



# Hierarchical clustering: Visualization, feature importance and model selection

Luben M.C. Cabezas<sup>a,\*</sup>, Rafael Izbicki<sup>a</sup>, Rafael B. Stern<sup>b</sup>

<sup>a</sup> Department of Statistics, Federal University of São Carlos, São Carlos, 13565-905, São Paulo, Brazil

<sup>b</sup> Institute of Mathematics and Statistics, University of São Paulo, São Paulo, 05508-090, São Paulo, Brazil

## ARTICLE INFO

### Article history:

Received 28 December 2021

Received in revised form 27 January 2023

Accepted 4 April 2023

Available online 20 April 2023

### MSC:

62H30

1111

### Keywords:

Hierarchical clustering

Unsupervised learning

Phylogenetic models

## ABSTRACT

We propose methods for the analysis of hierarchical clustering that fully use the multi-resolution structure provided by a dendrogram. Specifically, we propose a loss for choosing between clustering methods, a feature importance score and a graphical tool for visualizing the segmentation of features in a dendrogram. Current approaches to these tasks lead to loss of information since they require the user to generate a single partition of the instances by cutting the dendrogram at a specified level. Our proposed methods, instead, use the full structure of the dendrogram. The key insight behind the proposed methods is to view a dendrogram as a phylogeny. This analogy permits the assignment of a feature value to each internal node of a tree through an evolutionary model. Real and simulated datasets provide evidence that our proposed framework has desirable outcomes and gives more insights than state-of-art approaches. We provide an R package that implements our methods.

© 2023 Elsevier B.V. All rights reserved.

## 1. Introduction

Clustering methods have the goal of grouping similar sample points. They are used in applications such as pattern recognition [1], image analysis [2], bioinformatics [3], and information retrieval [4].

Clustering techniques can be divided into two categories: partition-based and hierarchical. An example of partition-based clustering is K-means [5], which creates  $K$  clusters based on a Voronoi partition of the feature space. Similarly, mode-based clustering [6] creates a partition by assigning each observation to a mode of a density estimate. On the other hand, hierarchical clustering yields a tree-based representation of the objects (dendrogram), such as in agglomerative clustering [7], DIANA [8] and partial least squares methods [9]. This paper focuses solely on hierarchical clustering.

Hierarchical clustering has some advantages when compared to partition-based clustering. For instance, it does not require the specification of the number of clusters. Also, as shown in Fig. 1, a dendrogram displays a general similarity structure of the data by providing a multiresolution view. The lower in the dendrogram that two observations are merged, the more similar they are. By using this similarity structure, Fig. 1 illustrates how to obtain a partition by cutting a dendrogram at a given height.

However, there exist challenging questions when implementing hierarchical clustering. For instance, how does one choose a

particular hierarchical clustering method among the many that are available? (question 1) Specifically, how can the performance of a dendrogram be evaluated? Also, after a clustering method is chosen, how do the available features relate to the dendrogram? For example, how can the position of each sample in the dendrogram be explained in terms of its feature values? (question 2) Also, which features are best segmented by the dendrogram? (question 3)

**Novelty.** To the best of our knowledge, our framework is the first one to address the above questions by utilizing the full hierarchical structure provided by the dendrogram. Unlike current approaches, which require the user to generate a partition by cutting the dendrogram at a specified level [3,10–17], our approach makes use of the richer structure provided by the dendrogram and does not restrict itself to a single partition, thus avoiding loss of information, as demonstrated in Section 5.


### 1.1. Overview of the method

We answer the questions in the last section through probabilistic evolutionary models commonly used in phylogenetic analysis [18–20]. The key insight in this framework is to view a dendrogram as a phylogeny. This metaphor is helpful since the phylogeny is associated to a probabilistic evolutionary model.

Probabilistic evolutionary models describe how each characteristic evolves over a phylogeny. In the context of hierarchical unsupervised learning, this translates into modeling how a feature evolves over a dendrogram using phylogeny information and

\* Corresponding author.

E-mail address: [lucruz45.cab@gmail.com](mailto:lucruz45.cab@gmail.com) (L.M.C. Cabezas).



**Fig. 1.** (Left) Dendrogram obtained from hierarchical clustering with average linkage. (Right) By cutting the dendrogram at height 4, two clusters are obtained.

feature data. These models predict feature values for leaves of the dendrogram and also for internal nodes, a task known as ancestral state reconstruction. This reconstruction allow one to visualize each feature's segmentation by the dendrogram, thus providing an answer to (2). Also the evolutionary model's predictions provide a measure for the accuracy of each dendrogram. By adapting ideas from cross-validation, we obtain a loss used for dendrogram selection, thus answering question (1). Finally, feature importance is associated to how well the probabilistic evolutionary model predicts that feature, which answers question (3).

We illustrate our approach using the **ceramic samples** dataset,<sup>1</sup> which contains information about the chemical composition of 88 ceramic samples. This dataset contains 17 features, each related to the percentage of certain chemical compound present in each sample (e.g.,  $\text{Al}_2\text{O}_3$  (Aluminium Oxide),  $\text{CaO}$  (Calcium Oxide), etc.). The table in Fig. 2(a) provides the hierarchical clustering prediction loss for a few clustering methods. This loss provides an immediate way of comparing the performance of the methods. For instance, since McQuitty has the lowest loss, it has the highest performance in this dataset, closely followed by average linkage. Also, Fig. 2(b) illustrates the feature importance score. While  $\text{Al}_2\text{O}_3$ ,  $\text{SrO}$  and  $\text{CaO}$  are the most relevant features in the construction of the dendrogram,  $\text{MgO}$ ,  $\text{CuO}$  and  $\text{ZnO}$  are the least important ones. Finally, Fig. 2(c) illustrates how the hierarchical structure explains the distribution behavior of each feature. The left dendrogram shows that  $\text{Al}_2\text{O}_3$  is well segmented by the dendrogram, as explained by the even color distribution among its major branches. That is, while 2 major clusters are characterized by low values for  $\text{Al}_2\text{O}_3$ , a third major cluster is characterized by high values of  $\text{Al}_2\text{O}_3$ . The right dendrogram shows that  $\text{MgO}$  is not so well explained by the dendrogram, since its major branches contain a high variability of colors. None of the illustrations required us to fix the number of clusters.

This work is organized as follows: Section 2 reviews hierarchical clustering and probabilistic evolutionary models, Section 3 formulates and explains our methodology, Section 4 outlines related approaches, and Section 5 shows experimental results.

## 2. The relation between hierarchical clustering and probabilistic evolutionary models

Hierarchical clustering and phylogenetics share a common trait: dendrograms and phylogenies are both formalized as trees. That is, the difference between a dendrogram and a phylogeny is given solely by the context in which it is used. While a dendrogram is used to specify the similarity structure between instances, the phylogeny describes evolutionary relationships. Since both types of analysis share a common formalism, it is possible to adapt the methods of one to the other.

Based on this shared formalism, we also use a common notation for describing the elements of a tree. The left panel in Fig. 1 depicts a tree. The nodes at the bottom of the tree are called leaves. These leaves often represent instances in the dataset. As one moves up the tree, the leaves merge into internal nodes, and then internal nodes merge with one another until the root is reached at the top of the tree. The descendants of a node are the leaves that stem from it.

The following subsections briefly review hierarchical clustering and phylogenetical evolutionary models.

### 2.1. Hierarchical methods

Hierarchical clustering methods yield a tree-based representation of the data named dendrogram, as illustrated in Fig. 1. In clustering, a node is often interpreted as the group of instances which are its descendants. The earlier a merge between two nodes, the more similar are the corresponding groups of descendants [21].

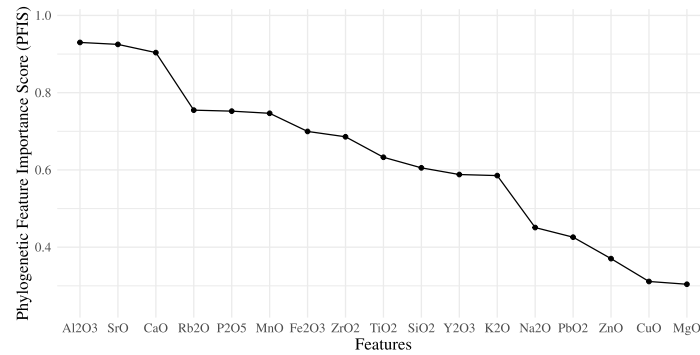
Using the dendrogram, it is possible to partition the instances according to their similarity. Such a partition is obtained by cutting the dendrogram with a horizontal line at a given height, as shown on the right panel of Fig. 1. While cutting at lower heights obtains groups with a greater similarity, cutting at higher heights obtains a smaller number of groups. A trade-off between these properties yields the most interpretable partition.

Hierarchical clustering methods can be divided into two paradigms: agglomerative (bottom-up) and divisive (top-down) [22]. Agglomerative strategies start at the leaves of the dendrogram, iteratively merging selected pairs of branches until the root of the tree is reached. The pair of branches chosen for merging


<sup>1</sup> <https://archive.ics.uci.edu/ml/datasets/Chemical+Composition+of+Ceramic+Samples>

Method	McQuitty	Average linkage	Diana	Ward D	Single linkage
Loss	0.545	0.546	0.561	0.578	0.651

(a) Cross-validated loss (CVL) indicates which hierarchical clustering method to use. In this case, McQuitty has the best performance.



(b) Phylogenetic feature importance score (PFIS) shows the importance of each feature in Mcquitty dendrogram.



(c) Evolutionary dendrograms gives insights about the clustering by describing how each feature behaves in each internal node.

**Fig. 2.** Illustration of the proposed hierarchical clustering methodology to the ceramic dataset. In panel (a), the prediction loss measures how well each hierarchical clustering method explains the data. In this example, McQuitty is the best method. In panel (b), the feature importance score determines which of the features are best explained by the McQuitty dendrogram. While  $\text{Al}_2\text{O}_3$  is the best explained by the dendrogram,  $\text{MgO}$  is the least explained. In panel (c), the evolutionary dendrograms for  $\text{Al}_2\text{O}_3$  (left) and  $\text{MgO}$  (right) provide insights on how the hierarchical structure explains the behavior of each of these feature.

is the one that has the smallest measurement of intergroup dissimilarity. Divisive methods start at the root of the tree. Such methods iteratively divide a chosen branch into smaller ones until the leaves are obtained. The splitting criterion maximizes a measure of between-group dissimilarity.

## 2.2. Evolutionary models and ancestral state reconstruction in phylogenetics

A phylogeny represents the evolutionary relationships among instances based on similarities in their features (traits) [23]. Each leaf of a tree often represents, for example, a specimen, a species or a family. In this context, an internal node is interpreted as the most recent common ancestor of its descendants.

It is often reasonable to posit that these ancestors have unobserved feature values. In this context, *ancestral state reconstruction* is the task of estimating unobserved feature values using *phylogenetic evolutionary models* based on the phylogeny and feature values in the dataset. This reconstruction is performed by probabilistic evolutionary methods like SIMMAP [24], for categorical

features and by methods based on Brownian motion [25–27], for continuous features. These models can also provide feature value predictions at the leaves.

The ancestral state reconstruction performed by these evolutionary methods are illustrated in Fig. 3, which was obtained using the *phytools* package [27].<sup>2</sup> Fig. 3(a) illustrates ancestral state reconstruction for a categorical binary feature. The pie chart over each node is the estimated probability distribution of its feature value. Fig. 3(b) illustrates ancestral state reconstruction for a continuous feature. The color of each node represents the expected value of its feature value. In this case, the expected value of the feature along the edges can be calculated using a Brownian motion model.

## 3. Methodology

Since both dendrograms and phylogenies are formally represented by a tree, it is possible to use common interpretations and

<sup>2</sup> <http://www.phytools.org/Cordoba2017/ex/15/Plotting-methods.html>

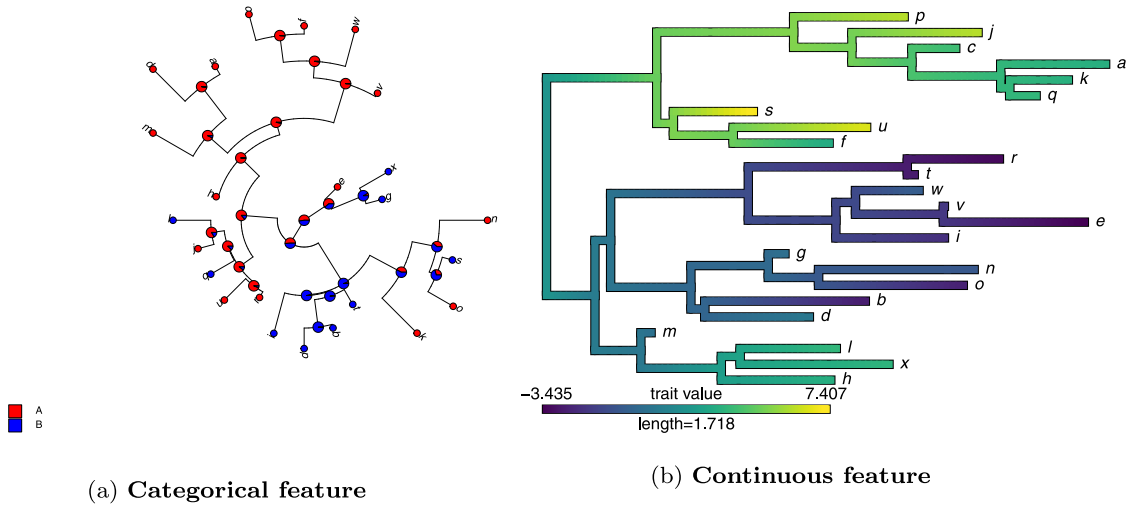


Fig. 3. Illustration of ancestral state reconstruction methods.

intuition of one in the other. In particular, it is useful to interpret an internal node of a dendrogram as an ancestor or an archetype of its descendants. Using this point of view, one can fit evolutionary models to hierarchical clustering and perform ancestral state recognition together with leaf prediction. The following subsections describe some of the benefits of this point of view. Section 3.1 shows useful graphical representations based on ancestral state reconstruction named **evolutionary dendrograms**. Sections 3.2 and 3.3 show how this connection leads, respectively, to a loss for clustering methods (CVL) and to a measure for feature importance (PFIS) in a dendrogram.

These tools can be combined into a complete hierarchical clustering analysis framework as shown in the flowchart of Fig. 4. First, we fit all hierarchical clustering methods we desire based on the input data. Next, for each dendrogram, an evolutionary model predicts each feature's values through a leave-one-feature-out strategy. These predictions yield a loss (CVL) for each dendrogram, thus allowing the selection of the one with best performance. For the selected dendrogram, we compute feature importance score using PFIS. This score uses the evolutionary model and a leave-one-observation-out strategy to determine how well each feature is explained by the dendrogram. We also visualize how the chosen dendrogram explains each feature using the probabilistic evolutionary model. This model uses ancestral state reconstruction to plot how each feature evolves over the dendrogram.

In what follows, we denote the dataset by

$$\mathbb{D} = [x_{i,j}]_{i=1,\dots,n}^{j=1,\dots,p},$$

where  $i$  indexes the instances, and  $j$  indexes the features. Also, we denote a dendrogram by  $\mathcal{T}$ . When we want to emphasize that  $\mathcal{T}$  is built using data  $\mathbb{D}$ , we denote it by  $\hat{\mathcal{T}}$ .

### 3.1. Evolutionary dendrograms

By interpreting a dendrogram as a phylogeny, one can fit a probabilistic evolutionary model to the dendrogram and obtain insights about the clustering through ancestral state reconstruction. This approach offers a multiresolution visualization of the data, which does not require a pre-specified number of clusters. More specifically, this approach describes how each feature behaves in each internal node of the dendrogram. Such a visualization provides insight about what partition size yields a more useful description of the data.

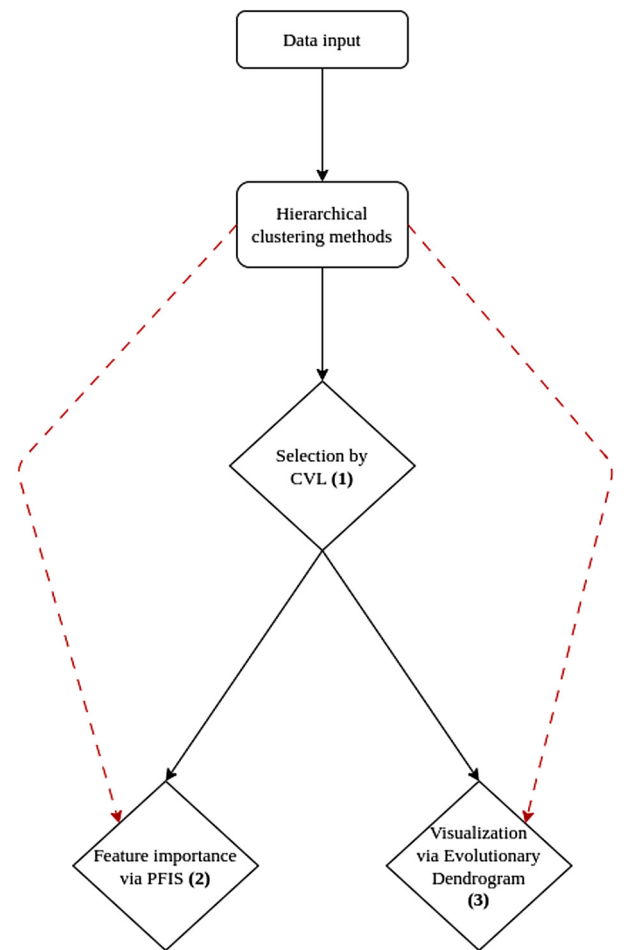
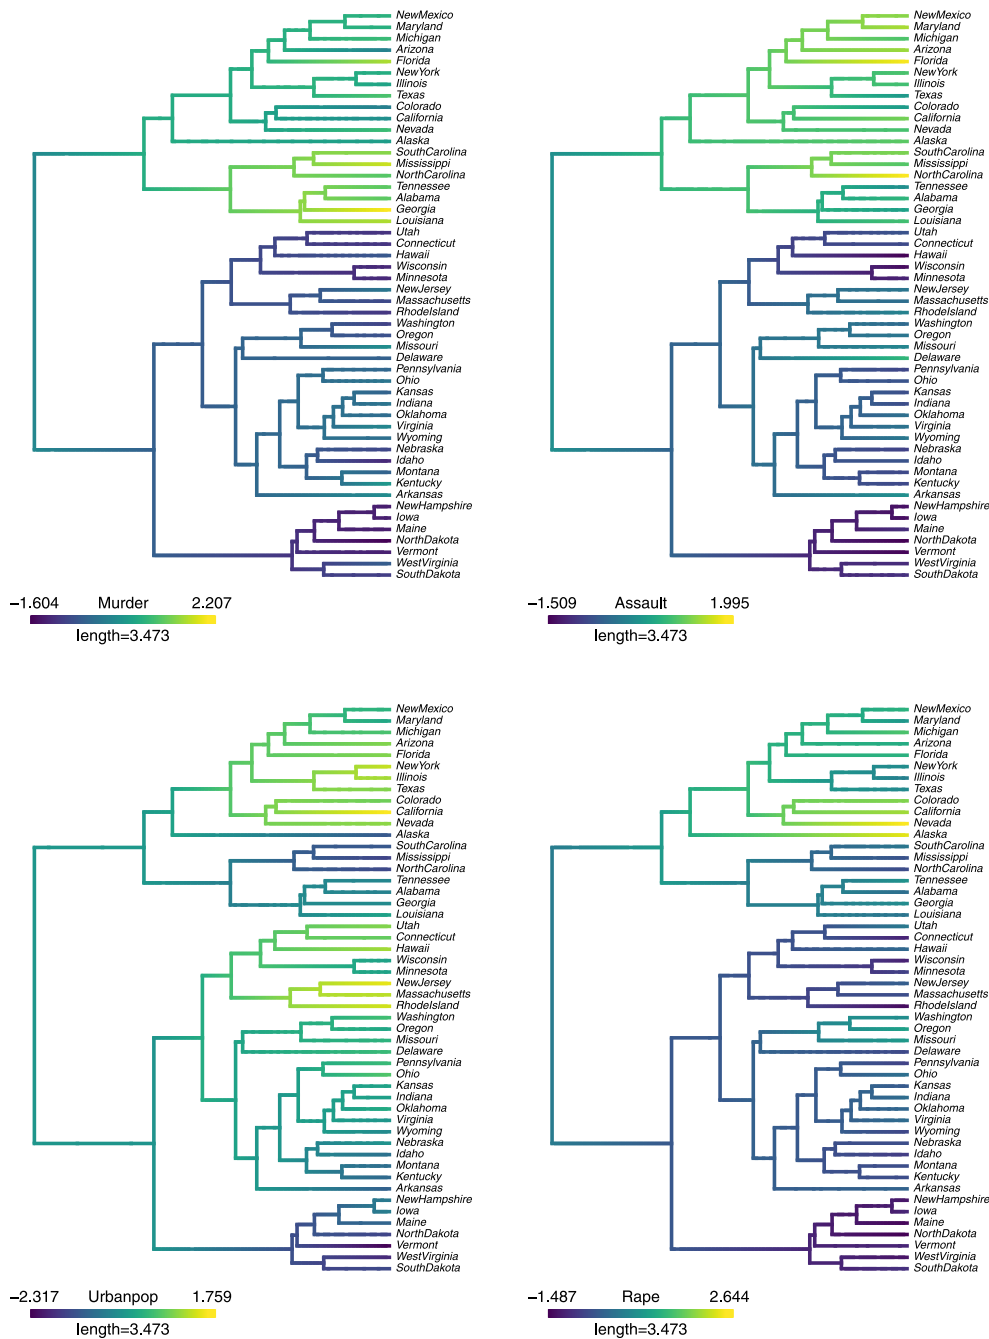


Fig. 4. Flowchart of our full (only continuous arrows) and partial (dashed arrows) framework. The numbers displayed in parenthesis show which question each step answers. The full framework selects the best performing hierarchical clustering method according to our proposed loss (CVL). Afterwards, both feature importance (PFIS) and evolutionary dendrograms are plotted to provide insights and to better explain the dendrogram that was selected by CVL. If we already have a single hierarchical method, the selection step can be skipped and we can proceed directly to plotting feature importance and evolutionary dendrograms, as indicated by the dashed arrows.



**Fig. 5.** Evolutionary dendrograms for standardized **Murder** (top left), **Assault** (top right) and **Rape** (bottom right) rates per 100,000 residents and standardized **Urban population** (bottom left) applied to a clustering obtained using the McQuitty method.

Because ancestral state reconstruction visualization tools are already implemented in phylogenetic analysis packages like *phytools* [27] and *ape* [28] both for continuous and categorical traits, it is straightforward to apply these methods to the hierarchical clustering context: one only needs to take the resulting dendrogram as the input for these procedures. We provide code for this simple adaptation in <https://github.com/Monoxido45/PhyloHclust> and illustrate this graphical analysis with an application to the USArrests [29] and Iris [30] datasets.

### 3.1.1. Evolutionary dendrograms applied to the USArrests dataset

The USArrests dataset [29] describes the urban population and murder, assault and rape rates per 100,000 residents for each of the USA states. Using this data, we obtain a dendrogram for

the USA states using the McQuitty method and perform ancestral state reconstruction applying a fitted Brownian Motion evolution model along the edges [23,27,31]. Fig. 5 shows the evolutionary dendrograms for standardized **Murder** (top left), **Assault** (top right) and **Rape** (bottom right) rates per 100,000 residents and standardized **Urban population** (bottom left). Through the evolutionary dendrogram, we can obtain valuable insights, such as the following:

- Since the major branches of the top dendrograms have similarly colored descendants in both cases, both **Assault** and **Murder** features are well segmented by the dendrogram.
- In the top left dendrogram, the major internal nodes vary from high (yellow) to low murder rates (dark blue), with



high **Murder** rates descendant states listed from South Carolina to Louisiana, medium murder rates (green) from New Mexico to Alaska, mild to low rates (light blue) from Utah to Arkansas, low rates (dark blue) from New Hampshire to South Dakota.

- In the right dendrogram, **Assault** rates are segmented in a similar way as **Murder**'s dendrogram. The main difference from the previous analysis is that high assault rates are not well separated from medium assault rates.
- One might use a coarse partition which divides the states into three groups: one with medium to high murder and attack rates, another with mild rates, and the last with low rates. Furthermore, one could also refine the partition by dividing the first group into two others with respectively high and medium murder rates.
- The deep internal nodes of **Rape** rate dendrogram are colored similarly to **Assault** rate dendrogram. That is, these features are segmented similarly by the dendrogram.
- **Urban population** is not well segmented by the dendrogram: the major branches of the **Urban Population** dendrogram have differently colored descendants, and almost all of its deep internal nodes are green. Thus, we could obtain more internally homogeneous groups by selecting internal nodes that are closer to the leaves, but these groups would be similar to others that are not in the same major branch. For instance, we can select internal nodes in the first major branch with descendants listed from South Carolina to Louisiana and internal nodes in the second major branch with descendants listed from New Hampshire to South Dakota.

It would be much harder to perform the above analysis directly through a standard partition-based boxplot analysis. One would need to cut the dendrogram at several different heights and then gather the collection of boxplots obtained for each feature and height. By doing so, the multi-resolution view offered by the dendrogram is lost and clustering cannot be visualized at an individual level. We further detail this comparison in Section 5.1.

### 3.1.2. Evolutionary dendrograms applied to the iris dataset

The Iris dataset [30] contains measurements related to 150 flowers. Specifically, it contains the width and length of the sepals and petals of flowers which are classified into three species: **setosa**, **versicolor**, and **virginica**. In this analysis, we performed clustering using complete linkage over the petal and sepal measurements.

Although the species were not used in the clustering, a graphical analysis shows how well the resulting dendrogram segments this label. Fig. 6 illustrates the evolutionary dendrogram for species using SIMMAP [24]. The Figure shows that, using a partition of size 5, the dendrogram segments the label adequately.

### 3.2. Cross-validated loss (CVL)

Before one can interpret a dendrogram, it is necessary to choose a clustering algorithm. In order to do this, we propose a new loss function for evaluating the performance of each algorithm.

The connection between dendrograms and phylogenies can also be explored for choosing between clustering algorithms. The main advantage of this relation is that a phylogeny induces predictions over instances. As a result, it is possible to adapt to clustering the methods for model selection that are used in supervised learning. More specifically, our proposal is an adaptation of a leave-one-feature-out cross-validation approach for selecting

phylogenies [20]. Let  $\hat{\tau}$  be a dendrogram fitted according to an arbitrary method. The score is given by

$$CVL(\hat{\tau}) := p^{-1} \sum_{j=1}^p L_j(\hat{\tau}_{-j}),$$

where

$$L_j(\hat{\tau}_{-j}) := n^{-1} \sum_{i=1}^n d_j(\hat{x}_{i,j}, x_{i,j}),$$

$\hat{\tau}_{-j}$  is the dendrogram fitted according to the same method, but without the  $j$ th feature,  $\hat{x}_{i,j}$  is a prediction  $\hat{\tau}_{-j}$  gives to  $x_{i,j}$ , the  $j$ th feature of the  $i$ th instance, and  $d_j(\hat{x}_{i,j}, x_{i,j})$  as the inaccuracy of  $\hat{x}_{i,j}$  with respect to  $x_{i,j}$ . The proposal is summarized in Algorithm 1. For Algorithm 1 to be operational, it is necessary to choose a prediction function and a measure of inaccuracy of predictions. The choice of these functions depends on whether the feature under analysis is categorical or continuous.

---

#### Algorithm 1 Cross-validation loss (CVL)

---

**Input:** Clustering algorithm,  $\hat{\tau}$ , and data  $x_{i,j}$ , for  $1 \leq i \leq n$  and  $1 \leq j \leq p$ .

**Output:**  $CVL(\hat{\tau})$ , the cross-validated predictive loss of  $\hat{\tau}$ .

---

- 1: **for** each feature,  $j$ , **do**
  - 2:   Let  $\hat{\tau}_{-j}$  be obtained by applying  $\hat{\tau}$  to the data with feature  $j$  removed.
  - 3:   **for** each instance,  $i$ , **do**
  - 4:     Let  $\hat{x}_{i,j}$  be a prediction for  $x_{i,j}$  based on  $\hat{\tau}_{-j}$  and on  $x_{i',j}$  for  $i' \neq i$ .
  - 5:     Let  $d_j(\hat{x}_{i,j}, x_{i,j})$  measure how inaccurate prediction  $\hat{x}_{i,j}$  is for  $x_{i,j}$ .
  - 6:   **end for**
  - 7:   Let  $L_j(\hat{\tau}_{-j}) := n^{-1} \sum_{i=1}^n d_j(\hat{x}_{i,j}, x_{i,j})$    ▷ Average inaccuracy for  $j$
  - 8: **end for**
  - 9: Let  $CVL(\hat{\tau}) := p^{-1} \sum_{j=1}^p L_j(\hat{\tau}_{-j})$    ▷  $\hat{\tau}$ 's average inaccuracy
- 

When the feature is categorical, we consider that the prediction is a list of probabilities for each category. Formally, if  $x_{i,j} \in \mathcal{A}$ , then  $\hat{x}_{i,j} = (\pi_{i,j,a})_{a \in \mathcal{A}}$ , where  $\pi_{i,j,a}$  is an estimate of the probability that  $x_{i,j} = a$ . Determining these probabilities based on the dendrogram is analogous to inferring missing characters on phylogenies, as performed for example by SIMMAP. As illustrated in Fig. 7, the main goal of such evolutionary models is to predict feature values that are missing using the data from other leaves and the information of the phylogeny (in our case, the dendrogram).

Hence, we propose using SIMMAP for computing  $\pi_{i,j,a}$ . Using such predictions,  $d_j$  is defined as the Brier score, that is,


$$d_j((\pi_{i,j,a})_{a \in \mathcal{A}}, x_{i,j}) = |\mathcal{A}|^{-1} \sum_{a \in \mathcal{A}} (\mathbb{I}(x_{i,j} = a) - \pi_{i,j,a})^2.$$

When the testing feature is continuous, we consider that the prediction is an estimate of a typical feature value. Drawing from the phylogenetical analogy,  $\hat{x}_{i,j}$  is computed using the Brownian Motion process [25,26], which was originally developed to infer continuous features for internal nodes and leaves of a phylogeny. In this case, each instance  $x_{i,j}$  is standardized (that is, it is normalized to have mean zero and variance one) before fitting the Brownian Motion process, and  $d_j$  is defined as the squared error,


$$d_j(x_{i,j}, \hat{x}_{i,j}) = (x_{i,j} - \hat{x}_{i,j})^2. \quad (1)$$

### 3.3. Phylogenetic feature importance score (PFIS)

Based on the idea that dendrograms equipped with an evolutionary model provide predictions over their leaves, it is also



**Fig. 6.** Evolutionary dendrogram obtained by applying ancestral state reconstruction of species based on the SIMMAP algorithm. The original dendrogram is obtained by performing complete linkage clustering to petal and sepal widths and heights.



**Fig. 7.** Example of a phylogeny with missing character on a leaf. Evolutionary models such as SIMMAP predict the value of the feature indicated with a question mark based on the phylogeny (dendrogram) and the other known characters.

possible to construct a score for feature importance. Let  $\hat{x}_{i,j}(\hat{\tau})$  be a prediction for  $x_{i,j}$  based on the fitted dendrogram  $\hat{\tau}$  and on  $x_{i',j}$ , for  $i' \neq i$ , obtained using an evolutionary model as in the last section. We define the Phylogenetic Feature Importance Score (PFIS) of feature  $j$  over  $\hat{\tau}$ ,  $\mathcal{I}_j(\hat{\tau})$ , as:

$$\mathcal{I}_j(\hat{\tau}) = 1 - \frac{\sum_{i=1}^n d_j(\hat{x}_{i,j}(\hat{\tau}), x_{i,j})}{n},$$

where  $n^{-1} \sum_{i=1}^n d_j(\hat{x}_{i,j}(\hat{\tau}), x_{i,j})$  measures the *inaccuracy* of  $\hat{\tau}$ 's in predicting  $j$ . Notice that, for standardized continuous features and the squared error (Eq. (1)),  $L_j(\hat{\tau})$  is exactly the ratio between sum of squares of residuals and total sum of squares, with  $SS_{res} = \sum_{i=1}^n d_j(\hat{x}_{i,j}(\hat{\tau}), x_{i,j})$  and  $SS_{tot} = n$ . It follows that, in this case,  $\mathcal{I}_j(\hat{\tau})$  is a coefficient of determination,  $R^2$ , often used to evaluate the explanation power of statistical models in regression analysis [32].

The main purpose of PFIS is to filter features according to how relevant they are in the segmentation of the dendrogram.

For instance, the features in Section 3.1 with small importance scores lead to evolutionary dendrograms without a clear structure. By knowing this in advance, one focus on looking at the more informative figures.

#### 4. Related work

To the best of our knowledge, our framework to graphical analysis and feature importance is the first one that uses the full hierarchical structure given by a dendrogram. All previous approaches involve mapping the dendrogram to a partition of instances, as we detail in the next subsections.

##### 4.1. Graphical analysis

Graphical analysis is a commonly used method for evaluating clustering algorithms with a fixed number of groups. This typically involves plotting the distribution of each feature on each cluster using various visualization techniques such as boxplots, barplots, histograms, or radar charts [13–15].

Another method for visualizing clusters is by creating a scatter plot of the features projected into a low-dimensional space using techniques such as principal component analysis [33] or t-SNE [34]. For a more in-depth review of these techniques and their variations, refer to Hennig [35].

It is worth noting that these visualization tools can also be used for hierarchical clustering methods, however, mapping the dendrogram into a partition can lead to the loss of important information (see Section 5.1 for a comparison).

We are not aware of visualization techniques specifically designed for hierarchical clustering methods that allow for the interpretation of how features vary across multiple resolution clusters.

##### 4.2. Feature importance

Several feature importance metrics have been developed for specific clustering algorithms. For instance, Pfaffel [36] proposes

a solution for  $k$ -means clustering, while Zhu [37] presents a framework for determining feature importance in model-based clustering. Similarly, Badih et al. [17] introduce a method for assessing feature importance in decision tree clustering.

However, there exist only a few feature importance scores that can be applied to arbitrary partition-based clustering algorithms. For instance, Ismaili et al. [16] trains a classifier that predicts the cluster of each instance based on its feature values. The feature importance for the clustering is identified with the feature importance for the classifier that was trained. For example, one might use the mean decrease accuracy for Random Forests [38], the absolute coefficient for lasso [39], or other importance metrics used in supervised learning [40–43]. Badih et al. [17] is closer to our approach, and is also based on a leave-one-variable out approach (LOVO). However, the feature importance is based on computing a within-cluster heterogeneity measure.

These approaches can only be used for hierarchical clustering by mapping the dendrogram into a partition, which leads to loss of information. We compare the above approaches against our method in Section 5.2.

#### 4.3. Clustering scores

Numerous metrics have been developed to assess the performance of partition-based clustering. Among such metrics are the non-overlap score [3], average distance between means [3], average distance [3], adjusted Rand index [10], V-measure [11], Silhouette coefficient [44], and Calinski–Harabasz index [45]. A comprehensive review of these metrics can be found in Vendramin et al. [46]. However, to the best of our knowledge, the only metric that performs cross-validation in a predictive framework, is Figure of Merit (FOM; Yeung et al. [12]), which we review in Section 5.3.

In order to apply the above approaches to a hierarchical clustering, one must map the dendrogram into a partition. Therefore, two dendrograms that are mapped to the same partition have the same score. For this reason, as demonstrated in Section 5.3, CVL commonly outperforms FOM.

A metric that is similar in concept to CVL is the cophenetic correlation coefficient [47–49]. This coefficient assesses the accuracy of the entire hierarchical structure by borrowing ideas from phylogenetic analysis. It measures the correlation between the euclidean distance between instances and the cophenetic distance induced by the phylogeny. However, it does not employ cross-validation, which can result in overfitting.

Section 5.3 compares CVL to the cophenetic correlation coefficient, Figure of Merit and Adjusted Rand Index.

## 5. Experiments

In this section, we compare our proposal for graphical analysis, feature importance and clustering loss to other methods based on partition clustering. In order to apply the latter methods, partitions are generated by cutting the dendrogram at varying heights.

#### 5.1. Evolutionary dendrograms vs. grouped boxplots

Boxplots are a common way of visualizing how a feature is segmented by a dendrogram. First, a partition is obtained by cutting the dendrogram at a given height. Next, a boxplot is drawn for each cluster in the partition. By successively cutting the dendrogram at different heights, one can fully visualize how the feature is segmented. However, much of the information in the dendrogram is lost in this procedure.

Fig. 8 compares such segmented boxplots to evolutionary dendrograms using the USArrrests dataset, which was analyzed in details (via evolutionary dendrograms) in Section 3.1.1. First, a dendrogram is obtained using the McQuitty method. Next, segmented boxplots and the evolutionary dendrogram are obtained for the **Murder** rate feature.

Fig. 8(a) presents boxplots grouped by cluster when 2, 3 and 4 clusters are used. Using these plots, one can see that the feature is well segmented for the chosen number of clusters. However, although boxplots bring measures of centrality and dispersion, they lose information about the internal structure of each cluster. In particular, one cannot find in the boxplots which instances belong to each cluster.

In contrast Fig. 8(b) displays the evolutionary dendrogram obtained for the same feature. By cutting the dendrogram at the root, one obtains the main top (B: green) and bottom (A: light blue) clusters. By looking at the color dispersion in both clusters, it is possible to conclude that the top cluster has a larger variance than the bottom cluster. Furthermore, by looking at the next internal nodes, one can observe that while the top cluster can be further segmented into dark green (F) and light green (E) sub-clusters, the bottom cluster can be segmented into purple (C) and dark blue (D) sub-clusters.

#### 5.2. PFIS vs. state-of-the-art importance scores

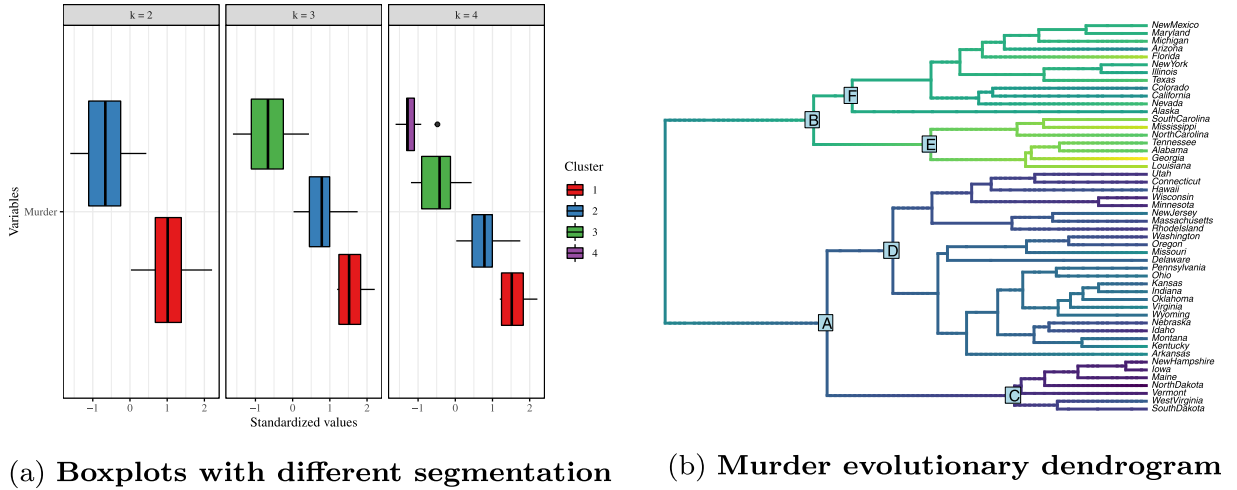
This subsection compares PFIS, presented in Section 3.3, to common alternatives for assessing feature importance in a dendrogram. We use the following baselines:

- [Classification-based importance – RF [16,17]]. We first cut the dendrogram into  $k$  clusters. Next, these clusters are used as labels for a random forest classifier [38] that uses the original features as inputs. Finally, the importance scores inferred by the random forest are used as the importance of the features for the dendrogram.
- [Classification-based importance – XGBoost [16,17]]. Same as before, but using a XGboost classifier [50] instead.
- [LOVO [17]]. A leave-one-variable-out (LOVO) approach that compares within-cluster heterogeneity by omitting each feature from the dataset. This approach also requires the dendrogram to be cut into  $k$  clusters.


In order to compare the performance of these approaches to PFIS, we need a dataset in which the ground-truth importance is known. Thus, we simulated data in a way that allows the true importance to be evaluated. Specifically, let  $X_{i,j}$  be the  $j$ th simulated feature value of the  $i$ th instance, where  $1 \leq j \leq 6$  and  $1 \leq i \leq 2^8$ . The data,  $X$ , is generated recursively according to a tree procedure described in Algorithm 2. In this algorithm,  $\theta_{i,j,k-1}$  represents an ancestor of  $\theta_{i,j,k}$  and of  $\theta_{2^{k-1}+i,j,k}$ . Also,  $z_{i,j,k}$  represents the noise that is added when generating a descendant. Algorithm 2 summarizes this procedure. Observe that  $z_{i,j,k}$  is of the order of magnitude of  $(\sigma_j)^k$ . That is, the larger the value of  $\sigma_j$ , the more noise is added to the leaf descendants proportionally to the internal nodes of the tree. As a result, the larger the value of  $\sigma_j$ , the less the structure of the tree is preserved by feature  $j$ . Since in the simulation  $\sigma_j = 2^{j-3}$ , one expects that the feature importance in the adjusted dendrogram should decrease from  $j = 1$  to 6.

Fig. 9 shows the various importance scores to the data described above. The dendrogram was obtained through Ward's method, and the scores that required a fixed number of clusters used  $k = 3$  groups, which is the optimal number of clusters according to the NbClust majority rule [51]. The right panel presents  $\sigma_j^{-1}$ , a measure of how much the tree structure is preserved by each feature. The figure shows that our method (PFIS) is the only one to preserve the ordering of the features given by  $\sigma_j^{-1}$ , and thus uncover the true feature importance values.





**Fig. 8.** Comparison between the boxplot visualization and the evolutionary dendrogram on the murder feature from the USArrests dataset. Evolutionary dendrograms provide more nuanced information as they take the full dendrogram hierarchy into account.



**Fig. 9.** Measures of feature importance: (far left) random forest [16], (middle left) XGboost [16], (middle) leave-one-variable-out [17], (middle right) PFIS, (far right)  $\sigma_j^{-1}$ , a gold standard (ground truth) that measures how much the tree structure that is used to generate the data is preserved by each feature. Our score preserves the ordering given by  $\sigma_j^{-1}$ , while all the others do not.

**Algorithm 2** Simulated data to compare feature importance methods

**Output:**  $X_{i,j}$ , the simulated data.

- 1: Let  $\sigma = (2^{-2}, 2^{-1}, 2^0, 2^1, 2^2, 2^3)$ .
- 2: Let  $\theta_{1,j,0} = 0$ , for  $1 \leq j \leq 6$ .
- 3: **for** each  $k \in \{1, \dots, 7\}$ , **do**
- 4:   **for** each  $i \in \{1, \dots, 2^{k-1}\}$  and  $j \in \{1, \dots, 6\}$ , **do**
- 5:     Let  $i_1 = i$  and  $i_2 = 2^{k-1} + i$
- 6:     Let  $\theta_{i_1,j,k} = \theta_{i_1,j,k-1} + z_{i_1,j,k}$ , where  $z_{i_1,j,k} \sim N(0, (\sigma_j)^k)$
- 7:     Let  $\theta_{i_2,j,k} = \theta_{i_1,j,k-1} + z_{i_2,j,k}$ , where  $z_{i_2,j,k} \sim N(0, (\sigma_j)^k)$
- 8:   **end for**
- 9: **end for**
- 10: Let  $X_{i,j} = \theta_{i,j,7}$ , for  $1 \leq i \leq 2^8$  and  $1 \leq j \leq 6$ .

**5.3. Cross-validated loss vs. state-of-the-art dendrogram selection tools**

We use simulated and real data to compare our CVL loss to other state-of-the-art dendrogram selection tools. For each experiment we will use different combinations of clustering methods and linkages in order to assess the quality and performance of each methodology. The combinations of methods and linkages are listed in Table 1.

This section evaluates the performance of cross-validated loss (CVL) by comparing it to the following baselines:

- [Cophenetic Correlation score – Coph [47–49]]. The idea of the cophenetic correlation is to compute the Pearson correlation between the Euclidean distance and the dendrogrammatic (or cophenetic) distance obtained for each different pair of observations. The dendrogrammatic distance is given by the height of the node of the obtained

**Table 1**

Clustering methods used for comparing CVL to FOM, COPH and ARI.

Method	Linkages
Agnes	weighted, average and ward
Agglomerative	ward.D, complete, single, ward.D2, average, mcquitty and median
Diana	Not applicable

**Table 2**

Datasets from UCI used to compare CVL to FOM.

Dataset	Sample size	Number of features	Label
Simulated	250	2	Y
Iris	150	4	Species
Diabetes	768	8	V9
Wheat seeds	210	7	V8
Ionosphere	351	34	label
Glass	214	10	V11
Haberman	306	3	V4
Wine	178	13	V1

dendrogram at which each pair of observations are first joined together.

- [Figure of Merit – FOM [12]]. First, for each feature, a dendrogram is adjusted holding out the feature from the data. Next, by cutting each dendrogram at a given height, it is transformed into a partition. Finally, the within-cluster similarity is computed for each held-out feature using its respective partition. The FOM score corresponds to the average of these similarities.
- [Adjusted Rand Index – ARI [10]]. The standard rand index [52] measures the similarity between a clustering of the data and a ground truth label by considering all pairs of observations and counting the pairs that are assigned to the same cluster. ARI adds a correction to the rand index that takes into account assignment to the same clustering by chance.

We compare CVL to these scores using simulated datasets and UCI datasets<sup>3</sup> that are available on R. The chosen UCI datasets are summarized in Table 2. We also use a simulated dataset with 250 instances such that the label is  $Y_i \sim \text{Bernoulli}(0.5)$  and the features are  $(X_1, X_2) | Y \sim N_2([2Y \quad 2Y], I)$ . Note that although clustering methods use only features, each dataset also has a label.

We compare all methods to a gold-standard which uses the labels in the datasets. The main idea is that the dendrogram is treated as a classifier for the label and the F1 score is the gold-standard for the dendrogram's performance. Such a classifier is obtained by cutting the dendrogram at an appropriate height so that it becomes a partition with size equal to the number of labels. Next, each cluster in the partition is associated to a single label that maximizes accuracy. Each instance is labeled equally as the partition to which it belongs. The gold-standard for the dendrogram's performance is the F1 score based on its assigned labels.

Based on the above information, it is possible to describe the experiment that compares CVL, FOM, Cophenetic correlation and ARI. For each dataset, we performed clustering using all of the methods in Table 1. Each of the 11 clustering methods was applied using the euclidean, manhattan and canberra distances. As a result, 33 dendrograms were obtained for each dataset. For each dendrogram, the CVL, FOM, COPH, ARI and gold-standard F1 scores were calculated. Since FOM and ARI require a pre-defined

**Table 3**

For each dataset, Spearman correlation between CVL and F1, CVL and FOM, COPH and F1 and ARI and F1 among the 33 clustering methods used in the experiment. In order to maintain the comparison sign for all losses/scores, we compared minus the value of COPH and ARI to F1, as indicated by the sign in front of both scores. Bold values indicate which of the losses/scores has the lowest correlation with  $F_1$  and, therefore, is in better agreement with this performance measure.

Dataset	Correlations			
	CVL( $\tau$ ) vs $F_1$	FOM vs $F_1$	–COPH vs $F_1$	–ARI vs $F_1$
Simulated	<b>–0.40</b>	0.23	–0.23	0.88
Iris	0.26	0.10	<b>–0.54</b>	–0.30
Pima indians	–0.53	0.52	<b>–0.63</b>	–0.06
Wheat seeds	–0.11	–0.48	0.28	<b>–0.83</b>
Ionosphere	<b>–0.43</b>	0.63	–0.12	0.60
Glass	–0.15	0.19	<b>–0.48</b>	0.14
Haberman	<b>–0.42</b>	0.38	–0.16	–0.09
Wine	–0.21	–0.84	–0.35	<b>–0.92</b>

number of clusters, we calculated both scores using the correct number of labels. Since the correct number of clusters is generally unavailable, this procedure gives an advantage to FOM and ARI.

Table 3 shows, for each dataset, the Spearman correlation between CVL and  $F_1$ , FOM and  $F_1$ , –COPH and  $F_1$  and between –ARI and  $F_1$  among the clustering methods that were used. Recall that, while CVL and FOM measure how bad is a given clustering, the gold standard  $F_1$  and the scores ARI and COPH are methods that measure whether the clustering is good. Therefore, a negative correlation between a loss or minus a score and  $F_1$  indicates that the measures are in agreement. One can observe in Table 3 that, while the Spearman correlation between CVL and  $F_1$  and –COPH and  $F_1$  is negative for all but one dataset, it is positive between FOM and  $F_1$  and –ARI and  $F_1$  except for two and five datasets respectively. Furthermore, both CVL and COPH outperformed the other methods in three out of the eight datasets each, while ARI performed better for two datasets. Thus, these results provide evidence that both the cophenetic correlation and CVL adequately assess the performance of clustering methods without the need to partition the dendrogram. In practice, they can complement each other to perform dendrogram selection.

## 6. Final remarks

We developed a novel paradigm for hierarchical clustering analysis. This framework allows the practitioner to choose between clustering algorithms, to interpret the resulting dendrograms, and to measure the importance of each feature. Contrary to existing approaches, this is done without losing the rich multi-resolution structure provided by the dendrograms. Indeed, our experiments show that our framework leads to easier-to-interpret and more accurate inference than competing approaches.

Our paradigm is based on viewing a dendrogram as a phylogeny. This connection opens possibilities that go beyond this work; hierarchical clustering analysis may now be enriched with other tools developed for phylogenetic analyses. The cross-validated loss, the importance score, and the evolutionary dendrogram are implemented in R at <https://github.com/Monoxido45/PhyloHclust>.

## CRedit authorship contribution statement

**Luben M.C. Cabezas:** Conceptualization, Investigation, Methodology, Software, Validation, Visualization, Writing – review & editing. **Rafael Izbicki:** Conceptualization, Investigation, Methodology, Supervision, Visualization, Writing – review & editing. **Rafael B. Stern:** Conceptualization, Investigation, Methodology, Supervision, Visualization, Writing – review & editing.

<sup>3</sup> <https://archive.ics.uci.edu/ml/datasets.php>

## Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Luben Miguel Cruz Cabezas reports financial support was provided by State of Sao Paulo Research Foundation. Rafael Izbicki reports a relationship with State of Sao Paulo Research Foundation that includes: funding grants. Rafael Izbicki reports a relationship with National Council for Scientific and Technological Development that includes: funding grants. Rafael Bassi Stern reports a relationship with State of Sao Paulo Research Foundation that includes: funding grants.

## Data availability

I have shared the link to my code/data repository in this paper.

## Acknowledgments

L. M. C. C. is grateful for the scholarship provided by Fundação de Amparo à Pesquisa do Estado de São Paulo (FAPESP), Brazil, grant 2020/10861-7. R. I. is grateful for the financial support of FAPESP, Brazil (grant 2019/11321-9) and CNPq, Brazil (grants 309607/2020-5 and 422705/2021-7). R. B. S. produced this work as part of the activities of FAPESP, Brazil Research, Innovation and Dissemination Center for Neuromathematics (grant 2013/07699-0). The authors are also grateful for the suggestions given by Leonardo M. Borges and Victor Coscrato.

## References

- [1] Y. Chen, J. Kim, H. Mahmassani, Pattern recognition using clustering algorithm for scenario definition in traffic simulation-based decision support systems, in: 2014 17th IEEE International Conference on Intelligent Transportation Systems, ITSC 2014, 2014, <http://dx.doi.org/10.1109/ITSC.2014.6957787>.
- [2] L.M. Rocha, F.A.M. Cappabianco, A.X. Falcão, Data clustering as an optimum-path forest problem with applications in image analysis, *Int. J. Imaging Syst. Technol.* 19 (2) (2009) 50–68, <http://dx.doi.org/10.1002/ima.20191>, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/ima.20191>.
- [3] S. Datta, S. Datta, Comparisons and validation of statistical clustering techniques for microarray gene expression data, *Bioinformatics* 19 (4) (2003) 459–466, <http://dx.doi.org/10.1093/bioinformatics/btg025>.
- [4] N. Jardine, C. van Rijsbergen, The use of hierarchic clustering in information retrieval, *Inf. Storage Retr.* 7 (5) (1971) 217–240, [http://dx.doi.org/10.1016/0020-0271\(71\)90051-9](http://dx.doi.org/10.1016/0020-0271(71)90051-9), URL <http://www.sciencedirect.com/science/article/pii/0020027171900519>.
- [5] J. MacQueen, Some methods for classification and analysis of multivariate observations, in: *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, University of California Press, Berkeley, Calif., 1967, pp. 281–297, URL <https://projecteuclid.org/euclid.bsmsp/1200512992>.
- [6] K. Fukunaga, L. Hostetler, The estimation of the gradient of a density function, with applications in pattern recognition, *IEEE Trans. Inform. Theory* 21 (1) (1975) 32–40.
- [7] J.H. Ward Jr., Hierarchical grouping to optimize an objective function, *J. Amer. Statist. Assoc.* 58 (301) (1963) 236–244.
- [8] P.J. Rousseeuw, L. Kaufman, *Finding Groups in Data*, Vol. 1, Hoboken: Wiley Online Librar, 1990.
- [9] J. Liu, Y. Bai, J. Kang, N. An, A new approach to hierarchical clustering using partial least squares, in: 2006 International Conference on Machine Learning and Cybernetics, 2006, pp. 1125–1131.
- [10] L. Hubert, P. Arabie, Comparing partitions, *J. Classification* 2 (1985) 193–218, <http://dx.doi.org/10.1007/BF01908075>.
- [11] A. Rosenberg, J. Hirschberg, V-Measure: A conditional entropy-based external cluster evaluation measure, in: *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, 2007, pp. 410–420, <http://dx.doi.org/10.7916/D80V8N84>.
- [12] K.Y. Yeung, D.R. Haynor, W.L. Ruzzo, Validating clustering for gene expression data, *Bioinformatics* 17 (4) (2001) 309–318, <http://dx.doi.org/10.1093/bioinformatics/17.4.309>.
- [13] A. Kassambara, *Practical Guide To Cluster Analysis in R: Unsupervised Machine Learning*, Vol. 1, Sthda, 2017.
- [14] J. Seo, B. Shneiderman, Interactively exploring hierarchical clustering results [gene identification], *Computer* 35 (7) (2002) 80–86, <http://dx.doi.org/10.1109/MC.2002.1016905>.
- [15] T. Galili, dendextend: an R package for visualizing, adjusting and comparing trees of hierarchical clustering, *Bioinformatics* 31 (22) (2015) 3718–3720, <http://dx.doi.org/10.1093/bioinformatics/btv428>.
- [16] O.A. Ismaili, V. Lemaire, A. Cornuéjols, A supervised methodology to measure the variables contribution to a clustering, in: C.K. Loo, K.S. Yap, K.W. Wong, A. Teoh, K. Huang (Eds.), *Neural Information Processing*, Springer International Publishing, Cham, 2014, pp. 159–166.
- [17] G. Badih, M. Pierre, B. Laurent, Assessing variable importance in clustering: a new method based on unsupervised binary decision trees, *Comput. Statist.* 34 (2019) 301–321, <http://dx.doi.org/10.1007/s00180-018-0857-0>.
- [18] J.B. Joy, R.H. Liang, R.M. McCloskey, T. Nguyen, A.F. Poon, Ancestral reconstruction, *PLoS Comput. Biol.* 12 (7) (2016) e1004763.
- [19] T. Pupko, I. Mayrose, A Gentle Introduction to Probabilistic Evolutionary Models, No commercial publisher| Authors open access book, 2020.
- [20] L.M. Borges, R. Izbicki, R.B. Stern, The overlooked role of predictiveness in phylogenetics: Data-splitting as a powerful method for model selection, 2021, (submitted for publication).
- [21] G. James, D. Witten, T. Hastie, R. Tibshirani, *An introduction to statistical learning*, Vol. 112, Springer, 2013.
- [22] T. Hastie, R. Tibshirani, J. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Second Edition, in: *Springer Series in Statistics*, Springer New York, 2009, URL <https://books.google.com.br/books?id=tVlJmNS3Ob8C>.
- [23] J. Felsenstein, Phylogenies and the comparative method, *Amer. Nat.* 125 (1) (1985) 1–15.
- [24] J.P. Bollback, SIMMAP: stochastic character mapping of discrete traits on phylogenies, *BMC Bioinformatics* 7 (1) (2006) 88.
- [25] B.C. O'Meara, C. Ané, M.J. Sanderson, P.C. Wainwright, Testing for different rates of continuous trait evolution using likelihood, *Evolution* 60 (5) (2006) 922–933.
- [26] J. Clavel, G. Escarguel, G. Merceron, mvMORPH: an R package for fitting multivariate evolutionary models to morphometric data, *Methods Ecol. Evol.* 6 (11) (2015) 1311–1319.
- [27] L.J. Revell, phytools: an R package for phylogenetic comparative biology (and other things), *Methods Ecol. Evol.* 3 (2) (2012) 217–223.
- [28] E. Paradis, K. Schliep, ape 5.0: an environment for modern phylogenetics and evolutionary analyses in R, *Bioinformatics* 35 (3) (2019) 526–528.
- [29] D.R. McNeil, *Interactive Data Analysis: a Practical Primer*, Wiley, 1977.
- [30] R.A. Fisher, The use of multiple measurements in taxonomic problems, *Ann. Eugen.* 7 (2) (1936) 179–188.
- [31] L.J. Revell, Two new graphical methods for mapping trait evolution on phylogenies, *Methods Ecol. Evol.* 4 (8) (2013) 754–759.
- [32] J. Neter, M.H. Kutner, C.J. Nachtsheim, W. Wasserman, et al., *Applied linear statistical models*, Irwin Chicago, 1996.
- [33] J.F. Hair, *Multivariate data analysis*, 2009.
- [34] L. Van der Maaten, G. Hinton, Visualizing data using t-SNE, *J. Mach. Learn. Res.* 9 (11) (2008).
- [35] C. Hennig, Clustering strategy and method selection, in: *Handbook of Cluster Analysis*, Vol. 9, Chapman and Hall/CRC Boca Raton, 2015, pp. 703–730.
- [36] O. Pfaffel, FeatureImpCluster: Feature importance for partitional clustering, 2021, URL <https://CRAN.R-project.org/package=FeatureImpCluster>, R package version 0.1.5.
- [37] X. Zhu, Variable diagnostics in model-based clustering through variation partition, *J. Appl. Stat.* 45 (16) (2018) 2888–2905.
- [38] L. Breiman, Random forests, *Mach. Learn.* 45 (1) (2001) 5–32.
- [39] R. Tibshirani, Regression shrinkage and selection via the lasso, *J. R. Stat. Soc. Ser. B Stat. Methodol.* 58 (1) (1996) 267–288.
- [40] T. Hastie, R. Tibshirani, J.H. Friedman, J.H. Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Vol. 2, Springer, 2009.
- [41] R. Izbicki, A.B. Lee, Converting high-dimensional regression to high-dimensional conditional density estimation, *Electron. J. Stat.* 11 (2) (2017) 2800–2831.
- [42] V. Coscrato, M.H.d.A. Inácio, T. Botari, R. Izbicki, NLS: an accurate and yet easy-to-interpret regression method, 2019, arXiv preprint arXiv:1910.05206.
- [43] G.Y. Shimizu, R. Izbicki, A.C. de Carvalho, Model interpretation using improved local regression with variable importance, 2022, arXiv preprint arXiv:2209.05371.
- [44] P.J. Rousseeuw, Silhouettes: a graphical aid to the interpretation and validation of cluster analysis, *Comput. Appl. Math.* 20 (1987) 53–65, [http://dx.doi.org/10.1016/0377-0427\(87\)90125-7](http://dx.doi.org/10.1016/0377-0427(87)90125-7).

- [45] T. Caliński, J. Harabasz, A dendrite method for cluster analysis, *Comm. Statist. Theory Methods* 3 (1974) 1–27, <http://dx.doi.org/10.1080/03610927408827101>.
- [46] L. Vendramin, R.J. Campello, E.R. Hruschka, Relative clustering validity criteria: A comparative overview, *Stat. Anal. Data Min. ASA Data Sci. J* 3 (4) (2010) 209–235.
- [47] R.R. Sokal, F.J. Rohlf, The comparison of dendrograms by objective methods, *Taxon* (1962) 33–40.
- [48] S. Saraçlı, N. Doğan, İ. Doğan, Comparison of hierarchical cluster analysis methods by cophenetic correlation, *J. Inequal. Appl.* 2013 (1) (2013) 1–8.
- [49] A. Timofeeva, Evaluating the robustness of goodness-of-fit measures for hierarchical clustering, *J. Phys.: Conf. Ser.* 1145 (1) (2019) 012049.
- [50] T. Chen, C. Guestrin, Xgboost: A scalable tree boosting system, in: *Proceedings of the 22nd Acm Sigkdd International Conference on Knowledge Discovery and Data Mining*, 2016, pp. 785–794.
- [51] M. Charrad, N. Ghazzali, V. Boiteau, A. Niknafs, NbClust: an R package for determining the relevant number of clusters in a data set, *J. Stat. Softw.* 61 (2014) 1–36.
- [52] W.M. Rand, Objective criteria for the evaluation of clustering methods, *J. Amer. Statist. Assoc.* 66 (336) (1971) 846–850.