Order preserving hierarchical agglomerative clustering

Daniel Bakkelund¹

Received: 17 November 2020 / Revised: 15 September 2021 / Accepted: 2 November 2021 /
Published online: 16 December 2021
© The Author(s) 2021

Abstract

Partial orders and directed acyclic graphs are commonly recurring data structures that arise naturally in numerous domains and applications and are used to represent ordered relations between entities in the domains. Examples are task dependencies in a project plan, transaction order in distributed ledgers and execution sequences of tasks in computer programs, just to mention a few. We study the problem of *order preserving hierarchical clustering* of this kind of ordered data. That is, if we have $a < b$ in the original data and denote their respective clusters by $[a]$ and $[b]$, then we shall have $[a] < [b]$ in the produced clustering. The clustering is similarity based and uses standard linkage functions, such as single- and complete linkage, and is an extension of classical hierarchical clustering. To achieve this, we develop a novel theory that extends classical hierarchical clustering to strictly partially ordered sets. We define the output from running classical hierarchical clustering on strictly ordered data to be *partial dendrograms*; sub-trees of classical dendrograms with several connected components. We then construct an embedding of partial dendrograms over a set into the family of ultrametrics over the same set. An optimal hierarchical clustering is defined as the partial dendrogram corresponding to the ultrametric closest to the original dissimilarity measure, measured in the p -norm. Thus, the method is a combination of classical hierarchical clustering and ultrametric fitting. A reference implementation is employed for experiments on both synthetic random data and real world data from a database of machine parts. When compared to existing methods, the experiments show that our method excels both in cluster quality and order preservation.

Keywords Hierarchical clustering · Order preserving clustering · Partial dendrogram · Unsupervised classification · Ultrametric fitting · Acyclic partition

Editor: Michelangelo Ceci.

✉ Daniel Bakkelund
daniel.bakkelund@ifi.uio.no

¹ Institute of Informatics, University of Oslo, P.O. Box 1080, 0316 Oslo, Norway

1 Introduction

Clustering is one of the oldest and most frequently used techniques for exploratory data analysis and unsupervised classification. The toolbox contains a large variety of methods and algorithms, spanning from the initial, but still popular ideas of k -means (Macqueen, 1967) and hierarchical clustering (Johnson, 1967), to more recent methods, such as density- and model based clustering (Kriegel et al., 2011; Fraley & Raftery, 2002), and semi-supervised methods (Basu et al., 2008), plus a large list of variants. All these methods have one thing in common: they try to extract hidden structure from the data, and make it visible to the analyst. But they also share another feature: if the analysed data is already endowed with some form of structure, the structure is lost in the clustering process; the clustering does not try to retain the structure.

In this paper, we show how to extend hierarchical clustering to relational data in a way that preserves the relations. In particular, if the input is a set X equipped with a strict partial order $<$, and if $a, b \in X$, we ensure that if $a < b$ then we will have $[a] <' [b]$ after clustering, where $[a]$ and $[b]$ are the respective clusters of a and b , and $<'$ is a partial order on the clusters naturally induced by $<$.

Since directed acyclic graphs (DAGs) correspond to partial orders, our method works equally well for DAGs. If the input is a DAG, then every clustering in the produced hierarchy is a DAG of clusters, and there exists a DAG homomorphism from the original DAG to the cluster DAG.

1.1 Motivating real-world use case

The motivation for our method comes from an industry database of machine parts that are arranged in part-of relations: parts are registered as sub-parts of other parts. For historical reasons, there have been incidents of copy-paste of machine designs, and the copies have been given entirely new identifiers with no links to the original design. In hindsight, there is a wish to identify these equivalent machine parts, but telling them apart is hard. Also, the metadata that is available has a tendency of displaying high similarity between a part and its sub-parts, leading to “vertical clustering” in the data.

Since the motivation is to identify equivalent machinery with the aim of replacing one piece of machinery with an equivalent part, and since a part and its sub-parts by no means can be interchanged, it is essential to maintain this parent-child relationship. Moreover, since a part and its sub-part are never equivalent, this is a strict order relation. The set of all machine parts thus makes up a strictly partially ordered set. By preserving these relations in the clustering process, we can eliminate the errors due to close resemblance between the part and the sub-part, resulting in improved over all quality of the clustering.

It is possible to imagine other use cases. We choose to mention two; citation network analysis and time series alignment:

Citation networks are partial orders, where the order is defined by the citations. If we perform order preserving clustering in the above sense on citation networks, the clusters will contain related research, and the clusters will be ordered according to appearance relative other related research. This differs from clustering with regards to time: when clustering with time as a parameter, you have to choose, implicitly or explicitly, a time interval for each cluster. When the citation graph is used for ordering, the clusters will contain research

Fig. 1 A dendrogram over the set $X = \{a, b, c, d, e\}$. The elements of X are the leaf nodes of the dendrogram, and, starting at the bottom, the horizontal bars indicate which elements are joined at which step in the process. The numbers on the y-axis indicate at which dissimilarity level the different clusters were formed

that occurred in parallel, citing similar sources, and being cited by similar sources, regardless to whether they occurred in some particular time interval.

A time series is a totally ordered set of events, so that a family of time series is a partially ordered set. Assume that you want to do time series alignment, matching events from one time series with events from another, but for some reason the time stamps are corrupted and cannot be used for this purpose. Given a measure of (dis-)similarity between events, we can cluster the events to figure out which events are the more similar. Since an optimal order preserving clustering is one that both preserves all event orders and matches the most similar events across the time series, ideally the result is a series of clusters with each cluster containing the events that correspond to each other across the time series.

1.2 Problem overview

Given a set X together with a notion of (dis-)similarity between the elements of X , a hierarchical agglomerative clustering can be obtained as follows (Jain & Dubes, 1988, §3.2):

1. Start by placing each element of X in a separate cluster.
2. Pick the two clusters that are most similar according to the (dis-)similarity measure, and combine them into one cluster by taking their union.
3. If all elements of X are in the same cluster, we are done. Otherwise, go to Step 2 and continue.

The result from this process is a dendrogram; a tree structure showing the sequence of the clustering process (Fig. 1).

Now, given a partially ordered set $X = \{a, b, c, d\}$ where $a < b$ and $c < d$, we can use arrows to denote the order relation, thinking of X as a directed acyclic graph with two connected components. If we want to produce a hierarchical clustering of X , while at the same time maintaining the order relation, our options are depicted in the Hasse diagram in Fig. 2.

Each path in this diagram, starting at the bottom and advancing upwards, represents a hierarchical clustering. But, since we are required to preserve the strict order relation, we cannot merge any more elements than what we see here. This means that we will never obtain dendrograms like the one in Fig. 1, that joins at the top when all elements are placed in a single cluster. Rather, the output of hierarchical agglomerative clustering would take the form of *partial dendrograms* like those of Fig. 3.

To complicate matters, if both (a, d) and (a, c) are pairs of minimal dissimilarity, then they are both candidates for the first merge. From Fig. 2 we can see that ad and ac are mutual exclusive merges, and that choosing one over the other leads to very different

Fig. 2 Possible order preserving hierarchical clusterings over the set $X = \{a, b, c, d\}$ with $a < b$ and $c < d$. Adjacent elements indicate clusters

Fig. 3 Partial dendrograms over the set $X = \{a, b, c, d\}$ with $a < b$ and $c < d$. Each partial dendrogram corresponds to a path in Fig. 2 starting at the bottom and advancing upwards to the ordered set depicted below the dendrogram

solutions. We therefore need a method to decide which candidate merge, or which candidate partial dendrogram, is the better.

1.3 Outline of our method and contributions

As our first contribution, to solve the problem of picking one candidate merge among a set of tied connections, we present a permutation invariant method for hierarchical agglomerative clustering. The method uses the classical linkage functions of single-, average- and complete linkage, but is optimisation based, as opposed to the algorithmic definition of classical hierarchical clustering. Recalling that every hierarchical clustering corresponds to a unique ultrametric (Jardine & Sibson, 1971), the optimisation criterion is that of minimising the matrix norm of the difference between the original dissimilarity and the ultrametric corresponding to the hierarchical clustering, a method known as *ultrametric fitting* (De Soete et al., 1987).

We have seen that order preserving hierarchical agglomerative clustering on strictly partially ordered sets leads to partial dendrograms. In order to evaluate the ultrametric fitting of a partial dendrogram, our next contribution is an embedding of partial dendrograms over a set into the family of ultrametries over the same set.

Our main contribution, *order preserving hierarchical agglomerative clustering of strictly partially ordered sets*, is the combination of the two. We define an optimal order preserving hierarchical clustering to be the hierarchical clustering with the partial dendrogram that has the best ultrametric fit relative the original dissimilarity measure.

In want of an efficient algorithm, we present a method of approximation that can be computed in polynomial time. We demonstrate the approximation on synthetic data generated as random directed acyclic graphs and random dissimilarity measures, as well as on data from the parts database motivating this research. We evaluate the quality of the obtained clustering by computing the adjusted Rand index relative a planted partition (Hubert & Arabie, 1985). We provide a novel method for comparing two induced order relations using a modified adjusted Rand index, which we believe is a first of its kind. We also provide simple method for computing the level of order preservation of a clustering of an ordered set by counting the number of induced loops.

Beyond our main contribution, we believe that the embedding of partial dendrograms into ultrametrics may be of interest to a larger audience. The embedding provides a means for treating partial dendrograms as complete dendrograms, offering access to the entire rack of tools that already exists in this domain. An obvious example candidate is that of hierarchical clustering with must-link and no-link constraints. The no-link constraints will necessarily lead to partial dendrograms that can be easily evaluated in our framework.

1.3.1 Summary of contributions

Our main contribution is the theory for order preserving hierarchical agglomerative clustering for strict posets. Further contributions we wish to highlight are:

- A theory for embedding partial dendrograms over a set into the set of complete dendrograms over the same set.
- An optimisation based, permutation invariant hierarchical clustering methodology for non-ordered sets that is very similar to classical hierarchical clustering.
- A polynomial time approximation scheme for order preserving hierarchical agglomerative clustering
- A novel method for comparison of induced order relations over a set based on the adjusted Rand index.
- A measure of the level of order preservation of a clustering of an ordered set.

1.4 Related work

Hierarchical agglomerative clustering is described in a plethora of books and articles, and we shall not try to give an account of that material. For an introduction to the subject, see (Jain & Dubes, 1988, §3.2).

1.4.1 Clustering of ordered data

There are quite a few articles presenting clustering of ordered data, placing themselves in one of two categories.

The first is clustering of sets where the (dis)similarity measure is replaced by information about whether one pair of elements is more similar than another pair of elements, for example based on user preferences. This is sometimes referred to as *comparison based clustering*. See the recent article by Ghoshdastidar et al. (2019) for an example and references. In this category, we also find the works of Janowitz (2010), providing a wholly order theoretic description of hierarchical clustering, including the case where the dissimilarity measure is replaced by a partially ordered set.

The second variant is to partition a family of ordered sets so that similarly ordered sets are associated with each other. Examples include the paper by Kamishima and Fujiki (2003), where they develop a variation of k -means, called k - o 'means, for clustering preference data, each list of preferences being a totally ordered set. Other examples in this category include clustering of times series, identifying which times series are alike (Łuczak, 2016).

Our method differs from all of the above in that we cluster elements inside one ordered set through the use of a (dis)similarity measure, while maintaining the original orders of elements.

1.4.2 Clustering to detect order

Another variant is the detection of order relations in data through clustering: In Carlsson et al. (2014), it is demonstrated how hierarchical agglomerative quasi-clustering can be used to deduce a partial order of “net flow” from an asymmetric network.

In this category, it is also worth mentioning dynamic time warping. This is a method for aligning time series, and can be considered as clustering across two time series that is indeed order preserving. See Łuczak (2016) for further references on this.

1.4.3 Acyclic graph partitioning problems

The problem of order preserving hierarchical agglomerative clustering can be said to belong to the family of *acyclic graph partitioning problems* (Herrmann et al., 2017). If we consider the strict partial order to be a directed acyclic graph (DAG), the task is to partition the vertices into groups so that the groups together with the arrows still makes up a DAG.

Graph partitioning has received a substantial attention from researchers, especially within computer science, over the last 50 years. Two important fields of application of this theory are VLSI and parallel execution.

In VLSI, short for Very Large Scale Integration, the problem can be formulated as follows: Given a set of micro processors, the wires that connect them, and a set of circuit boards, how do you best place the processors on the circuit boards in order to optimise a given objective function? Typically, a part of the objective function is to minimise the wire length. But other features may also be part of the optimisation, such as the amount or volume of traffic between certain processors etc. (Markov et al., 2015)

For parallel processing, the input data is a set of tasks to be executed. The tasks are organised as a DAG, where predecessors must be executed before descendants. Given a finite number of processors, the problem is to group the tasks so that they can be run group-wise on a processor, or running groups in parallel on different processors, in order to execute all tasks as quickly as possible. Typically additional information available is memory requirements, expected execution times for the tasks, etc. (Buluç et al., 2016)

It is not difficult to understand why both areas have received attention, being essential in the development of modern computers. The development of theory and methods has been both successful and abundant, and a large array of techniques are available, both academic and commercially.

Although both problems do indeed perform clustering of strict partial orders, their solutions are not directly transferable to exploratory data analysis. Mostly because they have very specific constraints and objectives originating from their respective problem domains.

The method we propose in this paper has as input a strict partial order (equivalently; a DAG) together with an arbitrary dissimilarity measure. We then use the classical linkage functions single-, average-, and complete linkage to suggest clusterings of the vertices from the input dataset, while preserving the original order relation.

Our method therefore places itself firmly in the family of acyclic graph partitioning methodologies, but with different motivation, objective and solution, compared to existing methods.

1.4.4 Hierarchical clustering as an optimisation problem

Several publications aim at solving hierarchical clustering in terms of optimisation. However, due to the procedural nature of classical hierarchical clustering, combined with the linkage functions, pinning down an objective function may be an impossible task. Especially since classical hierarchical clustering is not even well defined for complete linkage in the presence of tied connections. This leads to a general abandonment of linkage functions in optimisation based hierarchical clustering.

Quite commonly, optimisation based hierarchical clustering is done in terms of *ultrametric fitting*. That is, it aims to find an ultrametric that is as close to the original dissimilarity measure as possible, perhaps adding some additional constraints (Gilpin et al., 2013; Chierchia & Perret, 2019). It is well known that solving single linkage hierarchical clustering is equivalent to finding the so called *maximal sub-dominant ultrametric*. That is; the ultrametric that is pointwise maximal among all ultrametries not exceeding the original dissimilarity (Rammal et al., 1986). But for the other linkage functions, there is no equivalent result.

Optimisation based hierarchical clustering therefore generally present alternative definitions of hierarchical clustering. Quite often based on objective functions that originate from some particular domain. Exceptions from this are, for example, Ward's method (Ward, 1963), where the topology of the clusters are the focus of the objective, and also the recent addition by Dasgupta (2016), where the optimisation aims towards topological properties of the generated dendrogram.

Although our method is, eventually, based on ultrametric fitting, we optimise over a very particular set of dendrograms. Namely the dendrograms that can be generated through classical hierarchical clustering with linkage functions. It is therefore reasonable to claim that our method places itself between classical hierarchical clustering and optimised models.

1.4.5 Clustering with constraints

A significant amount of research has been devoted to the topic of clustering with constraints in the form of pairwise *must-link* or *no-link* constraints, often in addition to other constraints, such as minimal- and maximal distance constraints, and so on. Some work as also been done on hierarchical agglomerative clustering with constraints, starting with the works of Davidson and Ravi (2005). For a thorough treatment of constrained clustering, see Basu et al. (2008).

Order preserving clustering (as well as acyclic partitioning) can be seen as a particular version of constrained clustering, where the constraint is a *directed, transitive cannot-link* constraint. A type of constraint that is not found in the constrained clustering literature.

1.4.6 Clustering in information networks

A large amount of research has been conducted on the problem of clustering nodes in networks, and a more recent field of research is that of clustering data organised in *heterogeneous information networks*, or HINs for short (Pio et al., 2018). A HIN is an undirected graph where both vertices and edges may have different, or even multiple, types. RDF graphs (Lassila & Swick, 1999) is but one example of HINs. In a sense, we can say that the availability of multiple types allow HINs to model the real world more closely, but with the penalty of increased complexity. It is fair to consider HIN clustering a generalisation of classical network clustering, where, in the classical setting, all vertices and edges are of one common type.

However, the general case in clustering both classical networks and HINs is that although the network structure serves to influence the clustering, the structure is usually lost in the clustering. The most classical example is where connectedness between vertices contribute to vertex similarity, and then the most connected vertices (clique-like subgraphs) are clustered together. Although this can be seen as a type of relation preserving clustering, in order preserving clustering, the opposite is taking place: the more connected two vertices are, the more reason *not* to place them in the same cluster. Indeed, as we show in Sect. 4, for the theory we present in this paper, two elements can only be clustered together if there are *no* paths connecting them.

An example of HIN clustering that *is* structure preserving is Li et al. (2017). A HIN comes with a schema, or a *schematic graph*, describing which types are related to which other types. For Li et al. (2017), the goal is to cluster each set of same-type nodes according to a discovered similarity measure. The result is thus a schematic graph where each node is a clustering of vertices of the same type. This differs from the problem we study in that we do not know which elements are of the same type; to discover this is the goal of the clustering. Hence, the problems are similar but different; we could rephrase our problem as that of *deriving* a directed schematic graph from unlabeled vertices, where each vertex in the schematic graph is a set of equivalent machine parts, and the directed edges are the part-of relations.

1.5 Organisation of the remainder of this paper

Section 2 provides necessary background material.

In Sect. 3, we develop *optimised hierarchical agglomerative clustering* for non-ordered sets; our permutation invariant clustering model that is tailored especially to fit into our framework for agglomerative clustering of ordered sets.

In Sect. 4, we tackle the problem of order preservation during clustering: We define what we mean by order preservation, and classify exactly the clusterings that are order preserving. We also provide concise necessary and sufficient conditions for an hierarchical agglomerative clustering algorithm to be order preserving.

Section 5 defines partial dendrograms and develops the embedding of partial dendrograms over an ordered set into the family of ultrametrics over the same set.

Our main result, *order preserving hierarchical agglomerative clustering for strict partial orders*, is presented Sect. 6.

Section 7 provides a polynomial time approximation scheme for our method, and Sect. 8 demonstrates the efficacy of the approximation on synthetic data.

Section 9 presents the results from applying our approximation method to a subset of the data in the parts database, comparing with existing methods, and finally, Sect. 10 closes the article with some concluding remarks, and a list of future work topics.

2 Background

In this section we recall basic background material. We start by recollecting the required order-theoretical tools together with equivalence relations, before recalling classical hierarchical clustering.

2.1 Relations

Definition 1 A **relation** R on a set X is a subset $R \subseteq X \times X$, and we say that x and y are **related** if $(x, y) \in R$. The short hand notation aRb is equivalent to writing $(a, b) \in R$.

2.1.1 Strict and non-strict partial orders

A **strict partial order** on a set X is a relation S on X that is irreflexive and transitive. Recall that, an irreflexive and transitive relation is also anti-symmetric. A **strictly partially ordered set**, or a **strict poset**, is a pair (X, S) , where X is a set and S is a strict partial order on X . We commonly denote a strict partial order by the symbol $<$.

On the other hand a **partial order** on X is a relation P on X that is reflexive, asymmetric and transitive, and the pair (X, P) is called a **partially ordered set**, or a **poset**. The usual notation for a partial order is $\leq$.

We shall just refer to strict and non-strict partial orders as *orders*, unless there is any need for disambiguation: If R is an order on X , we say that $a, b \in X$ are **comparable** if either $(a, b) \in R$ or $(b, a) \in R$. And, if every pair of elements in X are comparable, we call X **totally ordered**. A totally ordered subset of an ordered set is called a **chain**, and a subset where no two elements are comparable is called an **antichain**. We denote **non-comparability** by $a \perp b$. That is, for any elements a, b in an antichain, we have $a \perp b$.

A **cycle** in a relation E is a sequence in $\overline{E}$ on the form $(a, b_1), (b_1, b_2), \dots, (b_n, a)$. The **transitive closure** of E is the minimal set $\overline{E}$ for which the following holds: If there is a sequence of pairs $(a_1, a_2), (a_2, a_3), \dots, (a_{n-1}, a_n)$ in E , then $(a_1, a_n) \in \overline{E}$.

Let (X, E) be an ordered set. An element $x_0 \in X$ is a **minimal element** if there is no element $y \in X - \{x_0\}$ for which $(y, x_0) \in E$. Dually, y_0 is a **maximal element** if there is no $x \in X - \{y_0\}$ for which $(y_0, x) \in E$. If (X, E) has a unique minimal element, then this is called the **bottom element** or the **least element**, and a unique maximal element is called the **top element** or the **greatest element**.

Finally, a map $f : (X, <_X) \rightarrow (Y, <_Y)$ is **order preserving** if $a <_X b \Rightarrow f(a) <_Y f(b)$, and if f is a set isomorphism (that is, a bijection) for which f^{-1} is also order preserving, we say that f is an **order isomorphism**, and that the sets $(X, <_X)$ and $(Y, <_Y)$ are **order isomorphic**, writing $(X, <_X) \approx (Y, <_Y)$.

2.1.2 Partitions and equivalence relations

A **partition** of X is a collection of disjoint subsets of X , the union of which is X . The family of all partitions of X , denoted $\mathfrak{P}(X)$, has a natural partial order defined by

partition-refinement: If $\mathcal{A} = \{A_i\}_i$ and $\mathcal{B} = \{B_j\}_j$ are partitions of X , we say that $\mathcal{A}$ is a **refinement** of $\mathcal{B}$, writing $\mathcal{A} \in \mathcal{B}$, if, for every $A_i \in \mathcal{A}$ there exists a $B_j \in \mathcal{B}$ such that $A_i \subseteq B_j$. The sets of a partition are referred to as **blocks**.

An **equivalence relation** is a relation $\mathcal{R}$ on X that is reflexive, symmetric and transitive. Let the family of all equivalence relations over a set X be denoted by $\mathfrak{R}(X)$. If $\mathcal{R} \in \mathfrak{R}(X)$ and $(x, y) \in \mathcal{R}$, we say that x and y are **equivalent**, writing $x \sim y$. The maximal set of elements equivalent to $x \in X$ is called the **equivalence class of x** , and is denoted $[x]$. $\mathfrak{R}(X)$ is also partially ordered, but by subset inclusion: that is, for $\mathcal{R}, \mathcal{S} \in \mathfrak{R}(X)$, we say that $\mathcal{R}$ is less than or equal to $\mathcal{S}$ if and only if $\mathcal{R} \subseteq \mathcal{S}$.

The **quotient of X modulo $\mathcal{R}$** , denoted $X/\mathcal{R}$, is the set of equivalence classes of X under $\mathcal{R}$. Notice that $[x]$ is an element of $X/\mathcal{R}$, but a subset of X . Since the equivalence classes are subsets of X that together cover X , $X/\mathcal{R}$ is a partition of X with equivalence classes being the blocks of the partition. The family of partitions of X is in a one-to-one correspondence with the equivalence relations of X , and the correspondence is order preserving; if $\mathcal{A} = X/\mathcal{A}$ and $\mathcal{B} = X/\mathcal{B}$, we have

$$\mathcal{A} \in \mathcal{B} \Leftrightarrow \mathcal{A} \subseteq \mathcal{B}.$$

Both $\mathfrak{P}(X)$ and $\mathfrak{R}(X)$ have top- and bottom elements: The least element of $\mathfrak{P}(X)$ is the singleton partition $S(X)$, where each element is in a block by itself: $S(X) = \{\{x\} \mid x \in X\}$. The singleton partition corresponds to the diagonal equivalence relation, given by $\Delta(X) = \{(x, x) \mid x \in X\}$, which is the least element of $\mathfrak{R}(X)$. The greatest element of $\mathfrak{P}(X)$ is the trivial partition $\{X\}$, corresponding to the equivalence relation $X \times X$, where all elements are equivalent. That is

$$S(X) = X/\Delta(X) \quad \text{and} \quad \{X\} = X/(X \times X).$$

If $\mathcal{A}$ and $\mathcal{B}$ are partitions of X with $\mathcal{A}$ being a refinement of $\mathcal{B}$, we say that $\mathcal{A}$ is **finer** than $\mathcal{B}$, and that $\mathcal{B}$ is **coarser** than $\mathcal{A}$. We use the exact same terminology for the corresponding equivalence relations.

For a subset $A \subseteq X$, let the notation X/A denote the partition of X where all of A is one equivalence class, and the rest of X remains as singletons. Formally, this corresponds to the equivalence relation $\mathcal{R}_A = \Delta(X) \cup (A \times A)$. And finally, the **quotient map** corresponding to an equivalence relation $\mathcal{R} \in \mathfrak{R}(X)$ is the unique map $q_{\mathcal{R}} : X \rightarrow X/\mathcal{R}$ defined as $q_{\mathcal{R}}(x) = [x]$. That is, $q_{\mathcal{R}}$ sends each element to its equivalence class.

2.2 Classical hierarchical clustering

In this section, we recall classical hierarchical clustering in terms of Jardine and Sibson (1971). Our theory builds directly on the theory for classical hierarchical clustering, so we need to provide a fair bit of detail, especially since there is a general lack of standardised notation for hierarchical clustering theory.

We start by recalling the formal definition of a dendrogram, before recalling dissimilarity measures and ultrametrics. Thereafter, we recall linkage functions, before finally tying all the concepts together to define classical hierarchical agglomerative clustering.

Definition 2 A **clustering** of a set X is a partition of X , and a **hierarchical clustering** is a chain in $\mathfrak{P}(X)$ containing both the bottom and top elements. A **cluster** in a clustering is a block in the partition.

Alternatively, a clustering of X is an equivalence relation $\mathcal{R} \in \mathfrak{R}(X)$, and a hierarchical clustering is a chain in $\mathfrak{R}(X)$ containing both the bottom- and top elements of $\mathfrak{R}(X)$. A cluster is, then, an equivalence class in $X/\mathcal{R}$. We will refer to clusters as equivalence classes, clusters or blocks depending on the context, all terms being frequently used in clustering literature.

For the remainder of the paper, let $\mathbb{R}_+$ denote the non-negative reals. We generally assume that $\mathbb{R}_+$ is equipped with the usual total order $\leq$.

Definition 3 Given a set X , let $\mathfrak{P}(X)$ be partially ordered by partition refinement. A **dendrogram** is an order preserving map $\theta : \mathbb{R}_+ \rightarrow \mathfrak{P}(X)$ for which the following properties hold:

- D1. $\forall t \in \mathbb{R}_+ \exists \varepsilon > 0$ s.t. $\theta(t) = \theta(t + \varepsilon)$;
- D2. $\theta(0) = S(X)$, the least element of $\mathfrak{P}(X)$;
- D3. $\exists t_0 > 0$ s.t. $\theta(t_0) = \{X\}$, the greatest element of $\mathfrak{P}(X)$.

We will use the term dendrogram to denote both the graphical and the functional representation. If $in(\theta) = \{\mathcal{B}_i\}_{i=0}^n$, we assume that the enumeration is compatible with the order relation on $\mathfrak{P}(X)$; in other words, that $\{\mathcal{B}_i\}_{i=0}^n$ is a chain in $\mathfrak{P}(X)$. We denote **the family of all dendrograms over X** by $\mathcal{D}(X)$.

A **dissimilarity measure** on a set X is a function $d : X \times X \rightarrow \mathbb{R}_+$, satisfying

$$d1. \forall x \in X : d(x, x) = 0,$$

$$d2. \forall x, y \in X : d(x, y) = d(y, x).$$

If d additionally satisfies

$$d3. \forall x, y, z \in X : d(x, z) \leq \max\{d(x, y), d(y, z)\},$$

we call d an **ultrametric** (Rammal et al., 1986). The pair (X, d) is correspondingly called a **dissimilarity space** or an **ultrametric space**. The family of all dissimilarity measures over X is denoted by $\mathcal{M}(X)$, and the family of all ultrametrics by $\mathcal{U}(X)$.

Example 1 (Ultrametric) Property $d3$ is referred to as the **ultrametric inequality**, and is a strengthening of the usual triangle inequality. In an ultrametric space $(X, \mathbf{u})$, every triple of points is arranged in an isosceles triangle: Let $a, b, c \in X$, and let the pair a, b be of minimal distance such that $\mathbf{u}(a, b) \leq \min\{\mathbf{u}(a, c), \mathbf{u}(b, c)\}$. The ultrametric inequality gives us

$$\left. \begin{array}{l} \mathbf{u}(a, c) \leq \max\{\mathbf{u}(a, b), \mathbf{u}(b, c)\} = \mathbf{u}(b, c) \\ \mathbf{u}(b, c) \leq \max\{\mathbf{u}(b, a), \mathbf{u}(a, c)\} = \mathbf{u}(a, c) \end{array} \right\} \Leftrightarrow \mathbf{u}(a, c) = \mathbf{u}(b, c).$$

Ultrametrics show up in many different contexts, such as p -Adic number theory (Holly, 2001), infinite trees (Hughes, 2004), numerical taxonomy (Sneath & Sokal, 1973) and also within physics (Rammal et al., 1986), just to cite a few. For hierarchical clustering, ultrametrics are relevant because the dendrograms over a set are in a bijective relation to the ultrametrics over the same set (Carlsson & Mémoli, 2010).

We shall also need the following terms, which apply to any dissimilarity space: The **diameter** of (X, d) is given by the maximal inter-point distance:

$$\text{diam}(X, d) = \max\{d(x, y) \mid x, y \in X\}.$$

And the **separation** of (X, d) is the minimal inter point distance:

$$\text{sep}(X, d) = \min\{d(x, y) \mid x, y \in X \wedge x \neq y\}.$$

It is a well known fact that there exists an injective map from dendrograms to ultrametrics (Jardine & Sibson, 1971):

$$\Psi_X : \mathcal{D}(X) \longrightarrow \mathcal{U}(X).$$

In Carlsson and Mémoli (2010) the map Ψ_X is shown to be a bijection. If $\theta \in \mathcal{D}(X)$, the map is defined as

$$\Psi_X(\theta)(x, y) = \min\{t \in \mathbb{R}_+ \mid \exists B \in \theta(t) : x, y \in B\}. \quad (1)$$

That is, the ultrametric distance is the least real number t for which θ maps to a partition where x and y are in the same block. The minimisation is well defined due to Axiom D1. The ultrametric can be read from the diagrammatic representation of the dendrogram as the minimum height you have to ascend to in order to traverse from one element to the other following the paths in the tree.

Before we provide a formal definition of classical hierarchical clustering, we need to recall linkage functions. Our definition follows the lines of Carlsson and Mémoli (2010):

Definition 4 Let $\mathcal{P}(X)$ denote the power set of X . A **linkage function** on X is a map

$$\mathcal{L} : \mathcal{P}(X) \times \mathcal{P}(X) \times \mathcal{M}(X) \longrightarrow \mathbb{R}_+,$$

so that for each partition $Q \in \mathfrak{P}(X)$ and dissimilarity measure $d \in \mathcal{M}(X)$, the restriction $\mathcal{L}|_{Q \times Q \times \{d\}}$ is a dissimilarity measure on Q .

The classical linkage functions are defined as

$$\begin{aligned} \text{Single linkage} & : \mathcal{SL}(p, q, d) = \min_{x \in p} \min_{y \in q} d(x, y), \\ \text{Complete linkage} & : \mathcal{CL}(p, q, d) = \max_{x \in p} \max_{y \in q} d(x, y), \\ \text{Average linkage} & : \mathcal{AL}(p, q, d) = \frac{\sum_{x \in p} \sum_{y \in q} d(x, y)}{|p| \cdot |q|}. \end{aligned}$$

Definition 5 (Classical $\mathcal{HC}$) Given a dissimilarity space (X, d) and a linkage function $\mathcal{L}$, if we follow the procedure outlined in Sect. 1.2, using $\mathcal{L}$ as the “notion of dissimilarity”, the result is a chain of partitions $\{Q_i\}_{i=1}^{|X|-1}$ together with the dissimilarities $\{\rho_i\}_{i=1}^{|X|-1}$ at which the partitions were formed. The sequence of pairs $\mathcal{Q} = \{(Q_i, \rho_i)\}_{i=1}^{|X|-1}$ corresponds uniquely to a dendrogram $\theta_{\mathcal{Q}}$ as follows:

$$\theta_{\mathcal{Q}}(x) = Q_{\max\{i \in \mathbb{N} \mid \rho_i \leq x\}}. \quad (2)$$

We define a **classical hierarchical clustering of (X, d) using $\mathcal{L}$** to be a dendrogram

$$\mathcal{HC}^{\mathcal{L}}(X, d) = \theta_{\mathcal{Q}}$$

obtained through this procedure.

Remark 1 Notice that (2) maps $\{(Q_i, \rho_i)\}_{i=1}^{|X|-1}$ to a dendrogram if and only if

$$\text{sep}(Q_i, \mathcal{L}) \leq \text{sep}(Q_{i+1}, \mathcal{L}) \quad \text{for } 0 \leq i < |X| - 1. \quad (3)$$

Otherwise, the ρ_i will not make up a monotone sequence, and the resulting function θ_Q will not be an order preserving map. Although all of $\mathcal{SL}$, $\mathcal{AL}$ and $\mathcal{CL}$ satisfy (3), it is fully possible to define linkage functions that do not.

Finally, two distinct pairs of elements $(p_1, q_1), (p_2, q_2) \in Q \times Q$ for which

$$\mathcal{L}(p_1, q_1, d) = \mathcal{L}(p_2, q_2, d) = \text{sep}(Q, \mathcal{L}),$$

are referred to as **tied**, since they are both eligible candidates for the next merge.

3 Optimised hierarchical clustering

In this section, we devise a permutation invariant version of hierarchical clustering based on the classical definition. The key to permutation invariance is in dealing with tied connections. If we consider the procedure for hierarchical clustering outlined in Sect. 1.2, we can resolve tied connections by picking a random minimal dissimilarity pair. The way the procedure is specified, this turns $\mathcal{HC}^{\mathcal{L}}$ into a *non-deterministic* algorithm; it may produce different dendrograms for the same input in the presence of ties, depending on which tied pair is selected. But more importantly, it is capable of producing *any* dendrogram that can be produced by *any* tie resolution order:

Definition 6 Given a dissimilarity space (X, d) and a linkage function $\mathcal{L}$, let $\mathcal{D}^{\mathcal{L}}(X, d)$ be the set of all possible outputs from $\mathcal{HC}^{\mathcal{L}}(X, d)$.

A dissimilarity measure d over a finite set X can be described as an $|X| \times |X|$ real matrix $[d_{i,j}]$. Hence, given an ultrametric $\mathbf{u} \in \mathcal{U}(X)$ we can compute the pointwise difference

$$\|\mathbf{u} - d\|_p = \sqrt[p]{\sum_{x,y \in X} |\mathbf{u}(x,y) - d(x,y)|^p}. \quad (4)$$

We suggest the following definition, recalling the definition of Ψ_X (1):

Definition 7 Given a dissimilarity space (X, d) and a linkage function $\mathcal{L}$, the **optimised hierarchical agglomerative clustering over (X, d) using $\mathcal{L}$** is given by

$$\mathcal{HC}_{opt}^{\mathcal{L}}(X, d) = \arg \min_{\theta \in \mathcal{D}^{\mathcal{L}}(X, d)} \|\Psi_X(\theta) - d\|_p. \quad (5)$$

That is; among all dendrograms that can be generated by $\mathcal{HC}^{\mathcal{L}}(X, d)$, optimised hierarchical agglomerative clustering picks the dendrogram that is closest to the original dissimilarity measure. In the tradition of ultrametric fitting, this is the right choice of candidate.

As $\mathcal{D}^{\mathcal{L}}(X, d)$ contains all dendrograms generated over all possible permutations of enumerations of X , the below theorem follows directly from Definition 7:

Theorem 1 $\mathcal{HC}_{opt}^{\mathcal{L}}$ is permutation invariant. That is, the order of enumeration of the elements of the set X does not affect the output from $\mathcal{HC}_{opt}^{\mathcal{L}}(X, d)$.

And since $\mathcal{HC}^{SC}$ is permutation invariant, we have $|\mathcal{D}^{SC}(X, d)| = 1$, yielding

Theorem 2 $\mathcal{HC}_{opt}^{SC}(X, d) = \mathcal{HC}^{SC}(X, d)$.

Since $\mathcal{HC}^{AC}$ and $\mathcal{HC}^{CL}$ are not permutation invariant, there is no corresponding result in these cases. For complete linkage, however, we have the following theorem. First, notice that due to the definition of complete linkage (Definition 4), if θ is a solution to $\mathcal{HC}_{opt}^{CL}(X, d)$ and $\mathbf{u} = \Psi_X(\theta)$ is the corresponding ultrametric, then

$$\mathbf{u}(x, y) \geq d(x, y) \quad \forall x, y \in X.$$

Hence, in the case of complete linkage we can reformulate (5) as follows:

$$\mathcal{HC}_{opt}^{CL}(X, d) = \arg \min_{\theta \in \mathcal{D}^{CL}(X, d)} \|\Psi_X(\theta)\|_p. \quad (6)$$

To see why this is the case, notice that if $u, u' \in \mathcal{M}(X)$ and both $d \leq u$ and $d \leq u'$ pointwise, then we can produce two non-negative functions δ, δ' on $X \times X$ so that $u = d + \delta$ and $u' = d + \delta'$. In particular, we have $u - d = \delta$, from which we deduce

$$\|u - d\|_p \leq \|u' - d\|_p \Leftrightarrow \|\delta\|_p \leq \|\delta'\|_p \Leftrightarrow \|d + \delta\|_p \leq \|d + \delta'\|_p \Leftrightarrow \|u\|_p \leq \|u'\|_p.$$

Theorem 3 Solving $\mathcal{HC}_{opt}^{CL}(X, d)$ is NP-hard.

Proof Let $G = (V, E)$ be an undirected graph with vertices V and edges $E \subseteq V \times V$. Recall the *clique* problem: Given a positive integer $K < |V|$, is there a clique in G of size at least K ? Equivalently: is there a set $V' \subseteq V$ with $|V'| \geq K$ for which $V' \times V' \subseteq E$? This is a known NP-hard problem (Karp, 1972).

To reduce *clique* to $\mathcal{HC}_{opt}^{CL}$, define a dissimilarity measure on V as follows:

$$d(v, v') = \begin{cases} 1 & \text{if } (v, v') \in E, \\ 2 & \text{otherwise.} \end{cases} \quad (7)$$

Then (V, d) is a dissimilarity space. Let θ be a solution of $\mathcal{HC}_{opt}^{CL}(V, d)$, and set $\mathbf{d} = \Psi_V(\theta)$.

An intrinsic property of $\mathcal{CL}$ is that if two blocks $p, q \in \mathcal{Q}_i$ are merged, then

$$\forall v, v' \in p \cup q : d(v, v') \leq \mathcal{CL}(p, q, d).$$

And since we have $d(v, v') = 1 \Leftrightarrow (v, v') \in E$, it means that for a subset $V' \subseteq V$, we have that

$$\forall v, v' \in V' : \mathbf{d}(v, v') = 1 \Leftrightarrow V' \text{ is a clique in } G. \quad (8)$$

It follows that a largest possible cluster at proximity level 1 is a maximal clique in G .

We claim that minimising the norm is equivalent to producing a maximal cluster at proximity level 1: Let $\mathbf{d}$ be the $|V| \times |V|$ distance matrix $[\mathbf{d}_{ij}]$. Due to the definition of $\mathcal{CL}$, we have $\mathbf{d}(v, v') \in \{0, 1, 2\}$. If $\theta(1) = \{V_i\}_{i=1}^s$, then these are exactly the blocks that are subsets of cliques, so each V_i contributes with $|V_i|(|V_i| - 1)$ ones in $[\mathbf{d}_{ij}]$.

Having more ones reduces the norm of $\mathbf{d}$. Let V_j be of maximal cardinality in $\{V_i\}_{i=1}^s$. Assume first that V_j has at least two elements more than the next to largest block, and let $|V_j| = P$.

Removing one element from V_j reduces the number of ones in the dissimilarity matrix by $P(P-1) - (P-1)(P-2) = 2(P-1)$. Let the next to largest block have Q elements. Transferring the element to this block then increases the number of ones by $(Q+1)Q - Q(Q-1) = 2Q$. Since $Q < P-1$, this means that the total number of ones is reduced by moving an element from the largest block to any of the smaller blocks. Hence, achieving the largest possible number of ones implies maximising the size of the largest block.

If now, V_j only has one element more than the next to largest block, moving an element as above corresponds to keeping the number of ones. Since each V_i for $1 \leq i \leq s$ is a subset of a clique in G , the maximal number of ones is achieved by producing a block V_j that contains exactly a maximal clique of G .

Therefore, if $\mathcal{I}_{\{1\}}(x)$ is the indicator function for the set $\{1\}$, the size of a maximal clique in G can be computed as

$$\max_{1 \leq i \leq |V|} \left\{ \sum_{j=1}^{|V|} \mathcal{I}_{\{1\}}(\mathbf{d}_{ij}) \right\},$$

counting the maximal number of row-wise ones in $[\mathbf{d}_{ij}]$ in $O(N^2)$ time. We therefore conclude that $\mathcal{HC}_{opt}^{CL}$ is NP-hard. $\square$

The computational hardness of $\mathcal{HC}_{opt}^{CL}$ is directly connected to the presence of tied connections: every encounter of $\mathbf{n}$ tied connections leads to $\mathbf{n}!$ new candidate solutions.

Since neither $\mathcal{HC}_{opt}^{AL}$ is permutation invariant, the authors strongly believe that this is also NP-hard, although that remains to be proven.

We cannot in general expect the mapping $\theta \mapsto \|\Psi_X(\theta) - d\|_p$ to be injective, meaning that the answer to (5) may not be unique. Recall that $\mathcal{P}(X)$ denotes the power set of X . We shall consider $\mathcal{HC}_{opt}^L(X, -)$ to be the function

$$\mathcal{HC}_{opt}^L(X, -) : \mathcal{M}(X) \longrightarrow \mathcal{P}(\mathcal{D}(X)),$$

mapping a dissimilarity measure over X to a set of dendrograms over X .

3.1 Other permutation invariant solutions

Carlsson and Mémoli (2010) offer an alternative approach to permutation invariant hierarchical agglomerative clustering. In their solution, when they face a set of tied connections, they merge all tied the pairs in one operation, resulting in permutation invariance.

In the case of order preserving clustering, a family of tied connections can contain several mutually exclusive merges due to the order relation. Using the method of Carlsson and Mémoli leads to a problem of figuring which blocks of tied connections to merge together, and in which combinations and order. This leads to a combinatorial

explosion of alternatives. The method we have suggested is utterly simple, but it is designed to circumvent this very problem.

4 Order preserving clustering

In this section, we determine what it means for an equivalence relation to be order preserving with regards to a strict partial order, and establish precise conditions that are necessary and sufficient for a hierarchical agglomerative clustering algorithm to be order preserving.

4.1 Order preserving equivalence relations

Recalling the definition of a clustering (Definition 2), let $(X, <)$ be a strict poset. If $\mathcal{R}$ is an equivalence relation on X with quotient map $q : X \rightarrow X/\mathcal{R}$, we have already established, in Sect. 1.1, that we require

$$\forall x, y \in X : x < y \Rightarrow q(x) <' q(y).$$

That is, we are looking for a particular class of equivalence relations; namely those for which the quotient map is order preserving.

Given a strict poset (X, E) , there is a particular induced relation on the quotient set $X/\mathcal{R}$ for any equivalence relation $\mathcal{R} \in \mathfrak{R}(X)$ (Blyth, 2005, §3.1):

Definition 8 Given a strict poset (X, E) and an equivalence relation $\mathcal{R} \in \mathfrak{R}(X)$, first define the relation S_0 on X by

$$([a], [b]) \in S_0 \Leftrightarrow \exists x, y \in X : a \sim x \wedge b \sim y \wedge (x, y) \in E. \quad (9)$$

The transitive closure of S_0 is called **the relation on $X/\mathcal{R}$ induced by E** . We denote this relation by S .

Example 2 An instructive illustration of what the relation S_0 looks like for a strict poset $(X, <)$ under the equivalence relation $\mathcal{R}$ is that of an $\mathcal{R}$ -**fence** (Blyth, 2005), or just fence, for short:

Triple lines represent equivalences under $\mathcal{R}$, and the arrows represent the order on $(X, <)$. The fence illustrates visually how one can traverse from a_1 to b_n along arrows and

through equivalence classes in $X/\mathcal{R}$, and in that case we say that the fence **links** b_1 to a_n . **The induced relation S has the property that $(a, b) \in S$ if there exists an $\mathcal{R}$ -fence in X linking a to b .**

Recall that a cycle in a relation R is a sequence of pairs starting and ending with the same element: $(a, b_1), (b_1, b_2), \dots, (b_n, a)$. The below theorem is an adaptation of Blyth (2005, Thm.3.1) to strict partial orders.

Theorem 4 *Let (X, E) be a strict poset, $\mathcal{R} \in \mathfrak{R}(X)$, and let S be the relation on $X/\mathcal{R}$ induced by E . Then the following statements are equivalent:*

1. S is a strict partial order on $X/\mathcal{R}$;
2. There are no cycles in S_0 ;
3. $q_{\mathcal{R}} : (X, E) \longrightarrow (X/\mathcal{R}, S)$ is order preserving.

Proof From the definition of strict posets, they contain no cycles, so $1 \Rightarrow 2$. Since a non-cyclic set is irreflexive, and since S is transitive by construction, $2 \Rightarrow 1$.

Let $q_{\mathcal{R}}$ be order preserving. Notice that if S_0 is the set defined in (9), we have $S_0 = q_{\mathcal{R}} \times q_{\mathcal{R}}(E)$. In particular, for all $x, y \in X$ for which $(x, y) \in E$, we have $([x], [y]) \in S_0$. Assume that S is not a strict order. Then there is a cycle in S_0 ; that is there are $x, y \in X$ for which $(x, y) \in E$, but $([y], [x]) \in S_0$ also. This yields

$$\exists a', b' \in X : a' \sim x \wedge b' \sim y \wedge (b', a') \in E.$$

But, since $([x], [y]) \in S_0$, we also have

$$\exists a, b \in X : a \sim x \wedge b \sim y \wedge (a, b) \in E.$$

This yields $a \sim a'$ and $b \sim b'$, so we have

$$(q_{\mathcal{R}}(a), q_{\mathcal{R}}(b)) \in S_0 \wedge q_{\mathcal{R}}(b) = q_{\mathcal{R}}(b') \wedge (q_{\mathcal{R}}(b'), q_{\mathcal{R}}(a')) \in S_0.$$

But, since we have both $q_{\mathcal{R}}(a) = q_{\mathcal{R}}(a')$ and $(a, b) \in E$, this contradicts the fact that $q_{\mathcal{R}}$ is order preserving, so our assumption that both $([x], [y])$ and $([y], [x])$ are elements of S_0 must be wrong. Hence, if $q_{\mathcal{R}}$ is order preserving, there are no cycles in S_0 , and S is a strict partial order on $X/\mathcal{R}$. This shows that $3 \Rightarrow 1$.

Finally, let S be a strict partial order, and assume that $q_{\mathcal{R}}$ is not order preserving. Then, there exists $x, y \in X$ where $(x, y) \in E$ and for which at least one of $([x], [y]) \notin S$ or $([y], [x]) \in S$ holds. Now, $([x], [y]) \in S$ by Definition 8. Therefore, $([y], [x]) \in S$ implies that S has a cycle, contradicting the fact that S is a strict partial order. $\square$

Definition 9 Let (X, E) be a strict poset. An equivalence relation $\mathcal{R} \in \mathfrak{R}(X)$ is **regular** if there exists an order on $X/\mathcal{R}$ for which the quotient map is order preserving.

We denote **the set of all regular equivalence relations** over an ordered set $(X, <)$ by $\mathfrak{R}(X, <)$. Likewise, the family of all **regular partitions** of $(X, <)$ is denoted $\mathfrak{P}(X, <)$.

In general, we will denote the induced order relation for a strict poset $(X, <)$ and a regular equivalence relation $\mathcal{R} \in \mathfrak{R}(X, <)$ by $<'$.

4.2 The structure of regular equivalence relations

We now establish a sufficient and necessary condition for an agglomerative clustering algorithm to be order preserving. Recall that, if $A \subseteq X$, X/A denotes the quotient for which the quotient map $q_A : X \rightarrow X/A$ sends all of A to a point, and is the identity otherwise. That is, for every $x, y \in X$, we have

$$q_A(x) = q_A(y) \Leftrightarrow x, y \in A.$$

Theorem 5 *If $A \subseteq X$ for a strict poset $(X, <)$, the quotient map $q_A : X \rightarrow X/A$ is order preserving if and only if A is an antichain in $(X, <)$.*

Proof If A is not an antichain, then X/A places comparable elements in the same equivalence class, so q_A is not order preserving.

Assume A is an antichain. If q_A is not order preserving, then there is a cycle in $(X/A, <')$, and since we have only one non-singleton equivalence class, the cycle must be on the form

$$b \xrightarrow{\quad} A \xrightarrow{\quad} c.$$

But this means we have $a, a' \in A$ for which $b < a$ and $a' < c$, but since $c < b$, this implies $a' < a$, contradicting the fact that A is an antichain. $\square$

Since a composition of order preserving maps is order preserving, this also applies to a composition of quotient maps for a chain of regular equivalence relations $\mathcal{R}_1 \subseteq \dots \subseteq \mathcal{R}_n$. Combining this with Theorem 5, we have the following:

A clustering of a strict poset will be order preserving if it can be produced as a sequence of pairwise merges of non-comparable elements.

We close the section with an observation about the family of all hierarchical clusterings over a strict poset:

Theorem 6 *For a strict poset $(X, <)$, the set $\mathfrak{P}(X, <)$ of regular partitions over $(X, <)$ has $S(X)$ as its least element. Unless $<$ is the empty order, there is no greatest element.*

Proof $S(X)$ is always a regular partition, so $S(X) \in \mathfrak{P}(X, <)$. And since $S(X)$ is a refinement of every partition of X , $S(X)$ is the least element of $\mathfrak{P}(X, <)$.

If the order relation is not empty, then there are at least two elements that are comparable, and, according to Theorem 5, they cannot be in the same equivalence class. Hence, there is no greatest element. $\square$

The situation of Theorem 6 is depicted in Fig. 2, and has already been discussed in Sect. 1.2: In the case of tied connections that represent mutually exclusive merges, choosing to merge one connection over the other may lead to very different results. We therefore need a strategy to select one of these solutions over the others. This will be the main focus of Sects. 5 and 6.

5 Partial dendrograms

In this section we formally define partial dendrograms, and then construct the embedding of partial dendrograms into ultrametrics.

Based on the discussion of partial dendrograms in Sect. 1.2 together with the definition of dendrograms (Definition 3), we suggest the following:

Definition 10 A **partial dendrogram over** $(X, <)$ is an order preserving map $\theta : \mathbb{R}_+ \rightarrow \mathfrak{P}(X, <)$ satisfying properties D1 and D2 of Definition 3.

The only difference between a dendrogram and a partial dendrogram is that for a partial dendrogram we do not require the existence of a greatest element in the image of θ . Partial dendrograms are clearly a generalisation of dendrograms. To distinguish between the two, we will occasionally refer to the non-partial dendrograms as **complete dendrograms**. We denote the family of partial dendrograms over $(X, <)$ by $\mathcal{PD}(X, <)$.

For a partial dendrogram θ , we will write $\theta(\infty)$ to denote the maximal partition in the image of θ . Since $\mathfrak{P}(X, <)$ is finite, a partial dendrogram $\theta \in \mathcal{PD}(X, <)$ is *eventually constant*; that is, there exists a positive real number t_0 for which

$$t \geq t_0 \Rightarrow \theta(t) = \theta(\infty).$$

We refer to this number as the **diameter** of θ ; formally,

$$\text{diam}(\theta) = \min\{x \in \mathbb{R}_+ \mid \theta(x) = \theta(\infty)\}.$$

We now turn to the task of constructing the embedding. Looking at the partial dendrograms of Fig. 3, each connected component in a partial dendrogram is a complete dendrogram over its leaf nodes. Since complete dendrograms map to ultrametrics, each connected component gives rise to an ultrametric on the subset of X constituted by the connected component's leaf nodes. That is, if $\theta(\infty) = \{B_j\}_{j=1}^k$, and if θ_j is the complete dendrogram over B_j for $1 \leq j \leq k$, we can define the ultrametrics $\mathbf{u}_j = \Psi_{B_j}(\theta_j)$ so that $\{(B_j, \mathbf{u}_j)\}_{j=1}^k$ is a disjoint family of ultrametric spaces, which union covers X .

Now consider the following general result.

Lemma 1 Given a family of bounded, disjoint ultrametric spaces $\{(X_j, d_j)\}_{j=1}^n$ together with a positive real number $K \geq \max_j \{\text{diam}(X_j, d_j)\}$, the map

$$d_{\cup} : \bigcup X_j \times \bigcup X_j \longrightarrow \mathbb{R}_+$$

given by

$$d_{\cup}(x, y) = \begin{cases} d_j(x, y) & \text{if } \exists j : x, y \in X_j, \\ K & \text{otherwise} \end{cases}$$

is an ultrametric on $\bigcup_j X_j$.

Proof To prove that the ultrametric inequality holds, we start by showing that $d_{\cup_{1,2}}$ is an ultrametric on the restriction to the disjoint union $X_1 \cup X_2$: Let $x, y \in X_1$ and $z \in X_2$, and choose a positive $K \geq \max\{\text{diam}(X_1, d_1), \text{diam}(X_2, d_2)\}$. We now have

$$d_{U_{1,2}}(x, z) = K \quad d_{U_{1,2}}(x, y) = d_1(x, y) \quad d_{U_{1,2}}(y, z) = K.$$

This means that every triple of points are either already contained in an ultrametric space, or they make up an isosceles triangle. In both cases, the ultrametric inequality holds, according to the observation in Example 1.

By induction, we can now prove that $((X_1 \cup X_2) \cup X_3, d_{U_{1,2,3}})$ is an ultrametric space, and so on, until all the (X_j, d_j) are included. $\square$

Hence, for our partial dendrogram θ with $\theta(\infty) = \{B_j\}_{j=1}^k$ and subspace ultrametrics $\{u_j\}_{j=1}^k$, pick a $K \geq \max_j \{\text{diam}(B_j, u_j)\}$, and define $u_\theta : X \times X \rightarrow \mathbb{R}_+$ by

$$u_\theta(x, y) = \begin{cases} u_j(x, y) & \text{if } \exists j : x, y \in B_j, \\ K & \text{otherwise.} \end{cases} \quad (10)$$

According to Lemma 1, equation (10) is an ultrametric on X .

Definition 11 Given an ordered space $(X, <)$ and a non-negative real number ε , the **ultrametric completion on ε** is the map $\mathfrak{U}_\varepsilon : \mathcal{PD}(X, <) \rightarrow \mathcal{U}(X)$ mapping

$$\mathfrak{U}_\varepsilon : \theta \mapsto u_\theta,$$

where u_θ is defined as in (10), setting $K = \text{diam}(\theta) + \varepsilon$.

Example 3 To illustrate how the ultrametric completion turns out in the case of the partial dendrograms of Fig. 3, we have the following figure:

The above discussion serves to show that the construction is well defined. Our next goal is two-fold. First, we wish to provide an (explicit) function from partial dendrograms to dendrograms that realises this map. And second, we wish to establish conditions for this function to be an embedding; that is, an injective map. Injectivity is not strictly required for the theory to work, but it increases its discriminative power. An example to the contrary is provided towards the end of the section.

We have the map $\Psi_X : \mathcal{D}(X) \rightarrow \mathcal{U}(X)$ from (1), mapping dendrograms to ultrametrics. We now seek a map $\kappa_\varepsilon : \mathcal{PD}(X, <) \rightarrow \mathcal{D}(X)$ making the following diagram commute:

$$\begin{array}{ccc} \mathcal{D}(X) & \xrightarrow{\Psi_X} & \mathcal{U}(X) \\ \kappa_\varepsilon \uparrow & \nearrow \mathfrak{U}_\varepsilon & \\ \mathcal{PD}(X, <) & & \end{array} \quad (11)$$

Seeing that κ_ε must map partial dendrograms to complete dendrograms, a quick glance at Fig. 4 suggests the following:

$$\kappa_\varepsilon(\theta)(x) = \begin{cases} \theta(x) & \text{for } x < \text{diam}(\theta) + \varepsilon \\ \{X\} & \text{otherwise.} \end{cases}$$

It is straightforward to check that $\kappa_\varepsilon(\theta)$ is a complete dendrogram.

Fig. 4 “Completed” dendrograms corresponding to the partial dendrograms of Fig. 3, using $\varepsilon = 0.5$. The completion is marked by the dashed lines

Theorem 7 $\Psi_X \circ \kappa_\varepsilon = \mathcal{U}_\varepsilon$. That is; diagram (11) commutes.

Proof Assume first that $\theta \in \mathcal{PD}(X, <)$ is a proper partial dendrogram, and that $\text{im}(\theta) = \{\mathcal{B}_i\}_{i=0}^n$. Let the coarsest partition in the image of θ be given by $\mathcal{B}_n = \{B_j\}_{j=1}^m$. That is, each block B_j corresponds to a connected component in the partial dendrogram. Pick a block $B \in \mathcal{B}_n$ and assume $x, y \in B$.

If

$$k = \min\{i \in \mathbb{N} \mid \exists B' \in \mathcal{B}_i : B = B'\},$$

then $\mathcal{B}_k$ is the finest partition containing all of B in one block. Since $B \subseteq X$, the partitions

$$\mathcal{B}_i^B = \{B \cap B' \mid B' \in \mathcal{B}_i\} \quad \text{for } 1 \leq i \leq k$$

constitute a chain in $\mathfrak{A}(B)$ containing both $S(B)$ and $\{B\}$. Hence, we can construct a complete dendrogram over B by defining

$$\theta_B(x) = \{B \cap B' \mid B' \in \theta(x)\}. \quad (12)$$

This is exactly the complete dendrogram corresponding to the connected component of the tree over X having the elements of B as leaf nodes. By Definition 11,

$$x, y \in B \Rightarrow \mathcal{U}_\varepsilon(\theta)(x, y) = \Psi_B(\theta_B)(x, y). \quad (13)$$

Due to (12), we have

$$\begin{aligned} x, y \in B &\Rightarrow (\exists B \in \theta_B(x) : x, y \in B \Leftrightarrow \exists B' \in \theta(x) : x, y \in B') \\ &\Rightarrow \min\{t \in \mathbb{R}_+ \mid \exists B \in \theta_B(t) : x, y \in B\} \\ &= \min\{t \in \mathbb{R}_+ \mid \exists B' \in \theta(t) : x, y \in B'\}. \end{aligned}$$

Hence, by the definition of Ψ_X in (1) we conclude that

$$x, y \in B \Rightarrow \Psi_B(\theta_B)(x, y) = (\Psi_X \circ \kappa_\varepsilon)(\theta)(x, y).$$

Combining this with (13), we get that whenever $x, y \in B$, we have $\Psi_X \circ \kappa_\varepsilon = \mathcal{U}_\varepsilon$.

On the other side, let $x \in B_i$ and $y \in B_j$ with $i \neq j$. By definition, we have $\mathcal{U}_\varepsilon(\theta)(x, y) = \text{diam}(\theta) + \varepsilon$. And, since there is no block in $\theta(\infty)$ containing both x and y , we find that the minimal partition in $\text{im}(\kappa_\varepsilon(\theta))$ containing x and y in one block is $\{X\}$. But this means that $\Psi_X(\kappa_\varepsilon(\theta))(x, y) = \text{diam}(\theta) + \varepsilon$, so $\Psi_X \circ \kappa_\varepsilon = \mathcal{U}_\varepsilon$ holds in this case too.

Finally, if θ is a complete dendrogram, we have $\kappa_\varepsilon(\theta) = \theta$, so $\Psi_X \circ \kappa_\varepsilon(\theta) = \Psi_X(\theta)$. But since $\theta(\infty) = \{X\}$, it follows that $\mathfrak{U}_\varepsilon(\theta)$ maps exactly to the ultrametric over X defined by $\Psi_X(\theta)$. $\square$

Theorem 8 *Let $(X, <)$ be a strict poset with a non-empty order relation. Then $\mathfrak{U}_\varepsilon$ is injective if $\varepsilon > 0$.*

Proof Injectivity follows if κ_ε is injective, so assume that $\kappa_\varepsilon(\theta) = \kappa_\varepsilon(\theta')$. Then, for every $x < \text{diam}(\theta) + \varepsilon$, we have $\kappa_\varepsilon(\theta)(x) = \kappa_\varepsilon(\theta')(x) \Leftrightarrow \theta(x) = \theta'(x)$. $\square$

Example 4 If ε is not chosen to be strictly positive, the map $\mathfrak{U}_\varepsilon$ will not necessarily be injective. Consider the below dendrograms.

Both of the partial dendrograms are mapped to the same complete dendrogram (on the right) for $\varepsilon = 0$. This illustrates what we mean by *reduced discriminative power* in the case of a non-injective completion. Since the partial dendrograms exhibit distinctively different information, it is desirable that the methodology can distinguish them.

6 Hierarchical clustering of ordered sets

We are now ready to embark on the specification of order preserving hierarchical clustering of ordered sets. We do this by extending our notion of optimised hierarchical clustering from Sect. 3. For the remainder of the paper, let an **ordered dissimilarity space** be denoted by $(X, <, d)$.

Consider the following modification of classical hierarchical clustering. The only difference is that for each iteration, we check that there are elements that actually can be merged while preserving the order relation. According to Theorem 5, this means merging a pair of non-comparable elements at each iteration. Recall that $S(X)$ denotes the singleton partition of X .

Let $(X, <, d)$ be given together with a linkage function $\mathcal{L}$.

1. Set $\mathcal{Q}_0 = S(X)$, and endow $\mathcal{Q}_0$ with the induced order relation $<_0$.
2. Among the pairs of non-comparable clusters, pick a pair of minimal dissimilarity according to $\mathcal{L}$, and combine them into one cluster by taking their union.
3. Endow the new clustering with the induced order relation.
4. If all elements of X are in the same cluster, or if all clusters are comparable, we are done. Otherwise, go to Step 2 and continue.

The procedure results in a chain of ordered partitions $\{(Q_i, <_i)\}_{i=0}^m$ together with the dissimilarities $\{\rho_i\}_{i=0}^m$ at which the partitions were formed. For an ordered set $(X, <)$, recall that non-comparability of $a, b \in X$ is denoted $a \perp b$. Let the **non-comparable separation of $(X, <, d)$** , be given by

$$\text{sep}_\perp(X, <, d) = \min_{x, y \in X} \{d(x, y) \mid x \neq y \wedge x \perp y\}.$$

The reader may wish to compare the following lemma to Remark 1.

Lemma 2 *The sequence of pairs $\{(Q_i, \rho_i)\}_{i=0}^m$ produced by the above procedure maps to a partial dendrogram through application of (2) if and only if*

$$\text{sep}_\perp(Q_i, <_i, \mathcal{L}) \leq \text{sep}_\perp(Q_{i+1}, <_{i+1}, \mathcal{L}).$$

Since the singleton partition Q_0 maps to a partial dendrogram, the algorithm will produce a partial dendrogram for any ordered dissimilarity space, and since there can be at most $|X| - 1$ merges, the procedure always terminates.

As for classical hierarchical clustering, the procedure is non-deterministic in the sense that given a set of tied pairs, we may pick a random pair for the next merge. Hence, the procedure is capable of producing partial dendrograms for all possible tie resolution strategies:

Definition 12 Given an ordered dissimilarity space $(X, <, d)$ and a linkage function $\mathcal{L}$, we write $\mathcal{D}^\mathcal{L}(X, <, d)$ to denote the set of all possible outputs from the above procedure

The set $\mathcal{D}^\mathcal{L}(X, <, d)$ differs from $\mathcal{D}^\mathcal{L}(X, d)$ in two important ways:

- $\mathcal{D}^\mathcal{L}(X, <, d)$ contains partial dendrograms, not dendrograms.
- The cardinality of $\mathcal{D}^\mathcal{L}(X, <, d)$ is at least that of $\mathcal{D}^\mathcal{L}(X, d)$, and often higher, due to mutually exclusive merges and the “dead ends” in $\mathfrak{A}(X, <)$ (see Fig. 2).

Even for single linkage we have $|\mathcal{D}^{\text{SL}}(X, <, d)| > 1$ if there are mutually exclusive tied connections.

In the spirit of optimised hierarchical clustering, we suggest the following definition, employing the ultrametric completion $\mathfrak{U}_\epsilon$ from Definition 11:

Definition 13 Given an ordered dissimilarity space $(X, <, d)$ together with a linkage function $\mathcal{L}$, let $\epsilon > 0$. An **order preserving hierarchical agglomerative clustering using $\mathcal{L}$ and ϵ** is given by

$$\mathcal{HC}_{\text{opt}, \epsilon}^{<\mathcal{L}}(X, <, d) = \arg \min_{\theta \in \mathcal{D}^\mathcal{L}(X, <, d)} \|\mathfrak{U}_\epsilon(\theta) - d\|_p. \quad (14)$$

The next theorem shows that if we remove the order relation, then optimised clustering and order preserving clustering coincide. Keep in mind that a dissimilarity space is an ordered dissimilarity space with an empty order relation; that is, $(X, d) = (X, \emptyset, d)$.

Theorem 9 *If the order relation is empty, then order preserving optimised hierarchical clustering and optimised hierarchical clustering coincide:*

$$\mathcal{HC}_{opt,\varepsilon}^{<\mathcal{L}}(X, \emptyset, d) = \mathcal{HC}_{opt}^{\mathcal{L}}(X, d).$$

Proof First, notice that

$$\forall (Q, <_Q) \in \mathfrak{P}(X, \emptyset) : \text{sep}_\perp(Q, <_Q, \mathcal{L}) = \text{sep}(Q, \mathcal{L}),$$

where $<_Q$ denotes the (trivial) induced order. Hence, we have $\mathcal{D}^{\mathcal{L}}(X, \emptyset, d) = \mathcal{D}^{\mathcal{L}}(X, d)$. Since $\mathfrak{U}_\varepsilon|_{\mathcal{D}(X)} = \Psi_X$, the result follows. $\square$

6.1 On the choice of ε

In $\mathcal{HC}_{opt,\varepsilon}^{<\mathcal{L}}(X, <, d)$ we identify the elements from $\mathcal{D}^{\mathcal{L}}(X, <, d)$ that are closest to the dissimilarity measure d when measured in the p -norm. The injectivity of $\mathfrak{U}_\varepsilon$ induces a relation $\preceq_{d,\varepsilon}$ on $\mathcal{PD}(X, <)$ defined by

$$\theta \preceq_{d,\varepsilon} \theta' \Leftrightarrow \|\mathfrak{U}_\varepsilon(\theta) - d\|_p \leq \|\mathfrak{U}_\varepsilon(\theta') - d\|_p,$$

and the optimisation finds the minimal elements under this order.

The choice of ε may affect the ordering of dendrograms under $\preceq_{d,\varepsilon}$. We show this by providing an alternative formula for $\|\mathbf{u} - d\|_p$ that better expresses the effect of the choice of ε . Assume θ is a partial dendrogram over $(X, <)$ with $\theta(\infty) = \{B_i\}_{i=1}^m$, and let $\mathfrak{U}_\varepsilon(\theta) = \mathbf{u}$. We split the sum for computing $\|\mathbf{u} - d\|_p$ in two: the intra-block differences and the inter-block differences. The **intra-block differences** are independent of ε , and are given by

$$\alpha = \sum_{i=1}^m \sum_{x,y \in B_i} |\mathbf{u}(x,y) - d(x,y)|^p. \quad (15)$$

On the other hand, the **inter-block differences** are dependent on ε , and can be computed as

$$\beta_\varepsilon = \sum_{\substack{(x,y) \in B_i \times B_j \\ i \neq j}} |\text{diam}(\theta) + \varepsilon - d(x,y)|^p. \quad (16)$$

This yields $\|\mathbf{u} - d\|_p = \sqrt[p]{\alpha + \beta_\varepsilon}$. If we think of $\mathbf{u}$ as an approximation of d , and saying that $|X| = N$, the mean p -th error of this approximation can be expressed as a function of ε :

$$E_d(\varepsilon|\theta, p) = \frac{1}{N} \|\mathbf{u} - d\|_p^p = \frac{\alpha}{N} + \frac{1}{N} \sum_{\substack{(x,y) \in B_i \times B_j \\ i \neq j}} |\text{diam}(\theta) + \varepsilon - d(x,y)|^p.$$

From the formula for $E_d(\varepsilon|\theta, p)$, we see that when ε becomes large, the inter-block differences dominate the approximation error. Thus, for increasing ε , having low error eventually equals having few inter-block pairs. As a result, large ε will lead to clusterings where as many elements as possible are placed in one block, since this is the most effective method for reducing the number of inter-block pairs.

On the other hand, a low value of ε will move the weight towards optimising the intra-block ultrametric fit and move the bias away from large block sizes.

From the authors' perspective, focusing on block sizes seems to be less in the spirit of ultrametric fitting, compared to optimising the intra-block ultrametric fit. As such, it is

the authors' opinion that this points towards selecting a low value for ε . In the process of choosing, we have the following result at our aid:

Theorem 10 *For any finite ordered dissimilarity space $(X, <, d)$ and linkage function $\mathcal{L}$, there exists an $\varepsilon_0 > 0$ for which*

$$\varepsilon, \varepsilon' \in (0, \varepsilon_0) \Rightarrow (\mathcal{D}^{\mathcal{L}}(X, <, d), \leq_{d,\varepsilon}) \approx (\mathcal{D}^{\mathcal{L}}(X, <, d), \leq_{d,\varepsilon'}).$$

That is; all $\varepsilon \in (0, \varepsilon_0)$ induce the same order on the partial dendrograms.

Proof Since X is finite, $\mathcal{D}^{\mathcal{L}}(X, <, d)$ is also finite. And according to $E_d(\varepsilon|\theta, p)$, if the cardinality of $\mathcal{D}^{\mathcal{L}}(X, <, d)$ is n , there are at most pn positive values of ε that are distinct global minima of partial dendrograms in $\mathcal{D}^{\mathcal{L}}(X, <, d)$. But this means there is a finite set of ε for which the order on $(\mathcal{D}^{\mathcal{L}}(X, <), \leq_{\varepsilon,p})$ changes. And since all these values are strictly positive, they have a strictly positive lower bound. $\square$

Since the value of ε_0 depends on $\mathcal{D}^{\mathcal{L}}(X, <, d)$, it is non-trivial to compute. For practical applications, we recommend to choose a very small positive number for ε , but not so small that it becomes zero due to floating point rounding when added to the diameter of the partial dendrograms.

6.2 Idempotency of $\mathcal{HC}_{opt,\varepsilon}^{<\mathcal{L}}$

A detailed axiomatic analysis along the lines of for example Ackerman and Ben-David (2016) is beyond the scope of this paper, and is considered for future work. We still include a proof of idempotency of $\mathcal{HC}_{opt,\varepsilon}^{<\mathcal{L}}$, since this is an essential property of classical hierarchical clustering.

Idempotency of hierarchical clustering necessarily depends on the linkage function. We introduce the following concept, that allows us to prove this property for a range of linkage functions: We say that $\mathcal{L}$ is a **convex linkage function** if we always have

$$S\mathcal{L}(p, q, d) \leq \mathcal{L}(p, q, d) \leq C\mathcal{L}(p, q, d).$$

Notice that if $\mathbf{u}$ is an ultrametric on X , the ultrametric inequality yields

$$\mathbf{u}(a, b) = \text{sep}(X, \mathbf{u}) \Rightarrow \forall c \in X : \mathbf{u}(a, c) = \mathbf{u}(b, c),$$

so if $\mathcal{L}$ is a convex linkage function and $\mathbf{u}(a, b) = \text{sep}(X, \mathbf{u})$, we have

$$\mathcal{L}(\{a, b\}, \{c\}) = \mathcal{L}(\{a\}, \{c\}) = \mathcal{L}(\{b\}, \{c\}) \quad \forall c \neq a, b.$$

This is to say that a convex linkage function preserves the structure of the original ultrametric when minimal dissimilarity elements are merged. As a result, for any $\mathbf{u} \in \mathcal{U}(X)$, the set $\mathcal{D}^{\mathcal{L}}(X, \mathbf{u})$ contains exactly one element, namely the dendrogram corresponding to the ultrametric, which is why classical hierarchical clustering is idempotent.

For ordered spaces, the case is different. It is easy to construct an ordered ultrametric space $(X, <, \mathbf{u})$ for which $\mathbf{u}(a, b) = \text{sep}(X, \mathbf{u})$ and $a < b$, in which case the ultrametric cannot be reproduced. Hence, all of $\mathcal{U}(X)$ cannot be fixed points under $\mathfrak{U}_\varepsilon \circ \mathcal{HC}_{opt,\varepsilon}^{<\mathcal{L}}(X, <, -)$, but the mapping is still idempotent:

Theorem 11 (Idempotency) *For an ordered dissimilarity space $(X, <, d)$ and a convex linkage function $\mathcal{L}$, we have $\theta \in \mathcal{HC}_{opt,\epsilon}^{<\mathcal{L}}(X, <, d) \Rightarrow \mathcal{HC}_{opt,\epsilon}^{<\mathcal{L}}(X, <, \mathfrak{U}_\epsilon(\theta)) = \{\theta\}$.*

Proof Let $\theta(\infty) = \{B_i\}_{i=1}^m$. Then each B_i is an antichain in $(X, <)$, so we have

$$\forall x, y \in B_i : \text{sep}(B_i, \mathfrak{u}|_{B_i}) = \text{sep}_\perp(B_i, \mathfrak{u}|_{B_i}) \quad \text{for } 1 \leq i \leq m.$$

Since $\epsilon > 0$, we also have

$$x, y \in B_i \Rightarrow \mathfrak{u}(x, y) < \text{diam}(X, \mathfrak{u}) \quad \text{for } 1 \leq i \leq m.$$

And, lastly, since every pair of comparable elements are in pairwise different blocks, we have

$$x < y \vee y < x \Rightarrow \mathfrak{u}(x, y) = \text{diam}(X, \mathfrak{u}).$$

Now, since $\mathcal{L}$ is convex, based on the discussion preceding the theorem, the intra-block structure of every block will be preserved. And, since every inter-block dissimilarity is accompanied by comparability across blocks, the procedure for generation of $\mathcal{D}^{\mathcal{L}}(X, <, \mathfrak{U}_\epsilon(\theta))$ will exactly reproduce the intra block structure of all blocks and then halt. Hence, $\mathcal{D}^{\mathcal{L}}(X, <, \mathfrak{U}_\epsilon(\theta)) = \{\theta\}$. $\square$

7 Polynomial time approximation

In the absence of an efficient algorithm for $\mathcal{HC}_{opt,\epsilon}^{<\mathcal{L}}$, this section provides a polynomial time approximation scheme. The efficacy as approximation is demonstrated in Sect. 8, and a demonstration on real world data is given in Sect. 9.

Recall the set $\mathcal{D}^{\mathcal{L}}(X, <, d)$ of partial dendrograms over $(X, <, d)$ from Definition 12. The algorithm for producing a random element of $\mathcal{D}^{\mathcal{L}}(X, <, d)$ is described at the beginning of Sect. 6; the key is to pick a random pair for merging whenever we encounter a set of tied connections.

The approximation model is deceptively simple; we generate a set of random partial dendrograms, and choose the one with the best ultrametric fit.

Definition 14 Let $(X, <, d)$ be given, and let N be a positive integer. For any random selection of N partial dendrograms $\{\theta_i\}_i$ from $\mathcal{D}^{\mathcal{L}}(X, <, d)$, an **N -fold approximation of $\mathcal{HC}_{opt,\epsilon}^{<\mathcal{L}}(X, <, d)$** is a partial dendrogram $\theta \in \{\theta_i\}_i$ minimising $\|\mathfrak{U}_\epsilon(\theta) - d\|_p$. We denote the N -fold approximation scheme by $\mathcal{HC}_{N,\epsilon}^{<\mathcal{L}}$.

7.1 Running time complexity

Assume that $|X| = n$. In the worst case, we may have to check $\binom{n}{2}$ pairs to find one that is not comparable, and the test for $a \perp b$ has complexity $O(n^2)$, leading to a complexity of $O(n^4)$ of finding a mergeable pair. Since there are up to $n - 1$ merges, the worst case estimate of the running time complexity for producing one element in $\mathcal{D}^{\mathcal{L}}(X, <, d)$ is $O(n^5)$.

A part of this estimate is the number of comparability tests we have to perform in order to find a mergeable pair. For a sparse order relation, we may have to test significantly less

than $\binom{n}{2}$ pairs before finding a mergeable pair: if K is the expected number of test we have to do, the expected complexity of finding a mergeable pair becomes $O(Kn^2)$. This yields a total expected algorithmic complexity of $O(Kn^3)$. If the order relation is empty, we have $K = 1$, and the complexity of producing a dendrogram becomes $O(n^3)$, which is the running time complexity of classical hierarchical clustering. Hence, if the order relation is sparse, we can generally expect the algorithm to execute significantly faster than the worst case estimate.

When producing an N -fold approximation, the N random partial dendrograms can be generated in parallel, reducing the computational time of the approximation. For the required number of dendrograms to obtain a good approximation, please see Sect. 8.

8 Demonstration of approximation efficacy on randomly generated data

The purpose of the demonstration is to check to which degree the approximation reproduces the order preserving clusterings of $\mathcal{HC}_{opt,\epsilon}^{<\mathcal{L}}$. We start by describing the random data model and the quality measures we use in assessing the efficacy of the approximation, before presenting the experimental setup and the results.

8.1 Random ordered dissimilarity spaces

To test the correctness and convergence ratio of the approximation scheme, we employ randomly generated ordered dissimilarity spaces. The random model consists of two parts: the random partial order and the random dissimilarity measure.

8.1.1 Random partial order

A partial order is equivalent to a transitively closed directed acyclic graph, so we can use any random model for directed acyclic graphs to generate random partial orders. We choose to use the classical Erdős-Rényi random graph model (Bollobás, 2001). Recall that a directed acyclic graph on n vertices is a binary $n \times n$ adjacency matrix that is *permutation similar* to a strictly upper triangular matrix; that is, there exists a permutation that, when applied to both the rows and the columns of one matrix, transforms it into the other. Let this family of $n \times n$ matrices be denoted by $\mathbb{A}(n)$. For a number $p \in [0, 1]$, the sub-family $\mathbb{A}(n, p) \subseteq \mathbb{A}(n)$ is defined as follows: for $A \in \mathbb{A}(n)$, let A' be strictly upper triangular and permutation similar to A . Then each entry above the diagonal of A' is 1 with probability p . The sought partial order is the transitive closure of this graph; we denote the corresponding set of transitively closed directed acyclic graphs by $\bar{\mathbb{A}}(n, p)$.

8.1.2 Random dissimilarity measure

If $|X| = n$, a dissimilarity measure over X with no tied connections consists of $\binom{n}{2}$ distinct values. Hence, any permutation of the sequence $\{1, \dots, \binom{n}{2}\}$ is a non-tied random dissimilarity measure over X .

To generate tied connections, let $t \geq 1$ be the **expected number of ties per level**. That is, for each unique value in the dissimilarity measure, that value is expected to have multiplicity t . In the case where t does not divide $\binom{n}{2}$, we resolve this by setting the multiplicity of the largest dissimilarity to $\left(\binom{n}{2} \bmod t\right)$.

We write $\mathbb{D}(n, t)$ to denote the family of random dissimilarity measures over sets of n elements with an expected number of t ties per level.

Definition 15 Given positive integers n and t together with $p \in [0, 1]$, the family of **random ordered dissimilarity spaces generated by (n, p, t)** is given by

$$\mathbb{O}(n, p, t) = \overline{\mathbb{A}}(n, p) \times \mathbb{D}(n, t).$$

8.2 Measures of cluster quality

In the demonstration, we start by generating a random ordered dissimilarity space. We then run the optimal clustering method on the space, finding the optimal order preserving hierarchical clustering. Finally, we run the approximation scheme on the space and study to which degree the approximation manages to reproduce the optimal hierarchical clustering. For this, we need a quantitative measure of clustering quality relative a known optimum.

A large body of literature exists on the topic of comparing clusterings [see for instance Vinh et al. (2010) for a brief review]. We have landed on the rather popular *adjusted Rand index* (Hubert & Arabie, 1985) to measure the ability of the approximation in finding a decent partition, comparing against the optimal result.

Less work is done on this type of comparison for partial orders and directed acyclic graphs. We suggest to use a modified version of the adjusted Rand index for this purpose too, based on an adaptation of the Rand index used for network analysis (Hoffman et al., 2015). For an introduction to the Rand index, and also to some of the versions of the adjusted Rand index, see Rand (1971), Hubert and Arabie (1985), Gates and Ahn (2017).

8.2.1 Adjusted Rand index for partition quality

The Rand index compares two clusterings by computing the percentage of corresponding decisions made in forming the clusterings; that is, counting whether pairs of elements are placed together in both clusterings or apart in both clusterings. An adjusted Rand index reports in the range $(-\infty, 1]$, where zero is equivalent to a random draw, and anything above zero is better than chance. We use the adjusted Rand index (*ARI*) to compute the efficacy of the approximation in finding a partition close to a given planted partition. This corresponds to what Gates and Ahn (2017) refers to as a *one sided* Rand index, since one of the partitions are given, whereas the other is drawn from some distribution. In the below demonstration, we assume that the approximating partition is drawn from the set of all partitions over X under the uniform distribution.

8.2.2 Adjusted Rand index for induced order relations

When comparing induced orders on partitions over a set, unless the partitions coincide, it is not obvious which blocks in one partition correspond to which blocks in the other. To overcome this problem, we base our measurements on the base space projection:

Definition 16 For an ordered set (X, E) and a partition Q of X with induced order E' , the **base space projection of (Q, E') onto X** is the order relation E_Q on X defined as

$$(x, y) \in E_Q \Leftrightarrow ([x], [y]) \in E'.$$

This allows us to compare the induced orders in terms of different orders on X . Notice that if the induced order E' is a [strict] partial order on Q , then E_Q is a [strict] partial order on X .

Hoffman et al. (2015) demonstrate that the adjusted Rand index can be used to detect missing links in networks by computing the similarity of edge sets. The concept relies on the fact that a network link and a link in an equivalence relation are not that different: Both networks and equivalence relations are special classes of relations, and the Rand index simply counts the number of coincidences and mismatches between two relation sets. While Hoffman et al. (2015) uses the *ARI* to compare elements within a network, we use the same method to compare across networks.

Let A and B be the adjacency matrices of two base space projections, and let A_i denote the i -th row of A , and likewise for B_i . If $\langle a, b \rangle$ is the inner product of a and b , we define

$$\begin{aligned} a_i &= \langle A_i, B_i \rangle & c_i &= \langle A_i, 1 - B_i \rangle \\ b_i &= \langle 1 - A_i, B_i \rangle & d_i &= \langle 1 - A_i, 1 - B_i \rangle. \end{aligned}$$

Here, a_i is the number of common direct descendants of i in both relations, b_i is the number of descendants of i found in A but not in B , c_i is the number of descendants of i in B but not in A , while d_i counts the common non-descendants of i in the two relations. Using this, we can compute the **element wise adjusted order Rand index**

$$\bar{ARI}_i = \frac{2(a_i d_i - b_i c_i)}{(a_i + b_i)(b_i + d_i) + (a_i + c_i)(c_i + d_i)} \quad \text{for } 1 \leq i \leq n,$$

measuring the element wise order correlation between the base space projections in the Hubert–Arabie adjusted Rand index (Hubert & Arabie, 1985; Warrens, 2008)¹. Notice that we compare the i -th row in A to the i -row in B since these rows correspond to the projections' respective descendant relations for the i -th element in X . In Hoffman et al. (2015), the above index is computed for each element pair *within* the network to produce the intra-network similarity coefficient.

Since we are interested in the overall match, we choose to report on the mean value, defining the **adjusted order Rand index for A and B** as

¹ This particular formulation of the adjusted Rand index relies on the networks having known and fixed labels, so that we know which vertices map to which vertices (Warrens, 2008), which indeed holds for the base space projections of two different induced order relations.

Table 1 Parameter settings for the demonstrations

	n	$\mathcal{L}$	Link probability (p)	Expected ties (t)	Reference
Figure 5	200	$\mathcal{SL}, \mathcal{AL}, \mathcal{CL}$	0.01, 0.02, 0.05	5	Yes
Figure 6	200	$\mathcal{SL}, \mathcal{AL}, \mathcal{CL}$	0.05	3, 7	Yes
Figure 7	500	$\mathcal{SL}, \mathcal{AL}, \mathcal{CL}$	0.01	10, 50, 100	No
Figure 8	500	$\mathcal{SL}, \mathcal{AL}, \mathcal{CL}$	0.05	50, 100	No
Figure 9	500	$\mathcal{SL}, \mathcal{AL}, \mathcal{CL}$	0.10	100	No

The right-most column indicates whether the reference clustering is available or not. The left-most column refers to the figure wherein the outcome of the corresponding experiment is presented. The parameters have been chosen to illustrate how the algorithm behaviour changes with changing expected number of ties, changing link probability in the random partial order, and choice of linkage function

$$\bar{ARI}(A, B) = \frac{1}{n} \sum_{i=1}^n \bar{ARI}_i.$$

8.2.3 Normalised ultrametric fit

A natural choice of quality measure is to report the **ultrametric fit** $\|\mathcal{U}_\epsilon(\theta) - d\|_p$ of the obtained partial dendrogram θ , especially if we can compare it to the ultrametric fit of the optimal solution. The scale of the ultrametric fit depends heavily on both the size of the space and the order of the norm, so we choose to normalise. Also, we invert the normalised value, so that the optimal fit has a value of 1, and a worst possible fit has value 0. This makes it easy to compare the convergence of the ultrametric fit to the convergence of the ARI and $\bar{ARI}$.

Definition 17 Given a set of partial dendrograms $\{\theta_i\}$ over $(X, <, d)$, let their respective ultrametric fits be given by $\delta_i = \|\mathcal{U}_\epsilon(\theta_i) - d\|_p$. The **normalised ultrametric fit** are the corresponding values

$$\hat{\delta}_i = 1 - \frac{\delta_i - \min_i \{\delta_i\}}{\max_i \{\delta_i\} - \min_i \{\delta_i\}}.$$

In the presence of a reference solution, we substitute $\min_i \{\delta_i\}$ with the ultrametric fit of the reference.

8.2.4 Ultrametric fit relative the optimal ultrametric

The reference partition can be reached through different sequences of merges, and neither $\mathcal{AL}$ nor $\mathcal{CL}$ are invariant in this respect. Neither ARI , $\bar{ARI}$ nor ultrametric fit captures the match between the optimal hierarchy and the approximated hierarchy. We therefore also include plots of the difference between the optimal ultrametric $\mathbf{u}_{opt}$ and the approximated ultrametric $\mathbf{u}_{N,\epsilon}$. Since both ultrametries are equivalent to their respective hierarchies, the magnitude $\|\mathbf{u}_{opt} - \mathbf{u}_{N,\epsilon}\|_p$ can be interpreted as a measure of difference in hierarchies. In

Fig. 5 Efficacy for $n = 200$ and $t = 5$ with $p \in \{0.01, 0.02, 0.05\}$. The first axis is the size of the drawn sample

the below plots, this is reported as $opt.fit$. As for the ultrametric fit, we normalise and invert the values for easy comparison.

8.3 Demonstration on randomly generated data

The experiments in the demonstration split in two. First, we demonstrate the efficacy of the approximation relative a known optimal solution, to see to which degree $\mathcal{HC}_{N,\epsilon}^{<\mathcal{L}}$ manages to approximate $\mathcal{HC}_{opt,\epsilon}^{<\mathcal{L}}$. Second, we study the convergence rate of the ultrametric fit for larger spaces with much larger numbers of tied connections; spaces for which the optimal algorithm does not terminate within any reasonable time.

For each parameter combination in Table 1, a set of 30 random ordered dissimilarity spaces are generated. For each space, 100 approximations are generated according to the prescribed procedure. We then bootstrap the approximations to generate N -fold approximations for different N .

We present the results in terms of convergence plots, showing the efficacy of the approximation as a function of the sample size N . For the results where a reference solution is available, the plots contain four curves:

Fig. 6 Efficacy for $n = 200$ and $p = 0.05$ with $t \in \{3, 7\}$. The first axis is the size of the drawn sample. The plots for $t = 5$ can be found in the bottom row of Fig. 5

Fig. 7 Polynomial approximation rate for $n = 500$, $P = 0.01$ and $t \in \{10, 20, 40\}$. The first axis is the size of the drawn sample

Fig. 8 Polynomial approximation rate for $n = 500$, $p = 0.05$ and $t \in \{50, 100\}$. The first axis is the size of the drawn sample

$\mathbb{E}(ARI)$	The expected adjusted Rand index of the approximated partition.
$\mathbb{E}(\partial ARI)$	The expected adjusted Rand index of the approximated induced order.
$norm.fit$	The mean of the normalised fit.
$opt.fit$	The mean of the normalised difference between the approximated ultrametric and the optimal ultrametric.

For the results where no reference solution is available, we present the distribution of the normalised fit.

The results are presented in Figs. 5, 6, 7, 8 and 9. The parameter settings corresponding to the figures are given in Table 1 for easy reference, and are also repeated in the figure text.

As we can see from the below results, the approximation generally performs very well. We also see that a large expected number of tied connections requires larger sample size for a good approximation, while a more dense order relation (higher value of p) seems to require a smaller sample compared to a more sparse relation. We also see that there is a seemingly strong correlation between the ultrametric fit of the approximation and the similarity between the approximation ultrametric and the optimal ultrametric.

Regarding choice of linkage function, the approximation only requires small samples for both SL and AL , while CL requires larger samples for larger numbers of tied connections.

8.3.1 First conclusions

The first thing that strikes the eye is that the approximations converge very rapidly. Even for moderately sized spaces (~ 500 elements), it appears to be sufficient with 20 samples for SL and AL , and for smaller spaces (~ 200 elements), even fewer samples are required. We also notice that there is a strong correlation between the ARI , ∂ARI and normalised fit.

Fig. 9 Polynomial approximation rate for $n = 500$, $p = 0.10$ and $t = 100$. The first axis is the size of the drawn sample

For the part of the demonstration where we have no reference clustering, we cannot know for sure whether the best reported fit is also optimal. However, from the convergent behaviour of the data, and the strong correlation between optimality and normalised fit in Figs. 5 and 6, this points in the direction of convergence to the true optimum.

Only $\mathcal{CL}$ displays convergence issues, indicating that if one wishes to use $\mathcal{CL}$ for large spaces or large numbers of tied connections, it may be wise to do so in conjunction with convergence tests.

On the other hand, since $\mathcal{SC}$ is independent of tie resolution order, every sequence of merges ending in the same maximal partition will produce the same partial dendrogram. This explains why the convergence rate of $\mathcal{SC}$ is less affected by the expected number of tied connections than, say, $\mathcal{CL}$.

The convergence rate is very high in some of the plots of Figs. 8 and 9. The authors believe this is due the high probability of two random elements being comparable (high p in $\mathcal{O}(n, p, t)$), since a dense relation leads to fewer candidate solutions. This in contrast to the larger set of candidates for a more sparse relation, such as in Fig. 7.

On the other hand, as we can see in Figs. 7 and 8, keeping p fixed and increasing the number of tied connections, and thereby the number of possible branch points, causes a slower convergence rate.

To summarise, we see that the approximation is both good and effective for $\mathcal{SC}$ and $\mathcal{AL}$. For $\mathcal{CL}$, although the approximation method seems good, the required sample size must be increased in the presence of large amounts of tied connections.

9 Demonstration on data from the parts database

While the above demonstration shows that $\mathcal{HC}_{N,\epsilon}^{<\mathcal{L}}$ performs well with respect to approximating $\mathcal{HC}_{opt,\epsilon}^{<\mathcal{L}}$, another question is how order preserving hierarchical clustering deals with the dust of reality. In this section, we present results from applying the approximation algorithm to subsets of the parts database described briefly in Sect. 1.1. As benchmark, we run classical hierarchical clustering on the same problem instances, comparing the performance of the methods using ARI , $\bar{ARI}$ and loop frequency (described below). As hierarchical methods for constrained clustering do not offer a no-link constraint, we also propose a simplified approach simulating no-link behaviour for $\mathcal{AL}$ and $\mathcal{CL}$ which we call $\mathcal{HC}^+$.

Table 2 Some key characteristics of the connected components selected for the demonstration

cc no.	cc size	In/out deg.	p	Expected ties
0	12	0.92	0.17	2.36
1	14	0.93	0.14	4.79
2	13	0.92	0.15	2.17
3	40	1.27	0.07	8.97
4	20	1.35	0.14	3.96
5	11	1.18	0.24	2.20
6	20	1.10	0.12	4.22
7	20	0.95	0.10	3.96

The in/out deg. column provides the directed average degree when the data is considered as a DAG. The column p shows the probability for two random elements to be connected in the transitive reduction

9.1 Demonstration dataset

To select data for the demonstration we proceeded as follows: We considered the part-of relations as a directed graph, and extracted all the connected components. As it turned out, there was one gigantic component and a large number of singleton elements, but also a hand-full of connected components of 11 to 40 elements each. We selected these smaller connected components as our demo dataset without any further consideration. Dissimilarities between the elements were obtained from a dissimilarity measure produced by an ongoing project in the company working on the very task of classifying equivalent equipment. Some key characteristics of the data is provided in Table 2

Due to limited labeling of the data, we do not know which elements are copies of other elements, so we have to fake copying to produce planted partitions. For the demonstration, we pick a connected component (X^0, E^0) where $X^0 = \{x_1^0, \dots, x_n^0\}$, and for some positive number m we make $m - 1$ copies of X^0 and E^0 , leading to m partially ordered sets $\{(X^k, E^k)\}_{k=0}^{m-1}$. We then form their disjoint union (X, E) where $|X| = m|X^0|$. X now consists of m connected components, each a copy of the others. If $x_i^0 \in X^0$, then **the set of elements equivalent to x_i^0** is the set $\{x_i^k\}_{k=0}^{m-1} \subseteq X$. Hence, the clusters we seek are the sets on this form.

If we denote the dissimilarity measure that comes with the data by d_0 , we define the extension to all of X as follows: First, if both elements are in the same component X^k for $0 \leq k \leq m$, then we simply use d_0 . And if they are in different components, indicating that they are in a copy-relationship, we increase their dissimilarity by an offset $\alpha \geq 0$. Concretely, the extended dissimilarity $d^\alpha : X \times X \rightarrow \mathbb{R}_+$ is given by

$$d^\alpha(x_i^r, x_j^s) = \begin{cases} d_0(x_i^0, x_j^0) & \text{if } r = s, \\ \alpha + d_0(x_i^0, x_j^0) & \text{otherwise.} \end{cases}$$

This means that if x and y are copies of each other, then $d^\alpha(x, y) = \alpha$, and if x and y are in the same component and if z is a copy of x , then $d(z, y) = \alpha + d_0(x, y)$. Furthermore, for each modified distance, we add a small amount of Gaussian noise to α to induce some variability. As a result, two copies x_i^r and x_i^s are offset by approximately α , and by varying the magnitude of α we can study how the offset affects the clustering.

9.2 Simulated constrained clustering

The available methods for hierarchical constrained clustering do not easily incorporate the partial order as a constraint. What we would like to compare against, is hierarchical constrained clustering with do-not-cluster constraints. For $\mathcal{CL}$ and $\mathcal{AL}$, we can obtain this by setting the dissimilarity between comparable elements to a sufficiently large number, causing all comparable elements to be merged towards the end. Indeed, for $\mathcal{CL}$ it is sufficient to set this dissimilarity to any value exceeding $\max\{d^a\}$, and as the below demonstration shows, this value works equally well for $\mathcal{AL}$. We denote hierarchical clustering with this kind of modified dissimilarity by $\mathcal{HC}^{+\mathcal{L}}$.

Since $d_0 < 1$ for all pairs of elements, we chose to use 1.0 as our maximum dissimilarity.

9.3 A measure of order preservation

While the $\bar{ARI}$ measures the correlation between the induced order of the planted partition and the induced order of the obtained clustering, the $\bar{ARI}$ does not convey information about whether the induced relation is a partial order or not. Since this is a key question for applications where order preservation is of high importance (such as acyclic partitioning of graphs), we suggest the following simple measure.

Let (Q, E') be a partition of (X, E) , and let E_Q be the base space projection of (Q, E') onto X (Definition 16). We say that (Q, E') **induces a loop** if there are elements on the form $(x, x) \in E_Q$. The number of loops induced by (Q, E') is thus the quantity $|\{(x, y) \in E_Q \mid x = y\}|$. There is at most one loop per element of X , and if E_Q contains a cycle, then every element of the cycle corresponds to a loop. In the name of normalisation, we measure the amount of loops as the fraction of elements in X that is a part of a cycle:

$$\text{loops}(Q, E') = \frac{|\{(x, y) \in E_Q \mid x = y\}|}{|X|}.$$

9.4 Picking a clustering in the hierarchy for comparison

Given a problem instance $(X, <, d)$ and a planted partition $Q \in \mathfrak{P}(X, <)$, the planted induced partial order is necessarily the induced relation $<'$. But in comparing a hierarchical clustering with a planted partition, we have to make a choice of clustering in the hierarchy. Given a hierarchical clustering, we choose to find the clustering in the hierarchy that has the highest ARI relative the planted partition. We then report all of ARI , $\bar{ARI}$ and loops with regards to this clustering.

9.5 Variance of the difference

In the below plots, we present the mean values of ARI , $\bar{ARI}$ and loops together with a visual indication of variability. For each instance of a random ordered dissimilarity space $(X, <, d)$, we run all of $\mathcal{HC}_{N,\varepsilon}^{<\mathcal{L}}$, $\mathcal{HC}^{\mathcal{L}}$ and $\mathcal{HC}^{+\mathcal{L}}$. Thus, we can analyse the performance of

Table 3 Parameters for execution of experiments

Parameter	Value(s)	Explanation
α	{0.10, 0.15, ..., 0.50}	Mean copy dissimilarity
σ	0.10	Variance of α
$\mathcal{L}$	{ $\mathcal{SL}$, $\mathcal{AL}$, $\mathcal{CL}$ }	Linkage models
m	$\lceil 200/ X^0 \rceil$	Number of copies (see below)
N	10	Sample size in the N -fold approximation
ϵ	10^{-12}	Ultrametric completion level
p	1	Choice of norm for ultrametric fitting

The number m of copies is the least number for which the total number of elements, $m|X^0|$, is at least 200

the methods by pairwise comparison on a problem instance level. That is, we choose to consider pairwise differences such as

$$ARI(\mathcal{HC}_{N,\epsilon}^{<\mathcal{L}}(X, <, d)) - ARI(\mathcal{HC}^{+\mathcal{L}}(X, <, d))$$

as one random variable, and likewise for $\bar{d}ARI$ and loops. The variance of this random variable shows the variance in the difference, and we can use this magnitude to analyse whether the sets of results are statistically distinguishable. For the below plots, we mark a region about each line corresponding to one standard deviation of this random variable. This means that the regions encompassing the lines will not overlap unless the difference between the mean values is less than two standard deviations.

To reduce the number of plots, we choose to plot the results of all three methods together. This is obviously impractical with respect to pairwise comparisons, so we employ the following convention: the indicated variance about the mean of $\mathcal{HC}_{N,\epsilon}^{<\mathcal{L}}$ and $\mathcal{HC}^{+\mathcal{L}}$ is the standard deviation of the differences between these methods. The indicated variance about the mean of $\mathcal{HC}^{\mathcal{L}}$ represents the standard deviation of the differences between $\mathcal{HC}^{\mathcal{L}}$ and $\mathcal{HC}_{N,\epsilon}^{<\mathcal{L}}$.

9.6 Execution and results

The parameters given in Table 3 define how the ordered dissimilarity spaces are constructed for each of the connected components. For each instance of an ordered dissimilarity space, $\mathcal{HC}_{N,\epsilon}^{<\mathcal{L}}$, $\mathcal{HC}^{\mathcal{L}}$ and $\mathcal{HC}^{+\mathcal{L}}$ are all run on the same instance with a choice of linkage function $\mathcal{L} \in \{\mathcal{SL}, \mathcal{AL}, \mathcal{CL}\}$. This allows us to compare the performance of the methods against each other on a per-instance basis. For each parameter combination in $\{\alpha\} \times \{\mathcal{SL}, \mathcal{AL}, \mathcal{CL}\}$, we repeated this process 50 times. The variance of the difference is based on these sets of 50 executions.

We present three families of plots, for ARI , $\bar{d}ARI$ and loops, respectively. We have picked three connected components for the presentation that we believe represent the span of observations. The full set of plots is provided in the Appendix.

First, connected component number 7 (*cc7*) is the sample on which we see the most clear benefit from using $\mathcal{HC}_{N,\epsilon}^{<\mathcal{L}}$, significantly outperforming both $\mathcal{HC}^{\mathcal{L}}$ and $\mathcal{HC}^{+\mathcal{L}}$ on all quality measures. Although *cc7* is not representative for the majority of observations, it is empirical evidence that there exist problem instances for which order preserving

Fig. 10 Performance of the different clustering methods with respect to ARI on connected components 7, 1 and 4. The shaded regions represent one standard deviation of the pairwise differences, as described in Sect. 9.5

clustering cannot be well approximated by hierarchical constrained clustering through do-not-cluster constraints.

Connected component number 1 ($cc1$) represents the majority of the instances. While $\mathcal{HC}_{N,\varepsilon}^{<L}$ still is best in class with respect to all quality measures, we see that for $\mathcal{AL}$ and $\mathcal{CL}$ the method $\mathcal{HC}^{+L}$ performs equally well with respect to ARI and sometimes also $\bar{o}ARI$.

At the other extreme of $cc7$ there is connected component number 4 ($cc4$), presented in the bottom row of Fig. 10. For this component, all the clustering models perform equally well in all quality measures, indicating that they produce the exact same clusterings. This can only be explained by the fact that the original dissimilarity measure d_0 , when restricted to this component, both is an ultrametric, and incorporates the order relation (Sect. 6.2).

The results are also summarised in Table 4 after the plots (Figs. 11, 12).

We summarise the experiment observations in Table 4. As can be seen from the table, $\mathcal{HC}_{N,\varepsilon}^{<L}$ is best in class in every category. However, $\mathcal{HC}^{+L}$ is also best in class in 81% of the cases when we restrict our attention to ARI and $L \in \{\mathcal{AL}, \mathcal{CL}\}$.

Fig. 11 Performance of the different clustering methods with respect to $\bar{d}ARI$ on connected components 7 and 1. The shaded regions represent one standard deviation of the pairwise differences, as described in Sect. 9.5

Fig. 12 Performance of the different clustering methods with respect to loops on connected components 7 and 1. The shaded regions represent one standard deviation of the pairwise differences, as described in Sect. 9.5

Table 4 The table presents for how many of the eight selected samples the different methods are best in class with regards to ARI , $\bar{a}RI$ and loops

	$\mathcal{HC}^{<\mathcal{L}}$			$\mathcal{HC}^{\mathcal{L}}$			$\mathcal{HC}^{+\mathcal{L}}$		
	ARI	$\bar{a}RI$	Loops	ARI	$\bar{a}RI$	Loops	ARI	$\bar{a}RI$	Loops
$\mathcal{SL}$	8	8	8	4	4	3	4	4	2
$\mathcal{AL}$	8	8	8	2	2	1	7	6	1
$\mathcal{CL}$	8	8	8	2	3	0	7	6	1
	100%	100%	100%	33%	37%	16%	75%	66%	16%

The scores are based on visual inspection of the plots. For ARI and $\bar{a}RI$, we count a one if there is less than one standard deviation to the best plot in at least half the sampled α values and zero otherwise. For loops, we count a one if the expected value is zero throughout. The full list of plots can be found in “Appendix 1”

To conclude, we see that if clustering is the sole objective, then $\mathcal{HC}^{+\mathcal{L}}$ is a good alternative to $\mathcal{HC}^{<\mathcal{L}}$ whenever $\mathcal{L} \in \{\mathcal{AL}, \mathcal{CL}\}$. If order preservation, or acyclic partitioning, is of any importance, then $\mathcal{HC}_{N,\epsilon}^{<\mathcal{L}}$ is the only viable method among those we have tested.

Moreover, as demonstrated by the top row of Fig. 10, although $\mathcal{HC}^{+\mathcal{L}}$ may be a good approximation of $\mathcal{HC}_{N,\epsilon}^{<\mathcal{L}}$ when $\mathcal{L} \in \{\mathcal{AL}, \mathcal{CL}\}$, there are problem instances on which the latter outperforms the former with significant margin, also for ARI .

10 Summing up

In this paper we have put forth a theory for order preserving hierarchical agglomerative clustering for strictly partially ordered sets. The clustering uses classical linkage functions such as single-, average-, and complete linkage. The clustering is optimisation based, and therefore also permutation invariant.

The output of the clustering process is partial dendrograms; sub-trees of dendrograms with several connected components. We have shown that the family of partial dendrograms over a set embed into the family of dendrograms over the set.

When applying the theory to non-ordered sets, we see that we have a new theory for hierarchical agglomerative clustering that is very close to the classical theory, but that is optimisation based rather than algorithmic. Differently from classical hierarchical clustering, our theory is permutation invariant. We have shown that for single linkage, the theory coincides with classical hierarchical clustering, while for complete linkage, the clustering problem becomes NP-hard. However, the computational complexity is directly linked to the number of tied connections, and in the absence of tied connections, the theories coincide.

We present a polynomial approximation scheme for the clustering theory, and demonstrate its convergence properties and efficacy on randomly generated data. We also provide a demonstration on real world data comparing against existing methods, showing that our model is best in class in all selected quality measures.

10.1 Future work topics

We suggest the following future work topics:

10.1.1 Complexity

While NP-hardness of $\mathcal{HC}_{opt,\epsilon}^{<\mathcal{CL}}$ follows from Theorem 3, the complexity classes of order preserving hierarchical agglomerative clustering for $\mathcal{SL}$ and $\mathcal{AL}$ remain to be established.

10.1.2 Order versus dissimilarity

Since the order relation is treated as a binary constraint it has a significant effect on the output from the clustering process, and may in some cases lead to undesirable outcomes. For example, if the dissimilarity measure associates “wrong” elements for clustering, the induced order relation may exclude future merges of elements correctly belonging together by erroneously identifying them as comparable. Also, if elements are wrongly identified as comparable to begin with, they can never be merged. Both due to Theorem 5.

Together, these observations indicate that “loosening up” the stringent nature of the order relation may be beneficial in applications where order preservation is not a strict requirement.

10.1.3 Other models for clustering

While we have chosen to develop our theory based on classical hierarchical clustering, it is likely that the theory we have presented can be extended or adjusted to apply to generalisations of hierarchical clustering too. As an example, we mention overlapping clusters and hierarchies of overlapping clusters (Jeantet et al., 2020). While hierarchies of overlapping clusters lead to DAGs, rather than trees, it still seems likely that a completion along the lines of Sect. 5 can be applied to the *partial DAGs* that necessarily arise when this type of hierarchical clustering is applied to strict partial orders in an order preserving fashion.

Also, as mentioned in the introduction, the framework we have presented can be applied directly to enable a theory for hierarchical agglomerative clustering in the presence of must-link and no-link constraints: The no-link constraints give rise to partial dendrograms that are easily evaluated via ultrametric completion.

Appendix 1: Plots from the part database demo

This section lists all the plots from the experiments described in Sect. 9. The plots are grouped per connected component, and present results for all clustering methods, quality measures and linkage models. Please see Table 2 for a list of statistical properties of the different connected components, and Table 3 for the parameter settings used during the experiments.

Results for connected component no. 0.

Results for connected component no. 1.

Results for connected component no. 2.

Results for connected component no. 3.

Results for connected component no. 4.

Results for connected component no. 5.

Results for connected component no. 6.

Results for connected component no. 7.

Appendix 2: Reference implementation

The implementation used for the experiments in Sects. 8 and 9 is available as open source at <https://bitbucket.org/Bakkelund/opnac>.

Acknowledgements I wish to thank each of the anonymous reviewers at Machine Learning for constructive feedback and comments greatly improving the exposition. I also owe a great thank you to Henrik Forssell, Department of Informatics (University of Oslo), and Gudmund Hermansen, Department of Mathematics (University of Oslo), for their comments, questions and discussions leading up to this work.

Funding Open access funding provided by University of Oslo (incl Oslo University Hospital). This work was funded by the DataScience@UiO innovation cluster (The Faculty of Mathematics and Natural Sciences, University of Oslo).

Code availability The code used for the demonstration part of the article is available as open source from <https://bitbucket.org/Bakkelund/opnac>.

Declarations

Conflict of interest None.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article

are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ackerman, M., & Ben-David, S. (2016). A characterization of linkage-based hierarchical clustering. *Journal of Machine Learning Research*, 17(231), 1–17.
- Basu, S., Davidson, I., & Wagstaff, K. (2008). *Constrained clustering: Advances in algorithms, theory, and applications* (1st ed.). Chapman & Hall/CRC.
- Blyth, T. (2005). *Lattices and ordered algebraic structures*. Universitext, Springer.
- Bollobás, B. (2001). *Random Graphs, Cambridge Studies in Advanced Mathematics* (2nd ed.). Cambridge University Press. <https://doi.org/10.1017/CBO9780511814068>.
- Buluç, A., Meyerhenke, H., Safro, I., Sanders, P., & Schulz, C. (2016). *Recent advances in graph partitioning* (pp. 117–158). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-49487-6_4.
- Carlsson, G., & Mémoi, F. (2010). Characterization, stability and convergence of hierarchical clustering methods. *Journal of Machine Learning Research*, 11, 1425–1470.
- Carlsson, G., Mémoi, F., Ribeiro, A., & Segarra, S. (2014). Hierarchical quasi-clustering methods for asymmetric networks. In E. P. Xing, & T. Jebara (Eds.), *Proceedings of the 31st international conference on machine learning* (Vol. 32, pp. 352–360). PMLR, Beijing, China, Proceedings of Machine Learning Research.
- Chierchia, G., & Perret, B. (2019). Ultrametric fitting by gradient descent. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, & R. Garnett (Eds.), *Advances in neural information processing systems* (Vol. 32, pp. 3181–3192). Curran Associates, Inc.
- Dasgupta, S. (2016). A cost function for similarity-based hierarchical clustering. In *Proceedings of the forty-eighth annual ACM symposium on theory of computing* (pp. 118–127). Association for Computing Machinery, New York, NY, USA, STOC '16, <https://doi.org/10.1145/2897518.2897527>.
- Davidson, I., & Ravi, S. S. (2005). Agglomerative hierarchical clustering with constraints: Theoretical and empirical results. In A. M. Jorge, L. Torgo, P. Brazdil, R. Camacho, & J. Gama (Eds.), *Knowledge discovery in databases: PKDD 2005* (pp. 59–70). Springer.
- De Soete, G., Carroll, J. D., & DeSarbo, W. S. (1987). Least squares algorithms for constructing constrained ultrametric and additive tree representations of symmetric proximity data. *Journal of Classification*, 4(2), 155–173. <https://doi.org/10.1007/BF01896984>.
- Fraley, C., & Raftery, A. E. (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association*, 97(458), 611–631. <https://doi.org/10.1198/016214502760047131>.
- Gates, A. J., & Ahn, Y. Y. (2017). The impact of random models on clustering similarity. *Journal of Machine Learning Research*, 18(87), 1–28.
- Ghoshdastidar, D., Perrot, M., & von Luxburg, U. (2019). Foundations of comparison-based hierarchical clustering. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, R. Garnett (Eds.), *Advances in neural information processing systems 32* (pp. 7456–7466). Curran Associates, Inc.
- Gilpin, S., Nijssen, S., & Davidson, I. (2013). Formalizing hierarchical clustering as integer linear programming. In *Proceedings of the twenty-seventh AAAI conference on artificial intelligence* (pp. 372–378). AAAI Press, AAAI'13.
- Herrmann, J., Kho, J., Uçar, B., Kaya, K., & Çatalyürek, Ü. V. (2017). Acyclic partitioning of large directed acyclic graphs. In *2017 17th IEEE/ACM international symposium on cluster, cloud and grid computing (CCGRID)* (pp. 371–380). <https://doi.org/10.1109/CCGRID.2017.101>.
- Hoffman, M., Steinley, D., & Brusco, M. J. (2015). A note on using the adjusted rand index for link prediction in networks. *Social Networks*, 42, 72–79. <https://doi.org/10.1016/j.socnet.2015.03.002>.
- Holly, J. E. (2001). Pictures of ultrametric spaces, the p-adic numbers, and valued fields. *The American Mathematical Monthly*, 108(8), 721–728.
- Hubert, L., & Arabie, P. (1985). Comparing partitions. *Journal of Classification*, 2(1), 193–218.
- Hughes, B. (2004). Trees and ultrametric spaces: A categorical equivalence. *Advances in Mathematics*, 189(1), 148–191. <https://doi.org/10.1016/j.aim.2003.11.008>.
- Jain, A. K., & Dubes, R. C. (1988). *Algorithms for clustering data*. Prentice-Hall, Inc.

- Janowitz, M. F. (2010). *Ordinal and relational clustering*. World Scientific. <https://doi.org/10.1142/7449>.
- Jardine, N., & Sibson, R. (1971). *Mathematical Taxonomy*. Wiley series in probability and mathematical statistics, Wiley.
- Jeanet, I., Miklos, Z., & Gross-Amblard, D. (2020). Overlapping hierarchical clustering (OHC). In *Intelligent data analysis (IDA 2020)*.
- Johnson, S. C. (1967). Hierarchical clustering schemes. *Psychometrika*, 32(3), 241–254. <https://doi.org/10.1007/BF02289588>.
- Kamishima, T., & Fujiki, J. (2003). Clustering orders. In G. Grieser, Y. Tanaka, & A. Yamamoto (Eds.), *Discovery science* (pp. 194–207). Springer. https://doi.org/10.1007/978-3-540-39644-4_17.
- Karp, R. M. (1972). *Reducibility among combinatorial problems* (pp. 85–103). Springer US. https://doi.org/10.1007/978-1-4684-2001-2_9.
- Kriegel, H. P., Kröger, P., Sander, J., & Zimek, A. (2011). Density-based clustering. *WIREs Data Mining and Knowledge Discovery*, 1(3), 231–240. <https://doi.org/10.1002/widm.30>.
- Lassila, O., & Swick, R. R. (1999). Resource description framework (RDF) model and syntax specification. W3c recommendation, W3C.
- Li, X., Wu, Y., Ester, M., Kao, B., Wang, X., & Zheng, Y. (2017). Semi-supervised clustering in attributed heterogeneous information networks. In *Proceedings of the 26th international conference on world wide web*, International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, WWW '17 (pp. 1621–1629), <https://doi.org/10.1145/3038912.3052576>.
- Łuczak, M. (2016). Hierarchical clustering of time series data with parametric derivative dynamic time warping. *Expert Systems with Applications*, 62, 116–130. <https://doi.org/10.1016/j.eswa.2016.06.012>.
- Macqueen, J. (1967). Some methods for classification and analysis of multivariate observations. In *In 5-th Berkeley symposium on mathematical statistics and probability* (pp. 281–297).
- Markov, I. L., Hu, J., & Kim, M. (2015). Progress and challenges in VLSI placement research. *Proceedings of the IEEE*, 103(11), 1985–2003.
- Pio, G., Serafino, F., Malerba, D., & Ceci, M. (2018). Multi-type clustering and classification from heterogeneous networks. *Information Sciences*, 425, 107–126. <https://doi.org/10.1016/j.ins.2017.10.021>.
- Rammal, R., Toulouse, G., & Virasoro, M. A. (1986). Ultrametricity for physicists. *Reviews of Modern Physics*, 58, 765–788. <https://doi.org/10.1103/RevModPhys.58.765>.
- Rand, W. M. (1971). Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*, 66(336), 846–850.
- Sneath, P., & Sokal, R. (1973). *Numerical taxonomy: The principles and practice of numerical classification*. A Series of books in biology. W. H. Freeman.
- Vinh, N. X., Epps, J., & Bailey, J. (2010). Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *Journal of Machine Learning Research*, 11(95), 2837–2854.
- Ward, J. H. (1963). Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association*, 58(301), 236–244.
- Warrens, M. J. (2008). On the equivalence of Cohen's Kappa and the Hubert-Arabie adjusted Rand index. *Journal of Classification*, 25, 177–183.