

RESEARCH ARTICLE

Peer groups for organisational learning: Clustering with practical constraints

Daniel W. Kennedy ^{1,2*}, Jessica Cameron ^{1,2}, Paul P. -Y. Wu ^{1,2}, Kerrie Mengersen^{1,2}

1 Science and Engineering Faculty, Queensland University of Technology, Brisbane, Queensland, Australia, 2 ARC Centre for Excellence in Mathematical and Statistical Frontiers, Brisbane, Queensland, Australia

* d8.kennedy@qut.edu.au



Abstract

Peer-grouping is used in many sectors for organisational learning, policy implementation, and benchmarking. Clustering provides a statistical, data-driven method for constructing meaningful peer groups, but peer groups must be compatible with business constraints such as size and stability considerations. Additionally, statistical peer groups are constructed from many different variables, and can be difficult to understand, especially for non-statistical audiences. We developed methodology to apply business constraints to clustering solutions and allow the decision-maker to choose the balance between statistical goodness-of-fit and conformity to business constraints. Several tools were utilised to identify complex distinguishing features in peer groups, and a number of visualisations are developed to explain high-dimensional clusters for non-statistical audiences. In a case study where peer group size was required to be small (≤ 100 members), we applied constrained clustering to a noisy high-dimensional data-set over two subsequent years, ensuring that the clusters were sufficiently stable between years. Our approach not only satisfied clustering constraints on the test data, but maintained an almost monotonic negative relationship between goodness-of-fit and stability between subsequent years. We demonstrated in the context of the case study how distinguishing features between clusters can be communicated clearly to different stakeholders with substantial and limited statistical knowledge.

OPEN ACCESS

Citation: Kennedy DW, Cameron J, Wu PP-Y, Mengersen K (2021) Peer groups for organisational learning: Clustering with practical constraints. PLoS ONE 16(6): e0251723. <https://doi.org/10.1371/journal.pone.0251723>

Editor: Renato Cordeiro de Amorim, University of Essex, UNITED KINGDOM

Received: October 18, 2020

Accepted: May 2, 2021

Published: June 1, 2021

Copyright: © 2021 Kennedy et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The raw data cannot be shared publicly due to participant confidentiality. However, all pre-processed and deidentified data required to replicate the study's results are uploaded to GitHub and publicly accessible via the following URL: <https://github.com/paul1010/peer-group-clustering>.

Funding: KM and DK are supported by an ARC Laureate Fellowship (FL150100150). The authors would like to acknowledge the agency which contributed financial support towards this research.

Introduction

Organisational benchmarking has become increasingly widespread across industries, including health, education and government [1]. An inherent consideration in any benchmarking exercise, which is suggested as a means for organisational learning and quality assurance, is that of comparing and contextualising an organisation's performance against its peers, known as peer-grouping [2, 3]. However, identifying relevant peer organisations is challenging due to the variability and uncertainty in the mix of clients and services provided by each organisation [4–6]. Although unsupervised clustering methods might appear as the solution to the peer-grouping problem, a key challenge to practical uptake lies in the management context where clustering outcomes need to be interpretable by stakeholders of varied backgrounds. In addition, the goals of organisational learning and benchmarking implicitly require peer groups to

Competing interests: The authors have declared that no competing interests exist.

remain sufficiently similar for long periods of time for such activities to occur. Although the clustering literature is very well developed [7], the unique needs of practical peer-grouping have yet to be specifically addressed.

Previous work

Peer groups have a strong similarity with the concept of strategic groups from industry organisation economics, which are defined as groups of companies with similar business models. They are operationally defined using variables relevant to industry, as groups with more in common than firms in different groups [8]. Study of strategic groups has a long history dating back to the 1970's, with clustering being among the original techniques for identifying strategic groups [9]. Clustering is characterised as the empirical method of identifying strategic groups against the theoretical methods which are expert-elicited, based on known qualitative differentiations between company strategies.

Inter-organisational learning. One of the key goals of peer-grouping is peer mentoring, often referred to as inter-organisational learning, for two-way transmission of ideas and knowledge to improve performance [10, p.54]. A study has shown that mentoring activities from external peers are viewed favourably in the local government space [11] and it has been argued that constructive feedback provided by peers on strengths and weaknesses of an organisation is seen as more independent and encouraging to organisations [12]. Although diversity in peer groups is recognised as being important for satisfactory peer mentoring, sufficient similarity is also required for real dialogue and trust [13, p. 105]. It has also been argued that to best facilitate the sharing of information and practices, peer groups must be stable (i.e. remain the same or similar) over time [14] and need to be a manageable size.

Clustering. Clustering has been used to identify peer groups in contexts such as hospitals [15] and educational institutions [16]. Clustering generally uses a measure of the similarity of data between pairs of observations, called the dissimilarity measure, to find groups where their pairwise dissimilarities are lower within the group than for pairs of observations in different groups [17, p.502]. For example, [18] employ k-means clustering to identify groups of similarly sized local government municipalities and socioeconomic demographics. In contrast, for the agglomerative hierarchical clustering method, organisations begin in separate clusters and are incrementally grouped into progressively larger clusters [7]. The result is a bifurcating 'tree' structure where each the successive merging of clusters into larger ones forms a node on the tree, resulting in a more flexible but less interpretable result when compared with the k-means approach. Hierarchical clustering has also been employed for peer-grouping [19, 20] and suggested for its flexibility and simple structure. In this approach, there is the potential for the number of observation in the cluster, referred to here as cluster size, to be directly controlled by a straightforward extension, since any merging step which would result in a cluster above the given limit can be blocked. In addition to the agglomerative approach described above, there are divisive forms of hierarchical clustering, where initially all observations are contained within a single group, and in successive steps the groups and subgroups are divided until all observations are in their own singleton groups. However, as with the k-means approach, the divisive approach lacks a means of enforcing cluster size constraints.

In general, clustering relies on discriminating characteristics such that organisations within a cluster are more similar and organisations in different clusters are more dissimilar (i.e. maximising discrimination between groups). Organisational benchmarking contexts tend to involve many variables, creating a high-dimensional clustering problem [1]. As a result, cluster interpretability can be challenging, not just due to the curse of dimensionality, but also to potential non-linearities and statistical dependence/correlation between variables.

Visualisation and communication is also more challenging, although techniques such as the co-plot [2] have been used to explore organisational clusters and their relationship to different variables.

Of critical importance for clustering is determining a quantitative measure of pairwise dissimilarity between observations [17, p.502]. Most often this dissimilarity is calculated based on differences for relevant variables, and then combined into a single values for each pair of observations. Mixture models are probability models describing data which are derived from latent subgroups with different underlying distributions [21]. Gaussian mixture models have multivariate Gaussian distributions with subgroup-specific means and covariances between variables. A popular non-parametric approach is the Dirichlet Process Mixture Model [22], which allows for an unknown number of subgroups, removing the need to pre-specify this number or fit multiple models with different numbers of subgroups. Using Markov Chain Monte Carlo (MCMC) sampling, the probability, given the data that a pair of observations is in the same subgroup, can be estimated for all pairs and then used as the dissimilarity measure [23].

Practical peer-grouping. A common practical consideration in establishing peer groups to facilitate organisational learning is constraining the number of organisations in each group. Minimum size constraints have been previously explored for k-means clustering [24]. Most research, however, has focused on clustering with must-link and cannot-link constraints (for a recent example see [25]). In addition, peer groups need to remain the same or similar over time [14] to facilitate inter-organisational learning. Some recent work has gone into monitoring cluster stability [26], and earlier work has sought to characterise specific transitional behaviours [27, 28] However, in the context of peer-grouping, the need is to control stability rather than observe or analyse it. Currently, peer-grouping that simultaneously addresses the need to produce clusters that reflect dissimilarities between organisations in the data, while maintaining consistency of clusters over time and achieving cluster size constraints, is an understudied problem.

In this paper, we develop a three-step methodology for aligning the statistical problem of clustering observations with the practical problem of peer-grouping, in the context of multivariate data. In the first step, data is pre-processed and then used to inform a Dirichlet Process Mixture Model (DPMM), from which the main outputs are probabilistic estimates of similarity between observations. In the second step, we use these similarity estimates and context-informed constraints to construct peer groups, which are both informed by data and usable in the application. Finally, we present methods for identifying key variables and features which distinguish between the peer groups, and visualisations to communicate these distinctions to both technical and non-technical audiences. Peer groups constructed through this methodology are justified statistically as having optimal goodness-of-fit while also being fit-for-purpose and understandable, which promotes trust in stakeholders.

Materials and methods

In this section, the three-step methodology is outlined, and then described in detail (see Fig 1). Since the main contributions of the work are the second step (construct peer groups) and third step (identify and communicate key features), more emphasis is placed on these steps. With cluster analysis as the tool used to construct the peer groups from data, the terms ‘cluster’ and ‘peer group’ are used interchangeably throughout the manuscript.

The case study and data are described, along with initial summarising, transforming and rescaling of the data. We use the well-known Dirichlet Process Mixture Model (DPMM) as the basis for computing a dissimilarity matrix as a representation of how similar organisations are

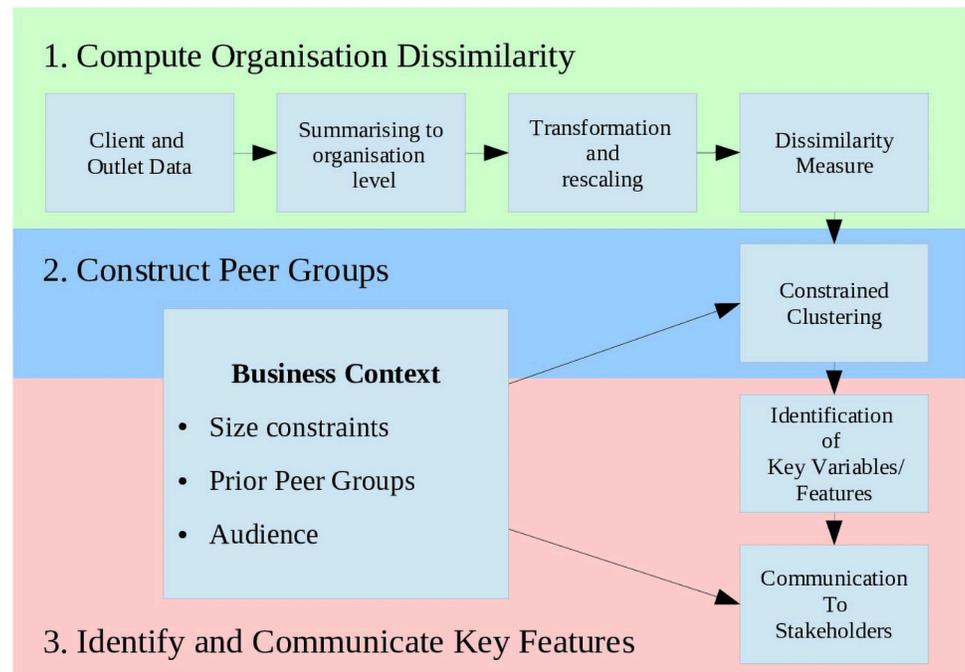


Fig 1. Diagram describing the 3-step methodology. The business context, which are the features of the peer-grouping application, have an influence over clustering process, as well as the communication of the peer group features to analysts, policy-makers, and stakeholders.

<https://doi.org/10.1371/journal.pone.0251723.g001>

to each other, based on the data. We present three novel methods for constructing peer groups informed by both the data (via the dissimilarity matrix), and constraints on size and the presence of prior peer groups. We use classification methods to automatically identify the important variables and groups of variables which define and discriminate between peer groups. Finally, we present some standard and novel visualisations for communicating these clusters to different audiences, which in the context of peer groups includes different levels of statistical and numerical proficiency.

Case study and data

The development of methodologies in this paper was motivated by a case study of a large organisation, with many outlets delivering a variety of services to clients. The organisations to be peer-grouped here are the outlets. The de-identified data comprise of many relevant factors for characterising differences between organisations and their client bases. These include different measures of client demographics, such as age, socioeconomic status, and regionality (e.g. metropolitan, remote, or regional), as well as organisational characteristics, such as client numbers and frequency of services. For reason of confidentiality, covariate names have been removed. Two of the covariates are labelled ‘response 1’ and ‘response 2’; however, these are not true responses and the clustering analysis is completely unsupervised.

Summary statistics such as the average client age, and demographic groups as represented by proportion of client base, were computed over the outlet’s clients for the financial year. Domain experts were consulted to reduce the number of variables down to 14 most important summary statistics. These were expressed as continuous variables or proportions. To improve comparability between variables, proportions were transformed onto the continuous scale using a logit transformation [29], with 0 and 1 values shifted slightly towards the centre. A

small number of variables were highly right-skewed, so a log-transformation was applied. Finally, all variables were scaled and centred post-transformation.

All analysis for this and subsequent steps was conducted in R [30]. Full Rmarkdown analysis files are available online: <https://github.com/paul1010/peer-group-clustering>.

Computing dissimilarity using the DPMM

Because the complexity of the data, a non-parametric Dirichlet Process Mixture Model (DPMM) method was fitted to cluster the data into peer groups for both years in which data were available: 2017 and 2018. Separate DPMMs were fitted to the data for each year. The output of this method was a posterior dissimilarity matrix (PDM), which gave the probability, under the DPMM, that a given pair of organisations were not in the same latent group. The DPMM was chosen to compute a dissimilarity matrix over a standard distance method (e.g. Euclidean) because of the complexity of the data. The DPMM's flexibility allows for groups with different concentrations of observations, and is better able to deal with the high dimensionality of the covariate space, multicollinearity, non-linearity and non-Gaussian distributions, and actually assumes latent grouping in the data. An additional benefit of the DPMM is that both continuous and discrete variables can be used, although in this work we only explore continuous variables.

An initial analysis from this dissimilarity matrix (based on the DPMM) used a standard clustering algorithm called Partitioning Around Medoids (PAM) [31] to identify the clusters, with the average silhouette width [32] used to achieve the optimal number of clusters. PAM was chosen as the data was known to be noisy, and it is known to be more robust to outliers compared with the k-means approach [17]. PAM was run multiple times with the highest silhouette width solution chosen to ensure the robustness of the result. However, these clusters were too large based on the end-users' requirements for less than 100 organisations per peer group. There were substantial differences between the clusters for 2017 and 2018, which would result in large numbers of organisations being moved. In response, we developed two methods based on hierarchical clustering to create peer groups that satisfied this practical constraint, and a single method to reallocate organisations to new clusters while respecting the need for stability.

The DPMM was fitted using a Metropolis-within-Gibbs sampling algorithm as part of the *PRemiuM* package [23] in R. The Markov Chain Monte Carlo (MCMC) sampler was run for 200,000 iterations on 6 parallel chains to fit DPMM. This was done separately for both years to compute different PDMs. The goal of MCMC sampling is to approximate the Bayesian posterior distribution, with the specific goal in this context to accurately estimate elements of the PDM. In order to assess whether convergence had been achieved, the MCMC chains were compared to assess whether (a) the chains had been run long enough to avoid Monte Carlo error, and (b) whether the chains had converged to the same stationary distribution. The PDM estimates were compared between pairs of chains using scatter plots, and the distributions of PDM estimates were also compared using empirical density plots. If convergence was achieved, then points in the scatter plots should fall on the line of equality and the densities would look very similar. For the Dirichlet Process parameter α , the trace, cumulative mean, autocorrelation, and density plots were used to assess convergence. The (unnormalised) log-Posterior was also assessed for convergence using the same method as α .

Clustering with business constraints

The data-set in the case study presented the opportunity for the large organisation to construct data-driven peer groups for the purposes of organisational learning between organisations,

developing defensible benchmarks, and tailoring policy to differentiated groups of organisations. For these purposes, the peer groups were required to be at most 100 organisations in size and sufficiently stable over time. The concern for stability arose because when organisations were reallocated to a new peer group for a new year, relationships developed for organisational learning would be compromised. If policy is based on peer groups, frequent changes to peer groups cause discontinuity and uncertainty for stakeholders. For benchmarking in particular, there is no size requirement, however, frequent changing of a stakeholder's peer group from year to year could reduce trust in any benchmark based on within-group comparison.

Maximum cluster size constraint. The first method, called *kirigami-1*, uses the optimal clusters as an initial partition, then iteratively bisects partitions using the nodes of the hierarchical clustering tree, until all clusters have been cut into sub-clusters with number of observations contained, or cluster size, less than 100.

The second method, called *kirigami-2*, starts with all observations in singleton clusters and then proceeds in a similar manner as standard hierarchical clustering. The difference is that if a merger results in a cluster which violates the upper threshold, this merger is blocked and the other possible mergers are considered instead. It proceeds until there are only two clusters or all possible mergers have been exhausted.

See [S1 Appendix](#) for the full algorithm pseudo-codes.

Temporal stability constraint. To update the clusters in response to data for a new reporting period, a DPMM was again fitted to the new data to produce a corresponding dissimilarity matrix. We found that applying the same algorithm over successive years did not necessarily produce highly similar sets of peer groups. Possible reasons for this are due to noise in the data, although also potentially due to organisational 'drift' in terms of client demographics. These causes may be exacerbated by having clusters very close together.

To address this issue, we introduced a reallocation algorithm, which produces clusters with controllable levels of change in the peer groups between years. The reallocation algorithm works similarly to *kirigami-2* (see [S1 Appendix](#) for the pseudo-code). It first identifies the top $p \times 100\%$ of observations with largest silhouette widths, reassigns them to singletons, then proceeds with constrained hierarchical clustering as normal. In reallocation, there are two competing objectives: goodness-of-fit and stability. Goodness-of-fit is measured by clustering indices (see the next section for descriptions), while stability can be measured using cluster similarity indices. However, the values of the similarity indices are difficult to interpret in context.

We considered that under a stable clustering, if two observations were in the same cluster (i.e. connected) in the previous year, then they should also be connected in the current year in most cases. Therefore, a measure of stability is the proportion of connections in the previous year that are retained in the current year. From this we developed a simple index called Proportion of Connections Retained (PCR):

$$\text{PCR} = \frac{\text{No. of pairs connected in both years}}{\text{No. of pairs connected in the previous year}} \quad (1)$$

The PCR is the proportion of connections present the previous year, which are retained in the current year. A connection is defined here between a pair of observations as being in the same cluster.

The PCR is simpler to understand than the Rand Index [33] or Meila Index [34], and in the context of peer-grouping, clarifies the cost of instability for a new set of peer groups, which may result in many connections between organisations being broken. Furthermore, it implicitly includes the temporal order of the cluster solutions. A negative aspect to the PCR is that it

implicitly favours larger partitioning of entities into larger clusters, since merging clusters results in fewer lost connections.

To determine an optimal reclustering based on the current data (represented by the new dissimilarity matrix) and the previous peer groups, we adopted a grid-search approach. The reallocation algorithm was run for a number of different evenly spaced values from $p = 0$ to $p = 0.95$. Any potential clustering whose PCR was lower than the threshold was removed, and the clustering with the best goodness-of-fit index was selected. A threshold of 90% was chosen to represent a reasonable penalty, since this means that 1 in 10 connections would be broken each year by an observation being reallocated.

Clustering specifications. The kirigami methods are both based on hierarchical clustering, so we examined how group dissimilarity (as opposed to a single entity) should be calculated and how the optimal number of clusters should be determined. We considered four methods for calculating group dissimilarity, known as linkage methods: average [35], Ward [36], complete [37], and single [38] (see [S2 Appendix](#) for descriptions). We were specifically interested in whether the kirigami methods, combined with a particular linkage method, would produce better peer groups. These methods can all be efficiently calculated using the Lance-Williams algorithm, where during hierarchical clustering the dissimilarity for a merged cluster is a function of the dissimilarities being merged.

For calculating the optimal number of clusters, there are many measures in the literature [39]. We considered three measures which are often used: the average silhouette [32], the Calinski-Haribaz (CH) Index [40], and the Pearson-Gamma measure [41]. These measures all indicate the optimal number of clusters with their maximum value.

The silhouette method measures for each observation the average distance between the observation's own cluster, and the minimum average distance to another cluster. Hence, it measures the specificity of each observation to its own cluster. The average is calculated as the overall goodness-of-fit measure. The CH index measures the ratio of the between-cluster and within-cluster sum-of-squares, so a good clustering by this measure has a large amount of variation explained by the clusters, similar to an Analysis-of-Variance problem. Finally, the Pearson-Gamma measure calculates the Pearson correlation for a vector of all pairwise distances against a vector of ones and zeros, with 0 if the pair are in the same cluster or 1 if the pair are in different clusters. This measures the linear correspondence between distance and the cluster equivalence relation. Full details and equations are described in [S2 Appendix](#).

Identification of key variables and features

Since there were 14 input variables, clusters were located in a high-dimensional space. It was therefore a complex task for an analyst to understand what specific variables and features were important for a given cluster in distinguishing it from other clusters. We present tools based on well-known statistical methods to identify these features, and demonstrate their use in the case study data. Variables identified in this step are then communicated in the next step for the simplified fingerprint plots.

Variables in the data are often correlated, resulting in the main components of variation being non-aligned with the main axes. These directions are well known as the main outputs of principal component analysis (PCA) and can be useful for identifying differences between clusters, since variation between clusters is likely to align along them. In addition to pairwise scatter-plotting of variables to examine cluster differences, this visualisation technique is used on the principal component scores of the data.

To accompany the visualisation approach described above, we also considered automatic methods of identifying variables which distinguish the clusters via classification algorithms. To

identify specific variables of importance, a random forest algorithm [42] was run on the data with the cluster allocation as the response and the covariates as the predictors. This ranked the covariates in order of importance in distinguishing the clusters. We then trained the random forest algorithm on subsets of the data made up of two clusters to determine which variables were important in distinguishing between pairs of clusters. The benefit of the random forest importance is that because of the flexibility of the tree-based model, non-linear and interactive effects can be captured, as opposed to a logistic regression or other classification method. For a small number of trees (< 100) the importance measure varied each time the random forest was fitted due to monte carlo error, so the number of trees was increased until the measure was sufficiently stable to reliably order the variables in order of importance.

Clusters can differ in terms of location or spread, and the differences can occur for a single variable, multiple variables or principal components. The random forest can identify an importance in the variable, however, we sought tools which could discriminate between these situations.

To identify whether location or spread were discriminative factors between two clusters, Linear Discriminant Analysis (LDA), Mahalanobis Distance (MD) and Quadratic Discriminant Analysis (QDA) were all fitted to the data and their predictive accuracy computed. LDA performs well when the locations of the clusters are distant compared to the spreads of the clusters, yet performs very poorly when the locations are similar as the model assumes equal spread for all clusters. The Mahalanobis Distance (MD) was computed to contrast with LDA, since using it for prediction assumes the same location for both clusters but different spreads. Finally, QDA allows for both differences in location and spread. Therefore, it should have the best prediction and forms an upper bound on accuracy for the other methods.

Communication to stakeholders using visualisation

We sought to provide insights into both the data and the peer-grouping results. The two stakeholder groups with whom we wanted to communicate were analysts and outlet representatives. We assume analysts are familiar with standard visualisation tools such as boxplots and scatterplots. In contrast, outlet representatives may prefer less statistical visualisations, and so graphics that summarised the information simply and succinctly were required.

Furthermore, the two groups likely have different questions about the data and the peer-grouping. Typical analyst question is: “what makes a given peer group distinct from the other peer groups?” To address this question, standard boxplots and scatterplots of the key variables and principal components were constructed to communicate distinctions between pair groups in one and two dimensions.

Typical outlet representative questions are more centred around the outlet and the relationship with its peer group:

1. What is it about my outlet which makes peer group X right for me?
2. What do the ‘typical’ organisations look like for my peer group?

To address the questions of this stakeholder group, visualisations of the key variables were constructed which required as little background knowledge as necessary in order to communicate distinctive features of peer groups.

Since the data are transformed prior to clustering, it may be difficult for a non-technical audience to understand the transformed scale. Instead of plotting variables on the original scale or the transformed scale, the probability inverse transform (PIT) function [43] was used to transform each key variable in the data to be uniformly distributed between 0 and 1, essentially converting data-points into percentiles. The overall distribution of the data in each

variable was uniform, but for the subset of data in each cluster, this is not necessarily true. Thus, the cluster distributions are represented in terms of the cohort's data-distribution percentiles, which are readily understandable to the user.

Further simplifying was done by splitting the data into 5 ordered bins which each contained 20% of the observations. We opted to use a 5-bin system as a good compromise between simplicity and representing the cluster's distribution, and a familiarity from its common usage in rating scales. The resulting plot provides a means of understanding the distribution of peer groups compared to the cohort, without needing to know the underlying metric scale.

Results

Initial analysis of the summary statistic data in 2017 found that a number of the variables were highly correlated with each other. These high correlations were also seen in the 2018 data. Therefore, three variables which were highly correlated with other variables (based on the Variance Inflation Factor) were removed, resulting in 11 variables for clustering.

Summary of results

Overall, kirigami-2 (bottom-up clustering) tended to produce better clusters according to goodness-of-fit metrics, compared to kirigami-1 (top-down clustering) for almost all upper-size thresholds (Section: Maximum cluster size constraint). The CH index produced more reasonable cluster sizes and indicated clear preference on deciding the number of clusters. There was not a clear distinction between different linkage methods, however, the single method tended to produce singleton clusters and the Ward method tended to produce larger clusters (Section: Linkage method and number of clusters). For the 2017 data, three clusters were found using the upper size threshold of 100. Principal components and discrimination analyses showed a clear difference in centre between cluster 3 and the two other clusters, but the difference between clusters 1 and 2 was related to both centre and spread (Section: Understanding high dimensional clusters). The fingerprint plot method was able to partially communicate these differences between clusters, but the differences between clusters 1 and 2 were not as clear (Section: Visualising peer groups). The reallocation algorithm was able to find a compromise clustering solution between stability and goodness-of-fit when constructing new clusters for the 2018 data, using the 2017 peer groups as the original groupings, yet the relationship between stability and goodness-of-fit was not strictly monotonic (Section: Reallocation).

Maximum cluster size constraint

Using the 2017 data, the behaviour of the constrained algorithms was observed for different upper-size thresholds with respect to cluster goodness-of-fit indices, cluster sizes and number (see Fig 2). For the three indices, there was a monotonic increase for both kirigami-1 and kirigami-2 as the size threshold was relaxed, with kirigami-2 being almost always better than kirigami-1 for all indices. Once the size threshold was above 160, the methods produced the exact same result, which was two clusters of size 168 and 32, with slight differences in these numbers depending on the linkage method used. This behaviour was expected because these clusters correspond to the optimal clusters without constraint, so both methods are able to find this within the constrained solution space.

The greatest discrimination between the methods was seen for the CH index, where the kirigami-1 index was much lower for thresholds lower than 170 the index. This was explained by the number of clusters, with kirigami-1 producing very large numbers of very small clusters. Since every extra cluster was another 'degree-of-freedom', the CH harshly penalised kirigami-

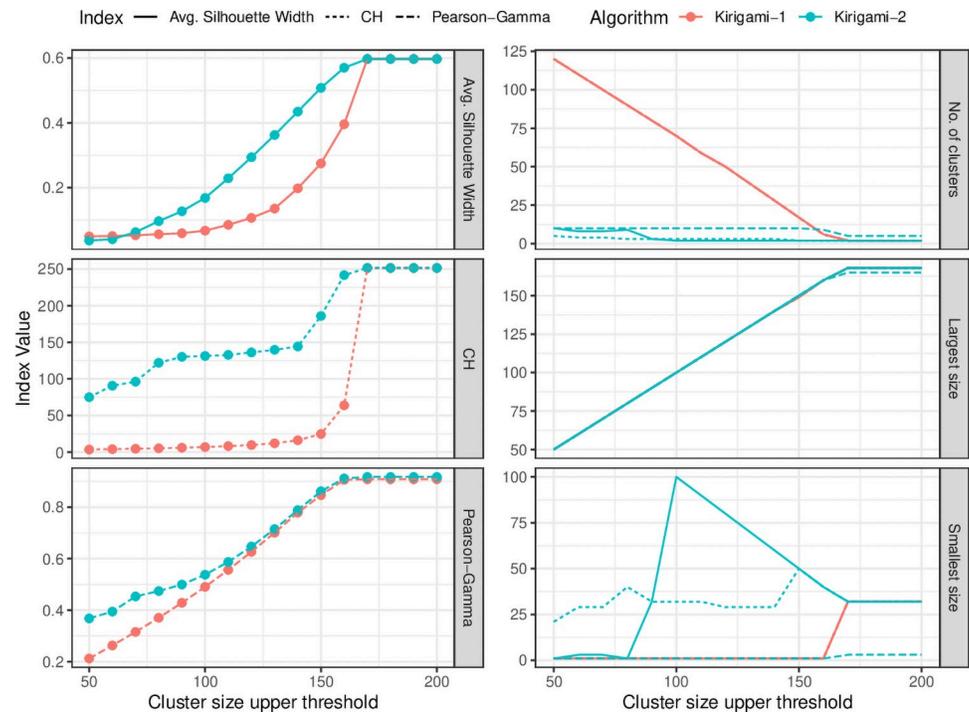


Fig 2. Statistics for constrained cluster solutions as a function of upper-size threshold. Index values (left column), cluster sizes and numbers (right column) for different values of the upper-size threshold for cluster size, for the two constrained clustering methods, kirigami-1 and kirigami-2. Cluster size and numbers are presented for solutions constructed using each of the three metrics, average silhouette width, CH, and Pearson-Gamma. Results here are presented for the average linkage method, however, the results did not vary markedly for other linkage methods.

<https://doi.org/10.1371/journal.pone.0251723.g002>

1. This contrasts with kirigami-2, which at most produced 10 clusters. There was a much lower discrimination for the Pearson-Gamma index, with only a small advantage for kirigami-2 when the threshold was less than 100.

In terms of size, the largest cluster was always equal to the threshold, except once the optimal clusters could be found within the constrained space. For kirigami-1, the smallest cluster size tended to be either 1, indicating one or more singleton clusters, or 32, corresponding to the optimal clusters, when the threshold was sufficiently high. For kirigami-2, the cluster numbers and smallest sizes were dependent on the goodness-of-fit index used. If the average silhouette width was used, then the size of the smallest cluster ranged from a singleton to 100, whereas if the Pearson-Gamma index was used, the smallest cluster only ranged between 1 and 6. When the CH index was used, smallest cluster size was between 21 and 50. This again shows the tendency for kirigami-2 and the CH index to prefer fewer, larger clusters, compared to the other indices and the kirigami-1 method, which promote many singleton or very small clusters.

Linkage method and number of clusters

For kirigami-1, the number of clusters was chosen by the goodness-of-fit index, whereas the goodness-of-fit index was only used in kirigami-2 to decide the initial clustering, and successive cuts increased the number of clusters to satisfy the upper-size threshold constraint. Therefore, it was of interest to investigate these indices' behaviours when changing the number of clusters for kirigami-1 specifically. The upper-size threshold was fixed at 100 for the analysis in this section.

Table 1. Comparison of clustering solution for the (agglomerative) hierarchical clustering algorithm and the two constraining algorithms kirigami-1 and kirigami-2. For brevity, only the optimal linkage methods are given for each of the three clustering indices. Solutions where the maximum cluster size exceeded the threshold of 100 are indicated in bold.

Method	Index	Best Linkage Method	Index Value	Max. Cluster size	Min. Cluster size	Number of Clusters
Hierarchical Clustering	CH	Average	251.44	168	32	2
	Average Silhouette Width	Ward	0.61	174	26	2
	Pearson-Gamma	Average	0.92	165	3	5
Kirigami-1	CH	Ward	9.84	100	1	75
	Average Silhouette Width	Ward	0.08	100	1	75
	Pearson-Gamma	Ward	0.51	100	1	75
Kirigami-2	CH	Average	131.32	100	32	3
	Average Silhouette Width	Ward	0.18	100	26	3
	Pearson-Gamma	Average	0.55	100	1	23

<https://doi.org/10.1371/journal.pone.0251723.t001>

As shown in Table 1, the clustering indices indicate that the constrained solutions from Kirigami-1 and Kirigami-2 are substantially worse than standard hierarchical clustering. As demonstrated in the previous section Kirigami-1 produces large numbers of very small clusters in order to constrain the size of the largest cluster to 100, resulting in poorer fits according to the clustering indices compared to Kirigami-2. By contrast, Kirigami-2 produces a much smaller number of clusters and larger clusters, depending on the index used to determine the optimal number of clusters.

As shown in Fig 3, the indices showed vastly different behaviours. The average silhouette width tended to reduce for larger numbers of clusters, with the highest value being 2 clusters for all linkage methods except for Ward, where the value was slightly higher for 3 clusters. The Pearson-Gamma index instead showed an increase for all linkage methods as the number of clusters increased, plateauing after around 6 clusters. The preference for 3 clusters was very clear in the CH index, with a sharp increase from 2 clusters to 3 and then a steady decline as more clusters were added.

The trees produced by the linkage methods were similar only for the first split and diverged quickly afterwards. The Ward linkage method tended to result here in larger clusters, whereas in the other linkage methods it was more common for very small or singleton clusters to break off from a larger cluster. Because the CH index penalises more for small clusters, the other linkage methods tended to produce poorer clusters according to this index. By contrast, the average silhouette width index indicated the Ward clustering was the worst for four or more clusters. The associated reduction from 0.17 to around 0.12 when the number of clusters was increased from 3 to 4 was actually the result of one cluster splitting into two new clusters, each of which was closer to the other clusters, resulting in a reduction in the silhouette widths of observations in the other two clusters.

Understanding high dimensional clusters

It was found to be very easy to distinguish cluster 3 from clusters 1 and 2 in terms of a few variables. Cluster 3's center was substantially separated from the other clusters, and it had relatively large spread (see Fig 4). It was much more difficult to identify differences between clusters 1 and 2, and standard plots were not useful for distinguishing between them. Using random forest importance, the most important covariates for discriminating between them were covariate 34 and covariate 10, and plotting these revealed only small differences in the centre but a large difference in spread. We termed this particular situation a 'shell structure', since the observations in cluster 1 are surrounded in multi-dimensional space by the

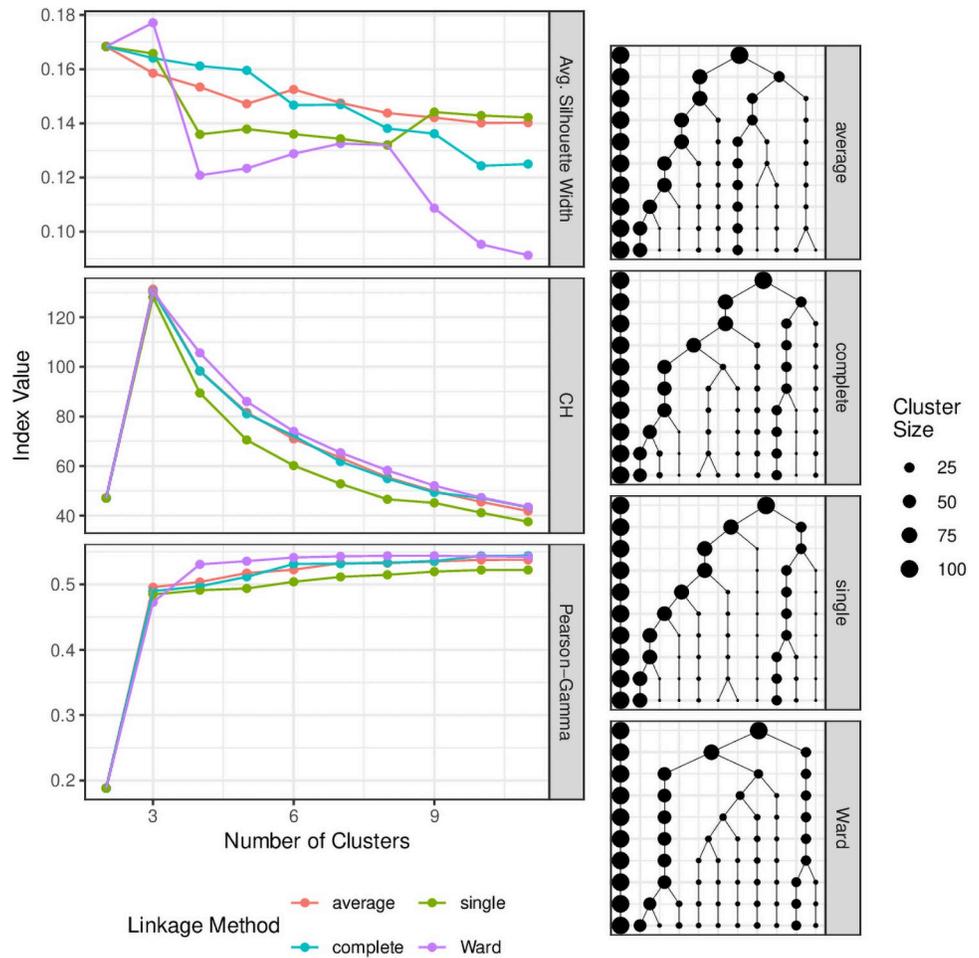


Fig 3. Goodness-of-fit and tree structures for different hierarchical clustering linkage methods. (left) Clustering goodness-of-fit indices plotted against the number of clusters for kirigami-2. The algorithm was performed using four different types of linkage method. (right) The hierarchical trees constructed using the kirigami-2 algorithm for the four linkage methods. The trees are represented vertically, with the highest level representing 2 clusters, and the lowest level representing 11 clusters after 10 bisections. Each point represents a cluster, with the point size representing the size of the clusters. Clusters are connected to their ‘parent cluster’ by vertical and diagonal edges.

<https://doi.org/10.1371/journal.pone.0251723.g003>

observations in cluster 2, which is much sparser. From a PCA analysis we found that clusters 1 and 2 had heavy overlap in the main directions of variation in the data, although cluster 1 was more concentrated for some principal components.

Computing the Mahalanobis distance from the centre of the two clusters, it was found that cluster 2’s observations tended to be further away than cluster 1’s (83.9% discrimination). Linear Discriminant Analysis also produced a similar level of discrimination (84.5%), indicating some difference between centre. Finally, a full quadratic discrimination analysis, which takes into account both centre and spread, was able to correctly classify between cluster 1 and 2 at 89% accuracy. This shows there are small differences in both centre and spread.

Visualising peer groups

The fingerprint plots (Fig 5) show clear differences between the clusters, especially for comparing clusters 1 and 2 with cluster 3. Clusters 1 and 2 are more similar, as the figure suggests, however, there is a clear difference for covariate 10 and response 1.

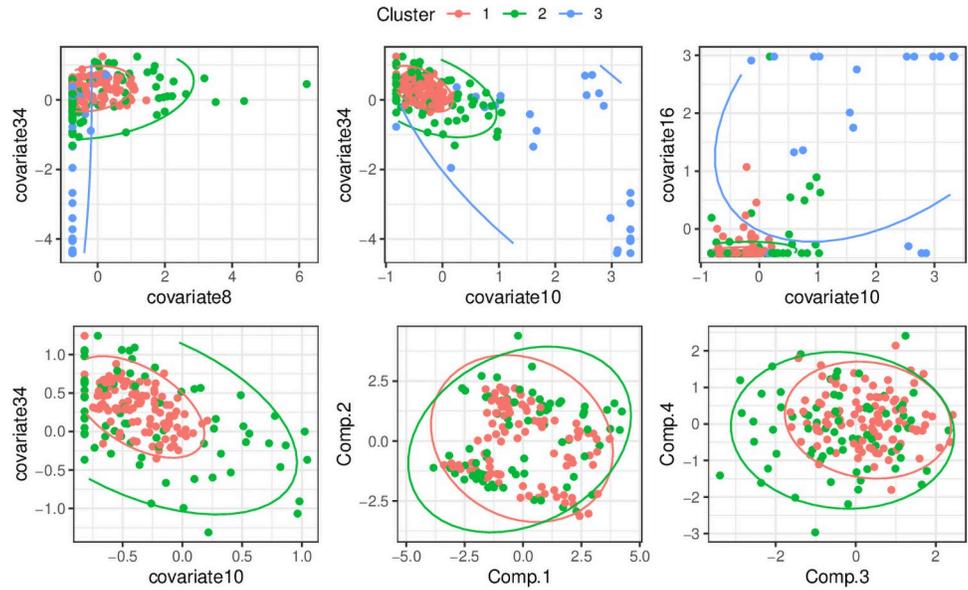


Fig 4. Pairwise scatter plots showing the difference between the clusterings for the 2017 data. The top three scatter plots show all three clusters for three pairs of variables from the data, selected by random forest importance. The bottom three show only clusters 1 and 2, and show (bottom-left) the top two variables for discriminating between the two clusters and the first four principal component scores (bottom middle and left). Ellipses represent the 95% probability regions, and are calculated based on the empirical means and covariance matrices of the clusters for a multivariate Gaussian distribution. The comparison of clusters 1 and 2 between covariate 10 and covariate 34 clearly shows the shell structure.

<https://doi.org/10.1371/journal.pone.0251723.g004>

Using the entire data distribution as a reference, the fingerprint plot can also show the individual cluster distributions. Although cluster 1 contains organisations in all quintiles for response 1, a greater proportion of its organisations tend to have high (highest 60–80%) and very high (highest 80–100%) response 1 values. Cluster 3 by contrast contains no large or very large organisations, but a large proportion of very small organisations (lowest 0–20%).

The graphic also provides justification to the stakeholder that they have been correctly assigned to their cluster. The featured observation, chosen at random, is seen to be in typical

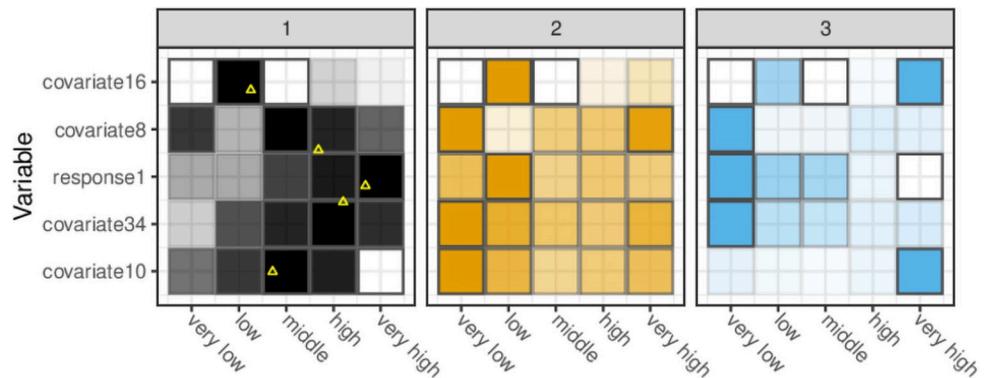


Fig 5. Fingerprint plots for the top 5 explanatory variables in the data-set, as chosen by the random forest discrimination. Data are separated out into quintiles, and then apportioned out into clusters. The relative proportion of observations in each quintile is depicted by the transparency of the colour. Quintiles are given descriptors to guide the reader. A single observation in Cluster 1 is depicted with a yellow triangle to demonstrate how the corresponding stakeholder can interpret the graphic.

<https://doi.org/10.1371/journal.pone.0251723.g005>

bins for that cluster. It is clear that it would be a poor fit for the other two clusters, particularly for covariate 34.

Reallocation

With the peer groups established for 2017, the posterior dissimilarity matrix for 2018 was used to reallocate observations to new peer groups when they were of sufficiently high silhouette distance from their 2017 peer group. As with 2017, Ward's linkage method was used along with the CH index.

As expected, the stability of the cluster solutions measured by PCR reduced as more observations were assigned to be reallocated, and the goodness-of-fit, measured by the CH index, increased (see Fig 6). This increase was not monotonic in either case, as might be expected given increasing the reallocation threshold relaxed the stability constraint. This was seen when the proportion to be reallocated was around 17%, where the PCR briefly rose before continuing to decrease. Investigation revealed this was because the number of clusters briefly reduced from 3 to 2, resulting in an increase in the PCR. The number of clusters then changed again

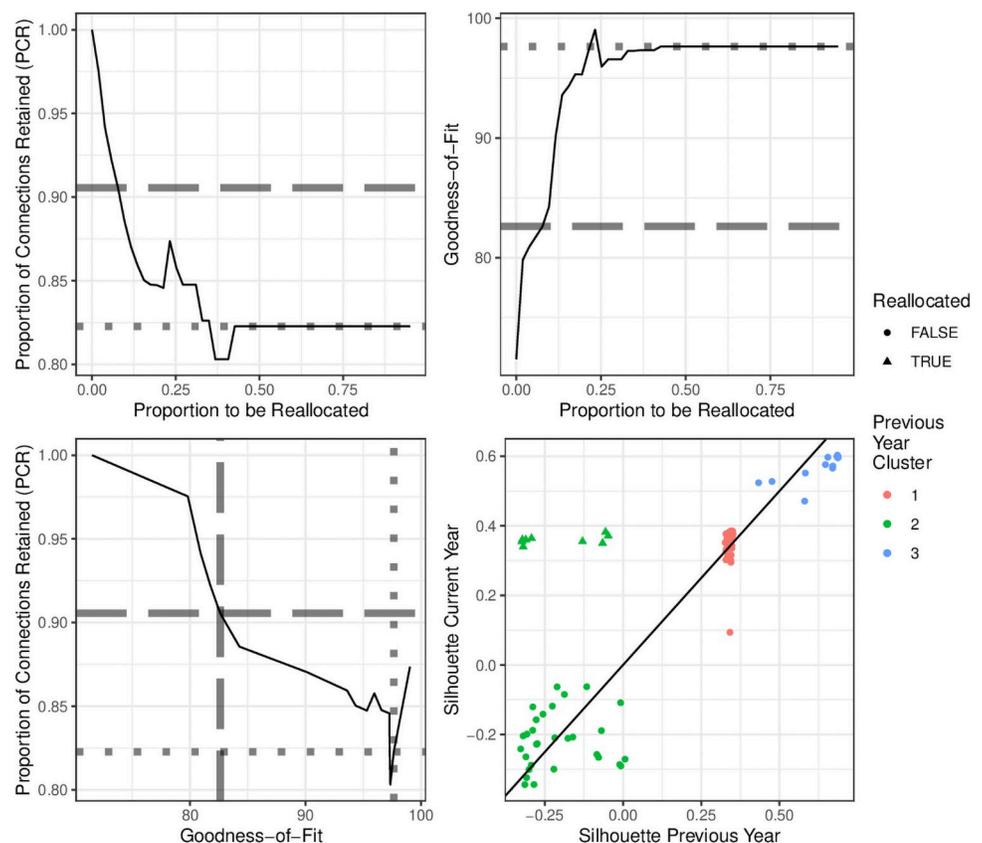


Fig 6. Goodness-of-fit and stability for different values of the reallocation threshold. (top-left) The stability, as measured by the proportion of connections retained (PCR). (top-right) The CH index goodness-of-fit value for partitions created using different proportions of reallocated entities (bottom-left). The stability plotted against the goodness-of-fit, providing a line to evaluate the trade-off between the two objectives. In each of these three plots, the horizontal and vertical guidelines are provided to show the PCR and goodness of fit values for the reallocation clustering solution (dashed) based on a lower stability threshold of 90% and a clustering solution with no stability constraint (dotted). (bottom-left) The silhouette distance in the previous year as a predictor of the silhouette distance in the following year.

<https://doi.org/10.1371/journal.pone.0251723.g006>

back to 3, with the original third cluster coalescing with a subset of observations from a larger cluster to form a new group.

The goodness-of-fit increased monotonically with the reallocation threshold with the exception of a single spike, which also coincided with the number of clusters changing from 3 to 2. The increases plateaued quickly, as once the reallocation threshold was above 50%, the optimal (size-constrained) solution could be reached and further relaxing of the stability constraint would not improve the fit. There was a strong penalty for the stability threshold of 90% based on the goodness-of-fit index. The CH-index for the reallocation cluster solution based on this penalty was 82.6, which was closer to the index for a clustering with no reallocation (71.5) than the solution with no constraint (97.6).

The reallocation solution for the $PCR < 90\%$, which was the solution chosen to be the basis for the 2018 peer groups, had 3 clusters that corresponded well to the 2017 clusters. There were a total of 9 reallocated observations, all of which were from the 2017 peer group 2, and moved to both peer group 1 (7) and to peer group 3 (2). As shown in Fig 6, the silhouette distances were similar between years for observations that remained in the same cluster. By contrast, the silhouette distances of the 9 reallocated observations, all from the 2017 peer group 2, increased substantially. The silhouette distances for peer group 2 in both 2017 and 2018 were negative, a result of the shell structure that peer group 2 adopts around peer group 1.

Discussion

The results show that the upper-size constraint forces a split to form clusters 1 and 2. This causes a substantial loss of goodness-of-fit, and the reallocation method's constrained solution was also lower than the unconstrained goodness-of-fit. While there may be improvements to be made on the kirigami and reallocation algorithms to find better constrained solutions, it does indicate that these are strong constraints which substantially reduce the distinction between clusters. It is therefore important to consider from the perspective of the stakeholders how important the business constraints are, relative to the importance of having well-defined, data-driven peer groups.

The CH index showed a clear maximum at its optimal number of clusters, so was the most useful index to use in conjunction with the constrained clustering for these data. The average silhouette was more sensitive to small differences in clustering solutions between the linkage methods. Given these were only a small proportion of the data, this indicates that the average silhouette is dominated by a minority of observations. The Pearson-Gamma method preferred many very small clusters, indicating it did not penalise for greater cluster number sufficiently. Very small clusters which are dissimilar to all other larger clusters represent a challenge for the Pearson-Gamma index, which is the correlation between the dissimilarity and the cluster membership over all pairs of observations. The merger of a very small cluster with the nearest cluster results in a reduction in the Pearson-Gamma index, but having very small clusters is undesirable for peer-grouping. The CH index is therefore recommended here, but it is likely to be data-dependent as to which index provides the clearest preference for number of clusters. Context may dictate that a given index is more appropriate.

The kirigami-2 was overall a better performing algorithm for these data than the kirigami-1, regardless of the goodness-of-fit index used. As shown in the results, the key reason for this was that the number of clusters produced by kirigami-1 was very high. This is because kirigami-1 needs to unfold many cuts in the hierarchical clustering tree in order to reduce the size of clusters, which resulted in many singleton clusters being created in the process. If the data were less noisy and clusters more clearly defined in the tree, then kirigami-1 would perform better, but this was not tested here. On the other hand, the kirigami-2 algorithm produces its

own tree, which sidesteps the need to use the unconstrained tree to find the clustering solution. In practice therefore, kirigami-2 is likely to produce better defined peer groups.

The associated tree constructed by kirigami-2 contains numerous potential peer-grouping solutions with different numbers of clusters, all of which are subject to the required size constraint. Providing multiple alternative solutions allows a practitioner greater flexibility in decision-making, given that peer-groups can often be subject to other requirements such as legal criteria, traditional business practices, or prior organisational relationships. Since these requirements may not be easily coded into logical constraints on the solution, the kirigami-2 method may be embedded within a larger peer-grouping decision framework or extended to allow for soft constraints, more complex objective functions, or greater human input and judgement.

It was expected that in reallocation of observations, the goodness-of-fit would increase monotonically as the stability threshold was relaxed, but this was not strictly the case. Instead there was some fluctuation when the stability threshold was close to being large enough to permit the optimal unconstrained clustering. It is therefore recommended that this relationship is considered when using reallocation, rather than simply automatically running the algorithm with a chosen threshold value of the PCR. This may be particularly important when the chosen threshold is small, meaning a greater proportion of observations are being reallocated.

Even with the tools demonstrated in this work, it remains difficult to explain complex, high dimensional clustering. Domain knowledge may help to construct terms for the principal components of the data. It may be possible to explain principal components in terms of abstract qualitative factors. For example, component 4 in the case study had high loadings for number of clients and the number of sessions per client, in the same direction. Therefore, component 4 might be described as 'organisation size' to a stakeholder. For example, the shell structure seen is very difficult to explain and justify as it defies basic intuition and 1D representation, as was seen for this case study.

The two peer-groups which make up the shell structure observed in the data were the result of the upper-size constraint, otherwise they would have been merged, but the constraint is not necessary for this behaviour to occur. Comparing the performances of discriminant analysis algorithms and principal component analysis can dissect the differences between clusters in centre and spread, providing simple numerical summaries of the discrepancies between complex multidimensional shapes. It nevertheless seems counter-intuitive when using single- and two-dimensional visualisations that two observations in the wider peer-group should be in the same peer-group, when they are both quite different to each other and more similar to observations in the more dense peer-group. The reasoning for the clustering solution is on the scale of the entire cohort rather than individual observations, because it is favourable to have a dense peer-group of close observations at the expense of a small number of observations in the sparser peer-group. This stands as an example of where clustering results, particularly in high dimensions, are very difficult to explain in the peer-grouping context where clear, non-technical explanations are desired.

A peer-group with a large spread may mean organisations within are highly dissimilar to each other, but this is not necessarily a negative outcome with respect to benchmarking and inter-organisational learning. For example, in the case of a shell structure, it could be reasonable to describe these organisations as atypical compared to the typical organisations in the denser, central peer-group. Here, being atypical may either be a source of similarity and facilitate exchange between peer-groups, or indicate that the peer-group model may not be a valid method for these organisations to learn from one another. Similarly, being atypical may also indicate that benchmarking is inappropriate for these organisations as they are not sufficiently similar to others. Therefore, practitioners should evaluate by other means whether a wide

peer-group can be used for a chosen purpose, but wideness does not preclude inter-organisational learning and may signal where benchmarking is inappropriate in the first place.

The critical challenges in merging the problem of peer-grouping with the problem of clustering are defining the business or contextual constraints on the clustering, and in the particular case of multi-dimensional data, developing a means for stakeholders to understand these constructed groups. The first challenge is more suited to be solved by computational tools, but the constraints imposed on the clustering solution must be translated into thresholds or rules. The second challenge is perhaps more difficult as it requires a complex, high-dimensional shape of a peer group to be reduced to single or two dimensions. While it is tempting to instead construct peer groups with simplicity in mind, fully data-driven clustering helps to overcome analyst biases about which variables are most important for distinguishing organisations. Instead, the second challenge presents the opportunity to develop novel visualisations, such as the fingerprint plot, which can bridge the gap between high-dimensional clustering and human understanding.

Conclusion

For organisational peer-grouping efforts to be effective, the resulting groups need to reflect real differences and similarities between organisations, and also conform to the requirements of the business context. Using a clustering approach instead of other methods such as expert elicitation has the clear benefit in being data-driven, thus reducing the risk of bias in deciding what separates organisations into peer groups. Methods developed in this work allow clustering solutions to be more amenable to the real-world peer-grouping context by meeting the business constraints of size and stability. The resulting clusters had complex, multidimensional shapes, so classification algorithms were utilised to automatically identify important features and combinations of features, and identify differences between clusters (in terms of centre and spread). Fingerprint plots, which are similar to multivariate heat maps, were developed to assist stakeholders in visualising peer group characteristics and comparing different groups. The methodology presented in this paper represents a novel contribution to the peer-grouping literature by incorporating both context and data in peer group construction, and visualising statistical results to be easily interpretable in the business context.

Supporting information

S1 Appendix. Pseudocode for constrained clustering algorithms. Complete descriptions of the steps in the three algorithms presented (kirigami-1, kirigami-2, and reallocation).
(PDF)

S2 Appendix. Descriptions of clustering specifications. Descriptions of linkage criteria and cluster goodness-of-fit indices used in this research.
(PDF)

Author Contributions

Conceptualization: Jessica Cameron, Paul P. -Y. Wu, Kerrie Mengersen.

Data curation: Jessica Cameron.

Formal analysis: Daniel W. Kennedy, Jessica Cameron, Paul P. -Y. Wu.

Funding acquisition: Kerrie Mengersen.

Investigation: Jessica Cameron, Paul P. -Y. Wu.

Methodology: Daniel W. Kennedy, Jessica Cameron.

Project administration: Kerrie Mengersen.

Supervision: Paul P. -Y. Wu, Kerrie Mengersen.

Validation: Daniel W. Kennedy, Jessica Cameron.

Visualization: Daniel W. Kennedy, Jessica Cameron.

Writing – original draft: Daniel W. Kennedy, Jessica Cameron.

Writing – review & editing: Daniel W. Kennedy, Jessica Cameron, Paul P. -Y. Wu, Kerrie Mengersen.

References

1. Klassen A, Miller A, Anderson N, Shen J, Schiariti V, O'Donnell M. Performance measurement and improvement frameworks in health, education and social services systems: a systematic review. *International Journal for Quality in Health Care*. 2009; 22(1):44–69. <https://doi.org/10.1093/intqhc/mzp057> PMID: 19951964
2. Segev E, Raveh A, Farjoun M. Conceptual maps of the leading MBA programs in the United States: core courses, concentration areas, and the ranking of the school. *Strategic Management Journal*. 1999; 20(6):549–565. [https://doi.org/10.1002/\(SICI\)1097-0266\(199906\)20:6%3C549::AID-SMJ39%3E3.0.CO;2-F](https://doi.org/10.1002/(SICI)1097-0266(199906)20:6%3C549::AID-SMJ39%3E3.0.CO;2-F)
3. Hanning BWT. Length of stay benchmarking in the Australian private hospital sector. *Australian Health Review*. 2007; 31(1):150. <https://doi.org/10.1071/AH070150> PMID: 17266499
4. Dimick JB, Staiger DO, Birkmeyer JD. Ranking Hospitals on Surgical Mortality: The Importance of Reliability Adjustment. *Health Services Research*. 2010; 45(6p1):1614–1629. <https://doi.org/10.1111/j.1475-6773.2010.01158.x> PMID: 20722747
5. Wilson D, Piebalga A. Performance Measures, Ranking and Parental Choice: An Analysis of the English School League Tables. *International Public Management Journal*. 2008; 11(3):344–366. <https://doi.org/10.1080/10967490802301336>
6. Pope DG. Reacting to rankings: Evidence from “America’s Best Hospitals”. *Journal of Health Economics*. 2009; 28(6):1154–1165. <https://doi.org/10.1016/j.jhealeco.2009.08.006> PMID: 19818518
7. Xu R, Wunsch D. *Clustering*. Wiley-IEEE Press; 2009.
8. Porter ME. *Competitive strategy: techniques for analyzing industries and competitors* / Porter Michael E. Free Press New York; 1980.
9. Cool KO, Schendel D. Strategic Group Formation and Performance: The Case of the U.S. Pharmaceutical Industry, 1963–1982. *Management Science*. 1987; 33(9):1102–1124. <https://doi.org/10.1287/mnsc.33.9.1102>
10. Atterton J, Thompson N, Carroll T. Mentoring as a mechanism for improvement in local government. *Public Money & Management*. 2009; 29(1):51–57. <https://doi.org/10.1080/09540960802617368>
11. Woods R, Artist S, O'Connor G. Learning in Australian local government: A roadmap for improving education & training. *Commonwealth Journal of Local Governance*. 2016; p. 108–126. <https://doi.org/10.5130/cjlg.v0i18.4845>
12. Wood R, Tan S, Ryan R. *Councils learning from each other: An Australian case study*, Australian Centre of Excellence for Local Government; 2015.
13. van Ewijk E. Mutual Learning in Dutch-Moroccan and Dutch-Turkish Municipal Partnerships. *Tijdschrift voor economische en sociale geografie*. 2012; 103(1):101–109. <https://doi.org/10.1111/j.1467-9663.2011.00698.x>
14. Bridges J, Mazevska D, Pearse J. Designing a nationally acceptable system of hospital peer grouping. *Australian Health Review*. 2000; 23(1):193. <https://doi.org/10.1071/AH000193> PMID: 10947602
15. Stefos T, LaVallee N, Holden F. Fairness in prospective payment: a clustering approach. *Health services research*. 1992; 27(2):239–261. PMID: 1592607
16. Hurley RG. Identification and Assessment of Community College Peer Institution Selection Systems. *Community College Review*. 2002; 29(4):1–27. <https://doi.org/10.1177/009155210202900401>

17. Hastie T, Tibshirani R, Friedman J. *The Elements of Statistical Learning*. Springer New York; 2009. Available from: <https://doi.org/10.1007/978-0-387-84858-7>.
18. Zafra-Gómez JL, López-Hernández AM, Hernández-Bastida A. Evaluating financial performance in local government: maximizing the benchmarking value. *International Review of Administrative Sciences*. 2009; 75(1):151–167. <https://doi.org/10.1177/0020852308099510>
19. Duan X, ming Jin Z. Positioning decisions within strategic groups. *Management Decision*. 2014; 52(10):1858–1887. <https://doi.org/10.1108/MD-08-2013-0415>
20. Gómez J, Orcos R, Palomas S. Do strategic groups explain differences in multimarket competition spillovers? *Strategic Organization*. 2016; 15(3):367–389. <https://doi.org/10.1177/1476127016665358>
21. Frühwirth-Schnatter S. *Finite Mixture and Markov Switching Models*. Springer Series in Statistics. Springer New York; 2006.
22. Sethuraman J. A CONSTRUCTIVE DEFINITION OF DIRICHLET PRIORS. *Statistica Sinica*. 1994; 4(2):639–650.
23. Liverani S, Hastie DI, Azizi L, Papatomas M, Richardson S. PReMiuM: An R Package for Profile Regression Mixture Models Using Dirichlet Processes. *Journal of Statistical Software*. 2015; 64(7):1–30. <https://doi.org/10.18637/jss.v064.i07> PMID: 27307779
24. Bradley PS, Bennett KP, Demiriz A. *Constrained K-Means Clustering*. Redmond, WA: Microsoft Research; 2000.
25. Okabe M, Yamada S. Clustering Using Boosted Constrained k-Means Algorithm. *Frontiers in Robotics and AI*. 2018; 5. <https://doi.org/10.3389/frobt.2018.00018> PMID: 33500905
26. Namitha K, Kumar GS. A Framework for Monitoring Clustering Stability Over Time. In: *Emerging Research in Computing, Information, Communication and Applications*. Springer Singapore; 2019. p. 467–477. Available from: https://doi.org/10.1007/978-981-13-5953-8_39.
27. Ntoutsis E, Spiliopoulou M, Theodoridis Y. Tracing cluster transitions for different cluster types. *Control Cybern*. 2009; 38:239–259.
28. Oliveira M, Gama Ja. MEC –Monitoring Clusters’ Transitions. In: *Proceedings of the 2010 Conference on STAIRS 2010: Proceedings of the Fifth Starting AI Researchers’ Symposium*. NLD: IOS Press; 2010. p. 212–224.
29. Warton DI, Hui FKC. The arcsine is asinine: the analysis of proportions in ecology. *Ecology*. 2011; 92(1):3–10. <https://doi.org/10.1890/10-0340.1> PMID: 21560670
30. R Core Team. *R: A Language and Environment for Statistical Computing*; 2020. Available from: <https://www.R-project.org/>.
31. Kaufman L, Rousseeuw PJ. Clustering by means of Medoids. In: Dodge Y, editor. *Statistical Data Analysis Based on the L_1 -Norm and Related Methods*. The address of the publisher: Elsevier Science Pub. Co.; 1987. p. 405–416.
32. Rousseeuw PJ. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*. 1987; 20:53–65. [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7)
33. Rand WM. Objective Criteria for the Evaluation of Clustering Methods. *Journal of the American Statistical Association*. 1971; 66(336):846–850. <https://doi.org/10.1080/01621459.1971.10482356>
34. Meilä M. Comparing clusterings—an information based distance. *Journal of Multivariate Analysis*. 2007; 98(5):873–895. <https://doi.org/10.1016/j.jmva.2006.11.013>
35. Sokal RR, Michener CD. A statistical method for evaluating systematic relationships. *University of Kansas Science Bulletin*. 1958; 38:1409–1438.
36. Ward JH. Hierarchical Grouping to Optimize an Objective Function. *Journal of the American Statistical Association*. 1963; 58(301):236–244. <https://doi.org/10.1080/01621459.1963.10500845>
37. Defays D. An efficient algorithm for a complete link method. *The Computer Journal*. 1977; 20(4):364–366. <https://doi.org/10.1093/comjnl/20.4.364>
38. Sibson R. SLINK: An optimally efficient algorithm for the single-link cluster method. *The Computer Journal*. 1973; 16(1):30–34. <https://doi.org/10.1093/comjnl/16.1.30>
39. Vendramin L, Campello RJGB, Hruschka ER. On the Comparison of Relative Clustering Validity Criteria. In: *Proceedings of the 2009 SIAM International Conference on Data Mining*. Society for Industrial and Applied Mathematics; 2009. Available from: <https://doi.org/10.1137/1.9781611972795.63>.
40. Calinski T, Harabasz J. A dendrite method for cluster analysis. *Communications in Statistics—Theory and Methods*. 1974; 3(1):1–27. <https://doi.org/10.1080/03610927408827101>

41. Halkidi M, Batistakis Y, Vazirgiannis M. On Clustering Validation Techniques. *Journal of Intelligent Information Systems*. 2001; 17(2/3):107–145. <https://doi.org/10.1023/A:1012801612483>
42. Breiman L. Random Forests. *Machine Learning*. 2001; 45(1):5–32. <https://doi.org/10.1023/A:1010933404324>
43. Angus JE. The Probability Integral Transform and Related Results. *SIAM Review*. 1994; 36(4):652–654. <https://doi.org/10.1137/1036146>