditions imply that the feature-wise Shapley formula takes the following simplifying form

$$\phi_j = \sum_{\mathcal{S} \subseteq \mathcal{G}_i \setminus \{j\}} \frac{|\mathcal{S}|!(|\mathcal{G}_i| - |\mathcal{S}| - 1)!}{|\mathcal{G}_i|!} (v_{\mathcal{G}_i}(\mathcal{S} + \{j\}) - v_{\mathcal{G}_i}(\mathcal{S})).$$

To do that, we will apply Lemma A.1, perform a series of combinatorial operations, before finally applying Lemma A.2. Recall that $j \in \mathcal{G}_i$. Using \mathcal{R} in place of the standard feature-wise subset notation \mathcal{S} , and decompose it through the partition \mathcal{R} , i.e. $\mathcal{S} \cup \mathcal{S}_0 = \mathcal{R}$. We may then write the feature-wise Shapley value formula as

$$\begin{aligned} \phi_j &= \sum_{\mathcal{R} \subseteq \mathcal{M} \setminus \{j\}} \frac{|\mathcal{R}|!(M - |\mathcal{R}| - 1)!}{M!} (v(\mathcal{R} \cup \{j\}) - v(\mathcal{R})) \\ &= \sum_{\mathcal{S} \subseteq \mathcal{G}_i \setminus \{j\}} \sum_{\mathcal{S}_0 \subseteq \mathcal{M} \setminus \mathcal{G}_i} \frac{(|\mathcal{S}| + |\mathcal{S}_0|)!(M - |\mathcal{S}| + |\mathcal{S}_0| - 1)!}{M!} (v(\mathcal{S}_0 \cup \mathcal{S} \cup \{j\}) - v(\mathcal{S}_0 \cup \mathcal{S})). \end{aligned} \tag{A.3}$$

Now, letting $\mathcal{S}_0 = \mathcal{S}_0$, $\mathcal{T}_0 = \mathcal{M} \setminus \mathcal{G}_i$ and $\mathcal{T}_{AB} = \mathcal{G}_i$, $\mathcal{S}_A = \mathcal{S} \cup \{j\}$, $\mathcal{S}_B = \mathcal{S}$ in Lemma A.1, we see that

$$v(\mathcal{S}_0 \cup \mathcal{S} \cup \{j\}) - v(\mathcal{S}_0 \cup \mathcal{S}) = v(\mathcal{S} \cup \{j\}) - v(\mathcal{S}).$$

As a consequence, we may write (A.3) as

$$\phi_j = \sum_{\mathcal{S} \subseteq \mathcal{G}_i \setminus \{j\}} (v(\mathcal{S} \cup \{j\}) - v(\mathcal{S})) \sum_{\mathcal{S}_0 \subseteq \mathcal{M} \setminus \mathcal{G}_i} \frac{(|\mathcal{S}| + |\mathcal{S}_0|)!(M - |\mathcal{S}| - |\mathcal{S}_0| - 1)!}{M!}. \quad (\text{A.4})$$

Now, let us focus on the inner sum of (A.4). We will first rewrite the summand using the binomial coefficient and then rewrite this as a sum over the size of \mathcal{S}_0 using combinatorics:

$$\sum_{\mathcal{S}_0 \subseteq \mathcal{M} \setminus \mathcal{G}_i} \frac{(|\mathcal{S}| + |\mathcal{S}_0|)!(M - |\mathcal{S}| - |\mathcal{S}_0| - 1)!}{M!} = \frac{1}{M} \sum_{\mathcal{S}_0 \subseteq \mathcal{M} \setminus \mathcal{G}_i} \frac{1}{\binom{M-1}{|\mathcal{S}| + |\mathcal{S}_0|}} = \frac{1}{M} \sum_{k=0}^{M-|\mathcal{G}_i|} \frac{\binom{M-|\mathcal{G}_i|}{k}}{\binom{M-1}{|\mathcal{S}|+k}}. \quad (\text{A.5})$$

To further simplify the right hand side of (A.5), we repeatedly utilize the identity

$$\int_0^1 z^a (1-z)^b dz = \frac{a! b!}{(a+b+1)!},$$

for positive numbers a and b . This is a well known property of the Beta function, and the integration constant of the Beta distribution. Let us consider the denominator of the summand in (A.5). Writing $a = |\mathcal{S}| + k$ and $b = M - 1 - |\mathcal{S}| - k$, we have that

$$\frac{1}{\binom{M-1}{|\mathcal{S}|+k}} = \frac{1}{\binom{b+a}{a}} = \frac{a! b!}{(b+a)!} = (a+b+1) \int_0^1 z^a (1-z)^b dz = M \int_0^1 z^{|\mathcal{S}|+k} (1-z)^{M-1-|\mathcal{S}|-k} dz.$$

Also utilizing that $\sum_{k=0}^q \binom{q}{k} r^k = (1+r)^q$, we then have that

$$\begin{aligned} \frac{1}{M} \sum_{k=0}^{M-|\mathcal{G}_i|} \frac{\binom{M-|\mathcal{G}_i|}{k}}{\binom{M-1}{|\mathcal{S}|+k}} &= \sum_{k=0}^{M-|\mathcal{G}_i|} \binom{M-|\mathcal{G}_i|}{k} \int_0^1 z^{|\mathcal{S}|+k} (1-z)^{M-1-|\mathcal{S}|-k} dz \\ &= \int_0^1 z^{|\mathcal{S}|} (1-z)^{M-1-|\mathcal{S}|} \sum_{k=0}^{M-|\mathcal{G}_i|} \binom{M-|\mathcal{G}_i|}{k} \left(\frac{z}{1-z}\right)^k dz \\ &= \int_0^1 z^{|\mathcal{S}|} (1-z)^{M-1-|\mathcal{S}|} \left(\frac{1}{1-z}\right)^{M-|\mathcal{G}_i|} dz \\ &= \int_0^1 z^{|\mathcal{S}|} (1-z)^{|\mathcal{G}_i|-|\mathcal{S}|-1} dz \\ &= \frac{|\mathcal{S}|!(|\mathcal{G}_i|-|\mathcal{S}|-1)!}{|\mathcal{G}_i|}. \end{aligned}$$

Inserting this simplified formula into (A.4) gives

$$\phi_j = \sum_{\mathcal{S} \subseteq \mathcal{G}_i \setminus \{j\}} \frac{|\mathcal{S}|!(|\mathcal{G}_i|-|\mathcal{S}|-1)!}{|\mathcal{G}_i|!} (v(\mathcal{S} + \{j\}) - v(\mathcal{S})).$$

Since this a sum over subsets only contained in \mathcal{G}_i , application of Lemma A.2 allows us to replace $v(\mathcal{S} + \{j\}) - v(\mathcal{S})$ by $v_{\mathcal{G}_i}(\mathcal{S} + \{j\}) - v_{\mathcal{G}_i}(\mathcal{S})$, and thereby completes the proof. ■

A.2 Proof of Theorem 2.2

We are now going to prove that under Conditions 2.1 and 2.2, *groupShapley* is equivalent to *post-grouped Shapley* and they take a simplified form, i.e.

$$\sum_{j \in \mathcal{G}_i} \phi_j = \phi_{\mathcal{G}_i} = v_{\mathcal{G}_i}(\mathcal{G}_i) - v_{\mathcal{G}_i}(\emptyset), \quad i = 1, \dots, G.$$

We will start by showing that $\sum_{j \in \mathcal{G}_i} \phi_j = v_{\mathcal{G}_i}(\mathcal{G}_i) - v_{\mathcal{G}_i}(\emptyset)$. From Theorem 2.1, we have that the collection $\{\phi_j : j \in \mathcal{G}_i\}$ corresponds to the Shapley values of $f_{\mathcal{G}_i}(\mathbf{x}_{\mathcal{G}_i})$. Thus, by the efficiency axiom, we have that $\sum_{j \in \mathcal{G}_i} \phi_j = v_{\mathcal{G}_i}(\mathcal{G}_i) - v_{\mathcal{G}_i}(\emptyset)$. This completes the first part of the proof.

What remains is to show that under the state conditions, we also have $\phi_{\mathcal{G}_i} = v_{\mathcal{G}_i}(\mathcal{G}_i) - v_{\mathcal{G}_i}(\emptyset)$, for which we will utilize Lemma A.1 once again. Letting $\mathcal{S}_0 = \mathcal{T}$, $\mathcal{T}_0 = \mathcal{G} \setminus \mathcal{G}_i$ and $\mathcal{S}_A = \mathcal{T}_{AB} = \mathcal{G}_i$, $\mathcal{S}_B = \emptyset$, we see that

$$v(\mathcal{T} \cup \mathcal{G}_i) - v(\mathcal{T}) = v(\mathcal{G}_i) - v(\emptyset). \quad (\text{A.6})$$

Furthermore, by writing out the two terms on the right hand side of (A.6), as in the proof of Lemma A.2, we have that $v(\mathcal{G}_i) = v_{\mathcal{G}_i}(\mathcal{G}_i) + \sum_{k \neq i} \mathbb{E}[f_{\mathcal{G}_k}(\mathbf{x}_{\mathcal{G}_k})]$ and $v(\emptyset) = v_{\mathcal{G}_i}(\emptyset) + \sum_{k \neq i} \mathbb{E}[f_{\mathcal{G}_k}(\mathbf{x}_{\mathcal{G}_k})]$. Thus, we get that $v(\mathcal{G}_i) - v(\emptyset) = v_{\mathcal{G}_i}(\mathcal{G}_i) - v_{\mathcal{G}_i}(\emptyset)$. Finally, since the Shapley value weights sum to 1, we have that

$$\begin{aligned} \phi_{\mathcal{G}_i} &= \sum_{\mathcal{T} \subseteq \mathcal{G} \setminus \mathcal{G}_i} \frac{|\mathcal{T}|!(|\mathcal{G}_i| - |\mathcal{T}| - 1)!}{|\mathcal{G}_i|!} (v(\mathcal{T} \cup \mathcal{G}_i) - v(\mathcal{T})) \\ &= (v_{\mathcal{G}_i}(\mathcal{G}_i) - v_{\mathcal{G}_i}(\emptyset)) \sum_{\mathcal{T} \subseteq \mathcal{G} \setminus \mathcal{G}_i} \frac{|\mathcal{T}|!(|\mathcal{G}_i| - |\mathcal{T}| - 1)!}{|\mathcal{G}_i|!} \\ &= v_{\mathcal{G}_i}(\mathcal{G}_i) - v_{\mathcal{G}_i}(\emptyset), \end{aligned}$$

which completes the proof. ■