

Flexible parametric bootstrap for testing homogeneity against clustering and assessing the number of clusters

Christian Hennig¹ · Chien-Ju Lin²

Accepted: 8 March 2015 / Published online: 11 June 2015
© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract There are two notoriously hard problems in cluster analysis, estimating the number of clusters, and checking whether the population to be clustered is not actually homogeneous. Given a dataset, a clustering method and a cluster validation index, this paper proposes to set up null models that capture structural features of the data that cannot be interpreted as indicating clustering. Artificial datasets are sampled from the null model with parameters estimated from the original dataset. This can be used for testing the null hypothesis of a homogeneous population against a clustering alternative. It can also be used to calibrate the validation index for estimating the number of clusters, by taking into account the expected distribution of the index under the null model for any given number of clusters. The approach is illustrated by three examples, involving various different clustering techniques (partitioning around medoids, hierarchical methods, a Gaussian mixture model), validation indexes (average silhouette width, prediction strength and BIC), and issues such as mixed-type data, temporal and spatial autocorrelation.

Keywords Cluster validation · Mixture model · Distance-based clustering · Markov chain · Mixed-type data · Spatial autocorrelation · Presence–absence data

Mathematics Subject Classification 62H30 · 62F03 · 62F40

1 Introduction

Cluster analysis is about finding groups of objects in data. Cluster analysis is a key area of data analysis with applications virtually everywhere where data arise. For example, the present paper features applications in social science, biogeography and medicine. Cluster analysis methods have been developed since the 1950s in various subject areas including statistics, mathematics, computer science and machine learning, biology, psychology, and geoscience. The field of cluster analysis is therefore characterised by a lack of unification. Some cluster analysis approaches are based on probability models for each cluster, others are based on density estimation, even others are based on distance measures and discrete mathematics and do not involve probability at all. As a result, the probabilistic behaviour of cluster analysis methods is often not well understood.

In the present paper, we treat two key issues in cluster analysis, namely the question whether a dataset is clustered at all, and the selection of an appropriate number of clusters. We present a general principle to address these issues, which can be applied to various approaches to cluster analysis.

A common approach to the selection of an appropriate number of clusters k is via cluster validation indexes. Cluster validation indexes are statistics that can be computed for a given clustering of a dataset and measure the quality or “validity” of the clustering. Various validation indexes have been proposed in the literature, for example the Calinski–Harabasz index, the average silhouette width (ASW), Sugar and James’s distortion, see, e.g. [Milligan and Cooper \(1985\)](#), [Sugar and James \(2003\)](#), [Arbelaitz et al. \(2012\)](#), [Xiong and](#)

✉ Christian Hennig
c.hennig@ucl.ac.uk

Chien-Ju Lin
chienju@mrc-bsu.cam.ac.uk

¹ Department of Statistical Science, University College London, Gower Street, London WC1E 6BT, UK

² MRC Biostatistics Unit, Cambridge Institute of Public Health, Forvie Site, Robinson Way, Cambridge Biomedical Campus, Cambridge CB2 0SR, UK

Li (2014). The indexes are computed for clusterings for a range of candidate values for k . It is usually recommended to select the k that optimises either the index or a change of the index between $k - 1$ and k clusters, depending on the index. These recommendations are often either purely heuristic, or based (often rather loosely) on theory using simple probability models for each cluster such as the Gaussian distribution. Some criteria for finding the number of clusters such as the Bayesian information criterion (BIC) for mixture model-based clustering (Fraley and Raftery 1998) are based on probability theory in a more consistent way, but for the purpose of the present paper they can be interpreted as cluster validity criteria as well.

We assume here that the researcher has decided which cluster analysis method and which cluster validity index to use. Our attitude regarding these decisions is that different cluster analysis methods and different validity indexes correspond to different “cluster concepts”, which may be of interest in different applications. There is no uniquely optimal choice of a combination of these, but the researcher rather needs to decide what cluster concept is required in a specific application. For example, it may be required that all objects are, on average, represented as precisely as possible by the centroid object of the cluster to which they are assigned, which can lead, depending on the distance concept involved, to the k -means or the “Partitioning Around Medoids” (PAM) clustering method, and the Calinski and Harabasz index, see Kaufman and Rousseeuw (1990), Calinski and Harabasz (1974), or the researcher may be interested in finding latent subpopulations distributed approximately according to Gaussian distributions, leading to model-based clustering with Gaussian mixtures and the BIC, Fraley and Raftery (1998). However, the approach taken here does not assume that clusters are in general identified with “clustering” probability models for data subsets such as Gaussian components of a mixture model [which do not always have characteristics that are expected of clusters such as small within-cluster distances and separation, see Hennig and Liao (2013)]. The question whether the data can be explained by a homogeneous probability model for “non-clustering”, or on the other hand whether there is an evidence for “real” clustering, is treated as separate from what constitutes a cluster. The philosophy of clustering involved here has been outlined in Hennig and Liao (2013).

The main idea of the present paper is that parametric bootstrap can be used to investigate the distribution of the given validation index, simulating from a model for homogeneous data, i.e. for the absence of “real” clustering. The validation index can then be used as a test statistic for testing homogeneity against a clustering alternative (this yields a test for each candidate k for which the index is computed, which need to be aggregated to a single homogeneity test), and the simulated null distribution can also be used to calibrate the

validity index by comparing its value on the dataset against what is expected under the null model. We argue that this is a better foundation for a decision about the number of clusters than the heuristics behind the standard recommendations in most of the literature.

Although it is rarely seen in practice, the idea of setting up a hypothesis test of a null hypothesis modelling “no clustering” for cluster validation indexes is not new. For example, it is mentioned in Chapter 4 of Jain and Dubes (1988). Jain and Dubes (1988) mention the “random graph”, “random cluster label” and “random position” (uniform/Poisson process distribution) hypothesis. Tests for some standard null hypotheses including the normal distribution are cited in Bock (1996) with a focus on a proper theoretical derivation of the distribution of the test statistics.

Often, however, rejecting such simple null hypotheses is not evidence for clustering, because there may be more structure in the data than what these null models assume, for example temporal or spatial dependence. “Parametric bootstrap” refers to sampling from a parametric model with parameters estimated from the data (Efron and Tibshirani 1994). In this paper, we propose using the parametric bootstrap to sample from null models that capture the non-clustering structure in the data for testing homogeneity against clustering, and for calibrating validity indexes. The parametric bootstrap allows us to use models that are more complex and less “theory-friendly” than the simple models mentioned above. Efron and Tibshirani (1994) treat the parametric bootstrap somewhat briefly, because they argue that a main advantage of the bootstrap is that inference can be constructed without parametric assumptions, for which the nonparametric bootstrap, i.e. exploring the distribution of a test statistic by sampling from the observed empirical distribution, was constructed. But for the aim of testing homogeneity against clustering, the empirical distribution is not suitable, because sampling from the empirical distribution will generate datasets with the same clustering characteristic as the original dataset to be analysed. Potential homogeneity can only be explored based on a model for non-clustering. Therefore, the non-parametric bootstrap is not an option here.

Using parametric bootstrap for testing homogeneity against a clustering alternative and calibration of cluster validity indexes is a very general principle. It can basically be used in every clustering problem together with any clustering method and any validation index (as long as there is enough computational power to run clustering and validation index lots of times). But every situation requires a new tailor-made null model, which means that there is no straightforward out-of-the-box way to run this approach. Readers who want to apply it need to design, implement and estimate the parameters of their own null model, capturing the structural features of their datasets that do not indicate clustering. The best way of demonstrating how to do this is to show

examples. After Sect. 2, in which the general idea is stated, it is applied to three different datasets. Section 3 is about mixed-type data for socio-economic stratification containing continuous, ordinal and nominal variables, the latter with categories carrying somewhat stronger than purely nominal information. The clustering method is PAM and the validation index used is the ASW (Kaufman and Rousseeuw 1990). Section 4 is about a dataset giving the methadone dosages taken by patients over 180 days, involving temporal autocorrelation. PAM was applied once more, but also compared with Complete and Average Linkage clustering, and cluster validity was assessed by the prediction strength (PS) (Tibshirani and Walther 2005), which explores cluster stability based on resampling. In Sect. 5, we analyse a presence–absence dataset of snail species on Aegean islands where the problem is to cluster the species distribution ranges. The null model takes into account spatial autocorrelation. Following Hennig and Hausdorf (2004), the dataset was clustered using Gaussian mixture model-based clustering with the BIC (Fraley and Raftery 1998) after defining a distance measure between distribution ranges and running a Multidimensional Scaling (MDS; Cox and Cox 2001). The example explores the use of the parametric bootstrap approach together with model-based clustering methods and demonstrates that the parametric bootstrap adds important information to the standard usage of model-based clustering and the BIC. Section 6 gives a concluding discussion.

Sections 3 and 5 give a nod to two predecessors of the current paper. Hennig and Hausdorf (2004) already introduced parametric bootstrap tests for homogeneity against clustering using the specific null model that will be applied in Sect. 5, although it did not consider the estimation of the number of clusters. The general principle proposed here was already applied in an ad hoc-fashion in Hennig and Liao (2013), where the dataset of Sect. 3 was analysed. Section 3 improves on the null model used in Hennig and Liao (2013).

2 The general setup

The general principle of the present paper is outlined theoretically in this section, and will then be illustrated by examples.

Given is a set of observations $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ from some set of possible objects \mathcal{X} . The observations can be characterised in various ways, normally either by p variables or by an $n \times n$ -dissimilarity matrix. Then there is a clustering method C so that $C(\mathbf{X}, k) = \{C_1, \dots, C_k\}$ with $k \in K \subseteq \mathbb{N}$, and, for $i = 1, \dots, k$: $C_i \subseteq \mathbf{X}$. In many cases, C will be a partitioning method assuming that $C_i \cap C_j = \emptyset$ for any $i \neq j$ and $\bigcup_{i=1}^k C_i = \mathbf{X}$, and $K = \{2, \dots, n\}$, but this is not required in general. Furthermore given is a validity index V , so that $V(\mathbf{X}, C(\mathbf{X}, k)) \in \mathbb{R}$ measures the quality of $C(\mathbf{X}, k)$

in some sense. We assume w.l.o.g. that a larger value of V implies a better cluster quality, at least as long as clusterings with the same k are compared.

The null model $\mathcal{P}_0 = \{P_\theta : \theta \in \Theta\}$ is a set of probability distributions P_θ on \mathcal{X}^n (equipped with a suitable σ -algebra) with the interpretation that the distributions P_θ model a situation that is interpreted as homogeneous in the sense of “absence of clustering”. The set Θ can also be rather general; in Sect. 5, for example, Θ involves the full empirical distributions of both the sizes of species and the numbers of species present in the regions. Basically Θ should capture all structural information as far as it cannot be interpreted as “clustering”, which may involve features that are usually referred to as nonparametric, such as full marginal distributions of some variables. Often the n observations will be modelled as i.i.d., but this again is not required.

Let $T_n : \mathcal{X}^n \mapsto \Theta$ be an estimator of θ . For a fixed number of clusters $k \in K$ and a fixed set of observations \mathbf{X} , a parametric bootstrap test is defined by estimating the distribution Q_k , which is the distribution of $V(\mathbf{X}, C(\mathbf{X}, k))$ under $P_{T_n(\mathbf{X})}$, by drawing m bootstrap datasets $\mathbf{X}_1, \dots, \mathbf{X}_m$ from $P_{T_n(\mathbf{X})}$. The bootstrapped p value for testing \mathcal{P}_0 is then

$$\hat{p}_k = \frac{|A_k| + 1}{m + 1}, \text{ where} \tag{1}$$

$$A_k = \{\mathbf{X}_i : V(\mathbf{X}_i, C(\mathbf{X}_i, k)) \geq V(\mathbf{X}, C(\mathbf{X}, k))\},$$

so that a low \hat{p}_k implies that it is very unlikely, under Q_k , that $V(\mathbf{X}^*, C(\mathbf{X}^*, k))$ is as large or larger as the observed validity $V(\mathbf{X}, C(\mathbf{X}, k))$, which therefore is an evidence for a stronger clustering than what is expected under Q_k .

This defines a homogeneity test for each $k \in K$. These need to be aggregated into a single homogeneity test. This can be done by defining an overall aggregated p value

$$\hat{p} = \frac{|A^*| + 1}{m + 1}, \text{ where} \tag{2}$$

$$A^* = \left\{ \mathbf{X}_i : \sum_{k \in K} \tilde{p}_k(\mathbf{X}_i) \leq \sum_{k \in K} \hat{p}_k \right\},$$

where $\tilde{p}_k(\mathbf{X}_i)$ is the analogue of \hat{p}_k based on $V(\mathbf{X}_i, C(\mathbf{X}_i, k))$, i.e. with $\mathbf{X}_{m+1} = \mathbf{X}$:

$$\tilde{p}_k(\mathbf{X}_i) = \frac{|\tilde{A}_i| + 1}{m + 1}, \text{ where}$$

$$\tilde{A}_i = \{\mathbf{X}_j : j \neq i, V(\mathbf{X}_j, C(\mathbf{X}_j, k)) \geq V(\mathbf{X}_i, C(\mathbf{X}_i, k))\}.$$

This means that aggregating the tests for all $k \in K$ is effectively based on averaging the p values or, equivalently, ranks of $V(\mathbf{X}, C(\mathbf{X}, k))$ among the bootstrapped samples over k .

An optimal value of k can be found by maximising a calibrated V :

$$\hat{k} = \arg \max_{k \in K} \frac{V(\mathbf{X}, C(\mathbf{X}, k)) - EV_k}{SV_k}, \quad (3)$$

$$EV_k = \frac{1}{m} \sum_{i=1}^m V(\mathbf{X}_i, C(\mathbf{X}_i, k)),$$

$$SV_k = \sqrt{\frac{1}{m-1} \sum_{i=1}^m (V(\mathbf{X}_i, C(\mathbf{X}_i, k)) - EV_k)^2}.$$

The interpretation is that this is the k for which $V(\mathbf{X}, C(\mathbf{X}, k))$ gives the best validity compared to what is expected under Q_k .

All the information from the parametric bootstrap can be visualised by plotting $V(\mathbf{X}, C(\mathbf{X}, k))$ together with all $V(\mathbf{X}_j, C(\mathbf{X}_j, k))$ against k (“bootstrap validity plot”), which will be done in the following sections. Actually, often this plot will be so expressive that computing the formal outcomes Eqs. (1), (2) and (3) does not add much information.

There are alternative ways to define the tests and the estimation of k based on the parametric bootstrap. Instead of Eq. (1), $V(\mathbf{X}, C(\mathbf{X}, k))$ could be standardised as in Eq. (3), and the p value could then be computed from a Gaussian distribution, although the Gaussian approximation cannot be proved to work in the generality required here. Instead of averaging p values in Eq. (2), one could also average raw values of V (implicitly assuming that these are meaningfully comparable over $k \in K$), or one could use a Bonferroni-adjustment of the lowest \hat{p}_k , $k \in K$, which can be very conservative but may work well if $V(\mathbf{X}, C(\mathbf{X}, k))$ for the best k is expected to stand out clearly. A comparison of these options is left to future work.

The parametric bootstrap tests proposed here require the specification of a null model, but they do not require the explicit specification of alternative models. The “clustering alternative”, against which the homogeneity null model is tested, is implicitly defined by the choice of V . The “effective alternative” are distributions P on \mathcal{X}^n under which the distribution of V is stochastically larger than under \mathcal{P}_0 . The choice of V therefore defines the meaning of “clustering” against which the homogeneity null hypothesis is tested.

A general limitation of parametric bootstrap is that the distribution $P_{T_n(\mathbf{X})}$ is usually interpreted to represent the whole of \mathcal{P}_0 . \hat{p} will be anti-conservative as a p value for testing \mathcal{P}_0 to the extent that other distributions in \mathcal{P}_0 exist that are both compatible with the observed data and tend to deliver larger values of V . Theoretical analysis of this problem is impossible in general and probably very tedious if possible at all in most specific situations and will therefore not be done here. The validity of significant test results therefore relies on the quality of the estimator T_n and the assumption that different

values of θ (at least as long as they are still compatible with the data) do not tend to yield vastly different values of V .

3 Socio-economic stratification (mixed-type data)

3.1 Data

Hennig and Liao (2013) analysed a dataset from the 2007 US Survey of Consumer Finances. There were $n = 17,430$ individuals and 8 variables (no missing values were in the dataset):

- $\log(x + 50)$ total amount of savings x as of of last month (treated as continuous),
- $\log(x + 50)$ total income x of 2006 (treated as continuous),
- years of education between 0 and 17; this is treated as ordinal (level 17 means “graduate school and above”),
- number of checking accounts that one has; this is ordinal with 6 levels (corresponding to no/1/2/3/(4 or 5)/(6 or more) accounts,
- number of savings accounts, coded as above,
- whether or not one has life insurance (binary, i.e. ordinal),
- housing, nominal with 9 levels: “neither owns nor rents”, “inapplicable”, “owns or is buying/land contract”, “pays rent”, “condo”, “co-op”, “townhouse association”, “retirement lifetime tenancy” and “own only part”,
- occupation class, nominal with 7 levels (from 0 to 6): “not working for pay”, “managerials and professionals”, “white collar workers and technicians”, “lower-level managerials and various professions”, “service workers and operators”, “manual workers and operators”, “farm and animal workers”.

The aim was to use clustering methods for socio-economic stratification [see Hennig and Liao (2013) for background]. An interesting issue, which is addressed by the approach of the present paper, is whether such a stratification is rather an artificial (although potentially useful) partition of a rather homogeneous population, in which case there are no clear boundaries between social classes or strata (note that the term “homogeneous” here does not refer to social equality).

3.2 Clustering method and validation index

Hennig and Liao (2013) settled for the PAM clustering method (namely for its large sample version CLARA) and the ASW validation index (Kaufman and Rousseeuw 1990), arguing that social strata should be defined by low within-cluster distances rather than components of a mixture model.

PAM tries to find k centroid objects in \mathbf{X} and assigns all objects in \mathbf{X} to the closest centroid so that the sum of distances

of every object to its cluster's centroid is minimised. The ASW averages standardised differences between the average distance of every object to the closest cluster to which it is not classified and the average distance to all objects of the cluster to which it is classified. The ASW can be between -1 and 1 . Values larger than 0 indicate that objects have on average lower distances within their own cluster than to their neighbouring cluster, which can be seen as minimum requirement for a distance-based clustering. Larger values of the ASW are better and Kaufman and Rousseeuw (1990) recommend to maximise the ASW for finding the best value of k , but this is not based on an analysis of what changes of the ASW can be expected when increasing k . This can be explored by the parametric bootstrap approach.

The distance measure was in principle a Euclidean distance for which the ordinal variables were used with standard Likert coding (1, 2, 3, ...), and the nominal variables were coded by binary dummies for the categories. However, variables were standardised with specific standardisation schemes in order to balance the contribution of the different types of variables to the clustering in an appropriate way. The account number variables were weighted down because they were comparably less important than the others, and the dummy variables for housing were weighted in such a way that the effective distance treated "owns" and "pays rent" as the extremes of the scale with the other categories in between, see Hennig and Liao (2013) for details.

3.3 Non-clustering structure

For running the parametric bootstrap, a null model needs to be defined that captures the non-clustering structure in the data. This requires some judgement by the researchers, because it depends on what constitutes a "significant clustering" in the given application. Before defining the null model, here is some discussion of what the non-clustering structure is.

A first thing to realise is that dependence between variables may lead certain clustering methods into building clusters with approximate local independence (i.e. independent within clusters). For standard latent class clustering of nominal variables, local independence is a standard assumption (e.g. Hennig and Liao 2013), as well as for k -means clustering if this is written down as a maximum likelihood (ML)-method for spherical Gaussian clusters. PAM and the ASW do not formally assume local independence. But if a number of variables are strongly correlated with a continuous transition without clear cluster boundaries between low values on all variables and high values on all variables on at least a subset of the observations, this subset does not make a suitable cluster according to most distance-based clustering methods including PAM/ASW, because there would be very large distances within the cluster between objects that are low or high, respectively, on all variables. Researchers

may well be interested in splitting up such subsets for practical purposes such as information reduction, but the resulting clustering may not be interpreted as "real" in the sense that there is no separation and therefore no "natural" cluster boundaries. Therefore, dependence between variables is seen as non-clustering structure here.

Furthermore, the information about the categories of the variables housing and occupation is somewhat stronger than nominal, which may cause some structure in the data which does not contribute to clustering that is interpretable as "real".

The modelling of the marginal distributions of the variables is a subtle issue. We regard the marginal distribution of nominal variables as not carrying clustering information, because a low relative frequency of certain categories cannot be interpreted as a cluster-defining "gap" between other categories. Therefore, the null model should reproduce the marginal distribution of the nominal variables as "non-clustering structure". But the situation is different for continuous distributions. The marginal distribution of a continuous variable can have various modes and gaps between them, which can be taken as indicating "real" clustering boundaries indeed. This means that a model is needed for a marginal distribution of the continuous variables which can be interpreted as non-clustering. Ordinality means that there is no metric distance between the categories, which indicates that low frequency-categories should not be interpreted as "gap between clusters", similar to the situation for nominal variables. Therefore, we will treat the marginal distribution of ordinal variables as non-clustering structure and therefore as a parameter.

In Hennig and Liao (2013), the Gaussian distribution was used as marginal null distribution for the continuous variables. A Gaussian distribution can properly be interpreted as "non-clustering", but it may be too restricted. The null model may be rejected not because there is a real clustering, but because the real marginal distribution is non-normal, for example skew. A more flexible way of modelling non-clustering structure is to use a general unimodal distribution. Furthermore, a special feature of the continuous variables is that there are a number of individuals with zero savings (7434) and/or zero income (5). Particularly, the large group of zero savings individuals causes strong non-normality of the marginal distribution. It is a matter of judgement whether such a group of individuals sharing the same value on a variable alone is seen as indicating a real clustering or not. For the present paper, we decided that we do not want to interpret this as indicating clustering, and therefore it needs to be incorporated in the null model. Note that using a Gaussian distribution as the null marginal as in Hennig and Liao (2013) implicitly amounts to interpreting this deviation from a Gaussian as an evidence for real clustering. Obviously the existence of a large number of Americans with zero savings is something real; whether this is interpreted as "clustering",

though, depends on whether one imagines the existence of such a group in a “classless” society. Here we take the point of view that even in such an idealised homogeneous classless society, people are still free to spend all their savings and a considerable number of them will do that.

3.4 Null model

The null model has to be defined in such a way that its parameters can be estimated by the data. Some aspects of the estimation will be ad hoc so that it is not needed to set up a full new estimation theory. It is therefore useful to think about the definition of the null model and the estimation of the parameters in one go.

We choose \mathcal{P}_0 to be based on a latent Gaussian model, the outcomes of which are transformed to match the marginal distributions mentioned above. The correlations of such a model can be estimated from ordinal data by using the technique of polychoric correlations (Drasgow 1986).

Assume $\mathbf{x}_1, \dots, \mathbf{x}_n$ to be i.i.d. distributed, with $\mathbf{x}_1 = (x_{11}, \dots, x_{1p})^t$ generated from a latent p variate ($p = 8$) Gaussian random variable $\mathbf{z} = (z_1, \dots, z_p)^t \sim \mathcal{N}(\mathbf{0}_p, \Sigma)$, Σ being a correlation matrix, i.e. with diagonals equal to one. For the continuous variables ($j = 1, 2$) assume $P\{x_{1j} = \log(50)\} = p_j > 0$ (remember that the continuous variables are $\log(50 + x)$ -transformed, so $\log(50)$ refers to zero savings or income) and the conditional distribution $\mathcal{L}(x_{1j}|x_{1j} > \log(50))$ to be unimodal with continuous density. Let G_j be the cdf of the full distribution of x_{1j} , and assume that $x_{1j} = G_j^{-1}(\Phi(z_j))$, where Φ is the cdf of the standard Gaussian distribution.

For an ordinal variable ($j > 2$; see below for nominal variables) x_{1j} with h categories c_1, \dots, c_h let $-\infty = u_{j0} < u_{j1} < \dots < u_{jh} = \infty$ be a sequence of Gaussian quantiles so that $x_{ij} = c_g \Leftrightarrow z_j \in (u_{j(g-1)}, u_{jg}]$, $g = 1, \dots, h$.

Furthermore, assume that the categories of the nominal variables ($j = 8, 9$) x_{1j} are ordered with an unknown ordering, and the true ordering is defined according to the average true correlations between the dummy variables indicating the categories of x_{1j} and the variables that were originally ordinal or continuous. Given this true ordering, z_j is related to x_{ij} as for the ordinal variables above.

3.5 Null model parameter estimation

In order to estimate the polychoric correlations, i.e. the matrix Σ , the continuous variables are treated as ordinal by splitting them up into 10 ordered categories each, of approximately the same size. The true orders of the nominal variables ($j = 8, 9$) can be estimated by computing the average sample correlations between the dummy variables indicating the categories of x_{1j} and the variables that were originally ordinal or continuous. With that, all the variables

are ordinal and Σ can be estimated as in Drasgow (1986). $-\infty = u_{j0} < u_{j1} < \dots < u_{jh} = \infty$ can be estimated so that they reproduce the observed marginal distributions of the ordinal and nominal variables.

The distributions G_j can be estimated by using the empirical probability for $x_{1j} = \log(50)$, and by fitting a kernel density estimator to the observations with $x_{ij} > \log(50)$ making sure that the estimated density is unimodal. In practice, this has been done by using the “density” function in R with default settings, and if this was not unimodal, by increasing the bandwidth by steps of the originally selected bandwidth divided by 20 until the resulting density is unimodal.

3.6 Parametric bootstrap

In order to explore the distributions $Q_k = P_{T_n(\mathbf{X})}(V(\bullet, C(\bullet, k)))$, repeat m times:

1. Generate n i.i.d. observations $\mathbf{z}_1^*, \dots, \mathbf{z}_n^*$ from $\mathcal{N}(\mathbf{0}_p, \hat{\Sigma})$.
2. Transform them into $\mathbf{X}^* = \{\mathbf{x}_1^*, \dots, \mathbf{x}_n^*\}$ according to Sect. 3.4, using the estimated distributions \hat{G}_j and Gaussian quantiles \hat{u}_{jg} .
3. Compute a distance matrix for the objects in \mathbf{X}^* as in Hennig and Liao (2013).
4. For $k \in K$, cluster \mathbf{X}^* by PAM and compute and store the ASW, $V_{qk} = V(\mathbf{X}^*, C(\mathbf{X}^*, k))$, $q = 1, \dots, m$.

3.7 Results

The results with $m = 500$ and $K = \{2, \dots, 10\}$ are shown in Fig. 1. The real dataset produces clearly outstanding ASW values overall except for $k = 2$ and $k = 5$. The average \hat{p}_k value is smaller than for all datasets generated from the null model, so $\hat{p} = \frac{1}{501}$, the smallest possible value, a strong rejection of the homogeneity model.

For $k = 2$ the raw ASW reaches its maximum, so according to the standard recommendation $k = 2$ should yield the best clustering. But for $k = 2$ the real dataset does not yield a significantly better clustering than the null model, as opposed to most other values of k . This means that for this dataset, the standard recommendation is misleading. A higher k gives a better ASW in comparison to what can be expected under the null model. The best calibrated ASW values as in Eq. (3) are 4.476 for $k = 3$ and 4.481 for $k = 8$, which leads to $\hat{k} = 8$ as recommended value. $k = 3$ is about as good.

In Hennig and Liao (2013), where the Gaussian distribution was used as the null marginal for the continuous variables, the conclusions were mainly the same, but the null model here improved the achieved ASW values to some extent, with the effect that now the real dataset no longer

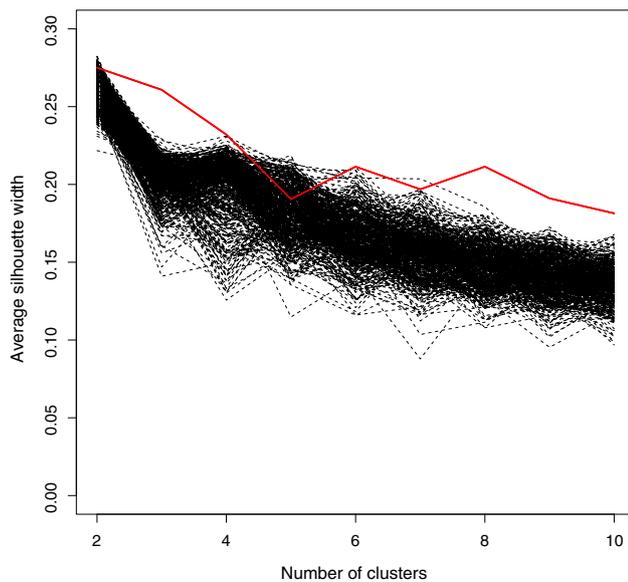


Fig. 1 ASW values for socio-economic stratification data (*red*) and 500 bootstrap samples (*black*) for $k = 2, \dots, 10$. (Color figure online)

produces the largest ASW-value for every single $k > 2$. This shows that assessment in Hennig and Liao (2013) was somewhat over-optimistic regarding the strength of the clustering, but the main conclusion is confirmed.

4 Methadone patients (Markov time series)

4.1 Data

Lin (2014) analysed dosage pattern data from 314 Taiwanese heroine addicts receiving methadone. For every methadone patient, the dataset contains records of the methadone dosage taken on each of the first 180 days from the beginning of the methadone therapy. There are six ordered dosage categories 1–6. Also there are missing values, meaning that the patient did not show up for obtaining methadone on a certain day. The data are shown in Fig. 2. Cluster analysis was done partly exploratory and partly for making the communication about the dosage patterns simpler. A clustering therefore can be useful regardless of whether there is some real clustering pattern in the data or not, but a really meaningful clustering, if it exists, is of medical interest.

Again, clusters should be characterised by small distances within clusters. Lin (2014), Lin et al. (2015) defined a distance measure in which missing values were treated as an additional category not having ordinal information, and the other categories were treated, following expert advice, in such a way that a change from one category to any other category was treated as fairly substantial, even between neighbouring categories, which means that categories were

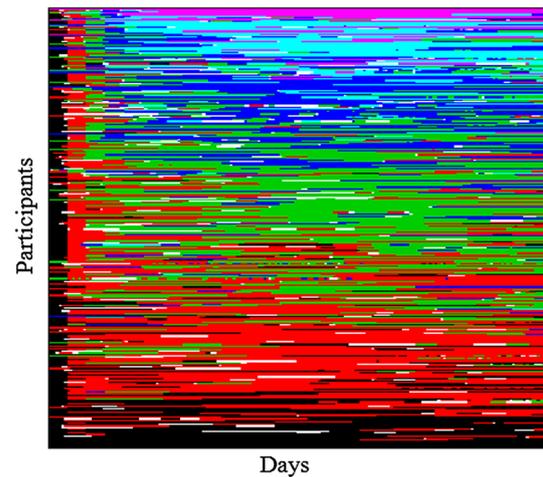


Fig. 2 Heatplot of methadone data. Dosage categories are 1 (*black*), 2 (*red*), 3 (*green*), 4 (*blue*), 5 (*light blue*), 6 (*violet*), missing (*white*). Patients are ordered according to average dosage. (Color figure online)

treated as carrying some compromise between ordinal and nominal information.

4.2 Clustering method and validation index

Several clustering methods were compared, namely PAM, Average Linkage and Complete Linkage hierarchical clustering (Kaufman and Rousseeuw 1990). Cluster validation was done by two methods, namely the ASW and the PS (Tibshirani and Walther 2005). We focus on the latter one here. The PS measures the stability of the clustering. The idea of the PS is as follows: The dataset is split into two equally sized parts b times. Every time, both halves are clustered into k clusters by the method that is also applied to the original dataset. For each of the two clusterings computed on the two halves, a prediction rule is created for the observations of the other half of the data. For any pair of observations in the same cluster in the same half it is then checked whether or not they are predicted into the same clustering by the clustering on the other half. If this is the case, their co-membership was correctly predicted. The PS is defined by averaging the proportions of correctly predicted co-memberships for the weakest clusters in each of the $2b$ halves. The prediction rule recommended in Tibshirani and Walther (2005) is to predict observations into the cluster with the closest cluster mean, which is appropriate for k -means. For PAM, we chose the closest centroid, for Average Linkage the minimum average distance and for Complete Linkage the minimax distance.

The PS is not calibrated for properly comparing values for different k ; it can be expected that larger values of k make it more difficult to achieve a high PS, particularly because the PS is determined by the least stable cluster. Because of this, Tibshirani and Walther (2005) suggest to choose the largest

k for which the PS is larger than 0.8. Instead, because PS values for fixed k are comparable, the parametric bootstrap idea can be applied. This amounts to b sample splits for each k and each of the m bootstrap samples, which makes this very computer-intensive.

4.3 Non-clustering structure

Here is some background knowledge about the methadone dosages. Obviously, for a given patient, the different days cannot be treated as independent. Once a week, every 7 days, the methadone patients get a new prescription. On all other days, the patients are free to use a smaller dosage than indicated on their prescription, so that there are occasional changes on these days, but most changes happen on day 1, 8, 15 etc. (“prescription days”). Some patients go to two different doctors and obtain two different prescriptions, which makes their dosages more flexible. The dataset does not include prescription data; we only know what dosage the patient took, but not what the prescription was; so it is not possible to use the prescription for setting up the null model. In any case, even on prescription days, old prescriptions are often renewed, and if there is change, it is mostly by only one category. There is much more change on the earliest prescription days when doctors do not yet have much experience with the patients than later. Furthermore, most patients start on the lowest dosage. There is no obvious connection between previous dosages and missing values; previous missing value behaviour is much more informative about missing values in the future than are observed dosages.

4.4 Null model

Modelling the time series of dosages of a single patient as a Markov chain ignores long range dependence. We choose this approach anyway, because other features of the data are more striking, and a Markov chain with different transition probabilities for prescription days and “normal days” already requires a large number of parameters. Some transition probabilities, particularly between non-neighbouring dosages, are very small and difficult to estimate with the limited amount of data.

The dosages of the n patients (ignoring missing values for the moment) are modelled in \mathcal{P}_0 as i.i.d. Markov chains with different transition probability matrices between the six dosages for (a) prescription day 2 (prescription day 1 defines the initial distribution of dosages), (b) prescription day 3, (c) all further prescription days combined, and (d) all normal days combined. Furthermore, we assume an unrestricted distribution of missingness patterns (i.e. a series of indicators of whether a patient is missing or not over all 180 days), from which one is drawn i.i.d. for each patient, independently of

the patient’s dosages. This effectively rules out missingness patterns as sources for “real” clustering; whatever is observed can be reproduced by the null model. This seems appropriate because real clusterings are only of medical use if they correspond to patterns of dosages. Missingness does not allow an interpretation in terms of the methadone needs of a patient.

4.5 Null model parameter estimation

The four transition probability matrices can be estimated in a straightforward manner by empirical transition probabilities in the given situations. The initial dosage of a patient can be drawn randomly from the empirical distribution of initial dosages. The distribution of missingness patterns can also be estimated directly as its empirical distribution.

4.6 Parametric bootstrap

Repeat m times:

1. For n i.i.d. observations in the bootstrap sample \mathbf{X}^* :
 - (a) Draw an initial dosage from their empirical distribution.
 - (b) Generate a sequence of 180 dosages using on each day the appropriate estimated transition probabilities.
 - (c) Independently of the sequence of dosages, draw a missingness pattern (i.e. a set of days with missing values) from the empirical distribution of missingness patterns, and make the corresponding days missing.
2. Compute a distance matrix for \mathbf{X}^* as explained in Lin (2014).
3. Repeat b times:
 - (a) Split \mathbf{X}^* into two equally large subsets.
 - (b) Cluster both subsets by PAM for $k \in K$, and by Average Linkage and Complete Linkage; for the latter two methods, clusterings for $k \in K$ can be obtained by cutting the trees at the appropriate height.
 - (c) For all pairs of observations within clusters, check whether the co-memberships are correctly predicted as explained in Sect. 4.2.
4. For the $2b$ clusterings, the three clustering methods, and $k \in K$, compute and store the PS, i.e. the minimum (over the k clusters) proportion of correctly predicted co-memberships in the cluster. Average these to get $V_{qk} = V(\mathbf{X}^*, C(\mathbf{X}^*, k))$, $q = 1, \dots, m$.

4.7 Results

The results based on $m = 500$, $b = 50$ are shown in Figs. 3 and 4. For Average Linkage and Complete Link-

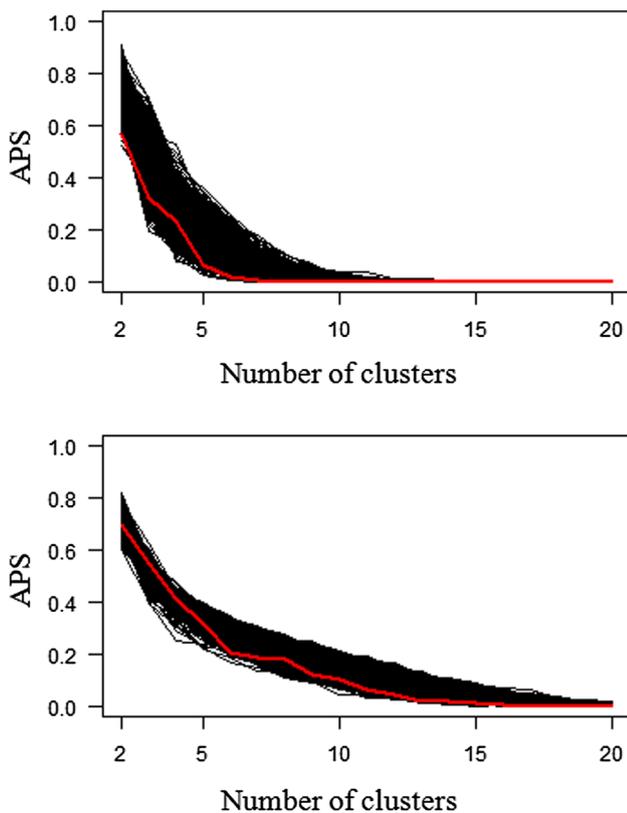


Fig. 3 PS values for methadone data (red) and 500 bootstrap samples (black) for $k = 2, \dots, 20$, Average Linkage clustering (above), Complete Linkage clustering (below). (Color figure online)

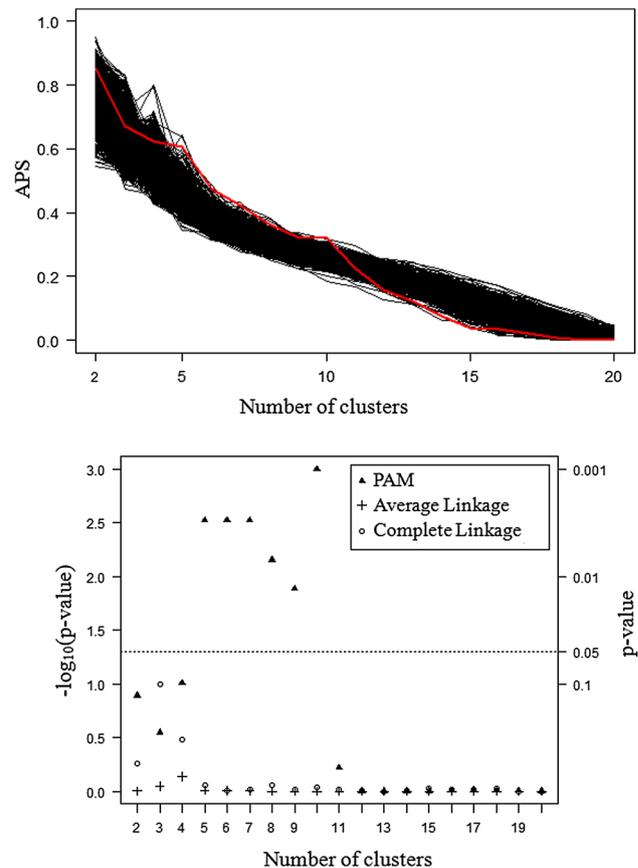


Fig. 4 PS values for methadone data (red) and 500 bootstrap samples (black) for $k = 2, \dots, 20$, PAM clustering (above); bootstrapped p values \hat{p}_k for all three clustering methods (below). (Color figure online)

age, the original methadone dataset yields stability (PS) values that are on average even below the values from the null model. Obviously, based on these clustering methods, there is no evidence for real clustering. For PAM, looking at some values of k , namely $k \in \{5, \dots, 10\}$ the PS values look significantly higher than those from the null model (with $k = 10$ looking best; see the right side of Fig. 4), although they do not clearly stick out. However, avoiding cherry-picking the best values of k , averaging \hat{p}_k values over k according to Eq. (2) gives $\hat{p} = 0.475$ due to some very low PS values for higher k , and therefore here again there is no evidence for real clustering (strictly speaking, having tried out three clustering methods, even a further adjustment for multiple testing would be required).

Once more, the default recommendation, which for the PS is “take the largest k for which $PS > 0.8$ ”, turns out to be misleading. The real dataset achieves this for PAM and $k = 2$ only, but this does not seem to be the best value compared to the null model, and $PS > 0.8$ can be achieved by the null model for $k = 2$ for all clustering methods, and for PAM occasionally even for $k = 3$.

5 Distribution ranges of snail species (spatial dependence)

5.1 Data

The third example dataset is a binary dataset giving presence–absence information on 80 species of snails for 34 Aegean islands (Cyclades; Hausdorf and Hennig 2005). The dataset is available under the name “kykladspecreg” in the R-package “prabclus”. Clustering of such species distribution ranges aims at finding “biotic elements”, groups of species sharing specific areas, which are connected to certain hypotheses about speciation (Hausdorf and Hennig 2003). We here interpret the observations \mathbf{x}_i , $i = 1, \dots, n$ as sets of islands for which the species is present.

5.2 Clustering method and validation index

Following Hausdorf and Hennig (2003) and Hennig and Hausdorf (2004), Kulczynski-dissimilarities were computed between different species:

$$d(\mathbf{x}_1, \mathbf{x}_2) = 1 - \frac{1}{2} \left(\frac{|\mathbf{x}_1 \cap \mathbf{x}_2|}{|\mathbf{x}_1|} + \frac{|\mathbf{x}_1 \cap \mathbf{x}_2|}{|\mathbf{x}_2|} \right).$$

Classical MDS (Cox and Cox 2001) was used to map the species onto 4-dimensional Euclidean space, and the resulting points were clustered by fitting a Gaussian mixture model with a uniform noise component using the R-package “mclust” (Fraley and Raftery 1998, 2002; Fraley et al. 2012). The reason for this was that experience with such datasets suggested that biotic elements may differ quite a bit regarding within-cluster variation, which could be captured better with the potentially different covariance matrices of Gaussian distributions than with standard distance-based clustering methods, see Hennig and Hausdorf (2004) for details.

A standard way to estimate the number of clusters (and also potential constraints of the covariance matrices) is the BIC, which here is defined as $2l_n(k) - r(k) \log(n)$ so that large values of the BIC are good, where $l_n(k)$ is the log maximized likelihood for k Gaussian components and $r(k)$ is the corresponding number of free parameters. The BIC is the default method in “mclust”. It can be computed for $k = 1$ and it can therefore also deliver a decision about whether $k = 1$ (interpreted as homogeneity of the distribution or absence of biotic elements) or not. The BIC is here used as cluster validation index, although this is not how it would normally be interpreted, because it is defined based on probability theory within a certain model. In Sect. 5.7 a modification is used, which compares the BIC value for $k > 1$ with the one for $k = 1$.

5.3 Results with plain Gaussian null model

Given that the BIC is proved to be consistent under certain (somewhat restrictive) conditions (Keribin 2000), and that it makes a decision involving a homogeneous “null model”, one may wonder whether parametric bootstrap adds anything to using the Gaussian mixture model combined with the BIC. Figure 5 was produced by parametric bootstrap with bootstrap data generated from a simple Gaussian null model, parameters of which were estimated in the standard way from the MDS output ($m = 200$). These data were clustered by “mclust” in the same way as the original snails data, and the BIC was computed for $k = 1, \dots, 10$. The BIC points to 8 clusters. Given that the expected BIC seems to go down linearly under the null model (it should indeed yield a maximum for $k = 1$, which is true under the null model), the clustering indeed seems to be highly significant, and the maximum of the BIC at $k = 8$ seems to be confirmed as optimal number of clusters. The only new thing that the parametric bootstrap shows is that the null model produces occasional outliers (caused by spurious clusters that generate a high log-likelihood) for the larger values of k ; thus, the

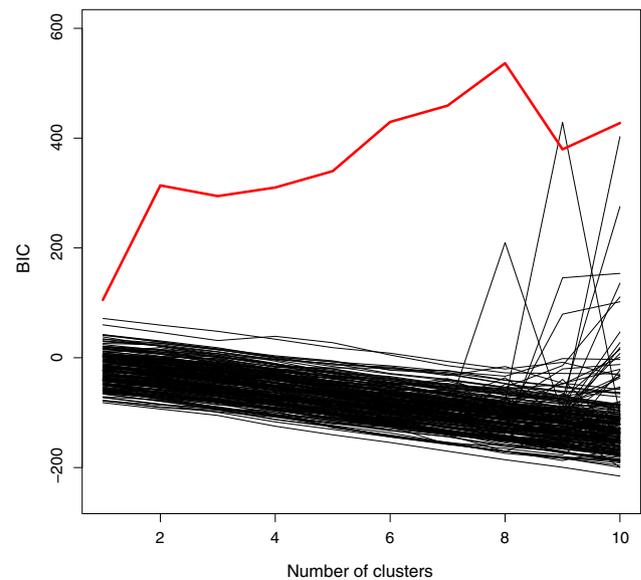


Fig. 5 BIC values for Aegean islands snails data (red) and 500 bootstrap samples (black) drawn from plain Gaussian null model for $k = 1, \dots, 10$. (Color figure online)

parametric bootstrap sheds some light on how the asymptotics of the BIC become unreliable if the number of parameters is too large. This, however, would probably not distract the researcher from declaring $k = 8$ to be the optimal number of clusters here.

5.4 Non-clustering structure

The reasoning in the previous subsection does not take into account the way the snails data were pre-processed, though. Instead, it takes the MDS output at face value. But there is some structure in the original data, which could explain an apparent clustering in the MDS output. The presences of the species are spatially autocorrelated. If it is known that a species is present on a certain island I , it is more likely to find the species also on neighbouring islands than if the species is known not to be present on I . Furthermore, some islands host more species than others (“island attractivity”) because of factors such as their size and vegetation, and there is a certain distribution of species sizes $|\mathbf{x}|$. We interpret all this structure as not indicating clustering; a clustering should be interpreted as groups of species being attracted significantly to certain specific islands, and other groups of species being attracted to other islands, which contradicts the idea that the species can be modelled as i.i.d., taking autocorrelation and island attractivity into account in the same way for all species. This involves some judgement; it is conceivable, for example, to interpret variations in island attractiveness as a consequence of real clustering rather than as a non-clustering feature of

the data, which, however, does not agree with the “biotic element”-concept in Hausdorf and Hennig (2003).

5.5 Null model for spatial autocorrelation

The null model used here has already been used in Hennig and Hausdorf (2004), although in a different way, not connected with the number of clusters. According to \mathcal{P}_0 , the species \mathbf{x}_i , $i = 1, \dots, n$ are modelled as i.i.d., having arisen from the following process. The parameters are an autocorrelation (“disjunction”) parameter p_d , the distribution of species sizes P_S and the island attractivity distribution P_I . Assume that there is a neighbourhood list indicating for each pair of island whether they are neighbours or not.

1. Draw $n_i = |\mathbf{x}_i|$ from P_S .
2. Draw an initial island $I_1 \in \mathbf{x}_i$ from P_I .
3. If $n_i > 1$, for $j \in 2, \dots, n_i$:
 - (a) Let N_0 be the set of non-neighbours of all of I_1, \dots, I_{j-1} (not including I_1, \dots, I_{j-1}). Let N_1 be the set of neighbours of all of I_1, \dots, I_{j-1} (not including I_1, \dots, I_{j-1}). If neither N_0 nor N_1 is empty:
 - (b) With probability p_d , draw an island I_j from the set of non-neighbours N_0 of all of I_1, \dots, I_{j-1} (not including I_1, \dots, I_{j-1}) according to P_I conditionally on N_0 .
 - (c) Otherwise, draw I_j from the set of neighbours N_1 of any of I_1, \dots, I_{j-1} (not including I_1, \dots, I_{j-1}) according to P_I conditionally on N_1 .
 - (d) If either $N_0 = \emptyset$ or $N_1 = \emptyset$, draw I_j from the remaining non-empty set N^* according to P_I conditionally on N^* .
 - (e) Put $I_j \in \mathbf{x}_i$.

5.6 Null model parameter estimation

P_S can be estimated by the empirical distribution of species sizes. P_I has a straightforward empirical counterpart as well, namely choosing the probability for each island proportional to the number of species on that island. The estimation of p_d is a bit more subtle. A naive estimator for p_d would be $q_d = \frac{\sum_{i=1}^n (a_i - 1)}{\sum_{i=1}^n (n_i - 1)}$, where a_i is the number of connectivity components of the species distribution range \mathbf{x}_i . However, q_d may not work very well because of situations with $N_0 = \emptyset$ or $N_1 = \emptyset$ in Sect. 5.5, and because initially separated connectivity components may grow together in the process of generating a species. We use the recommendation of Hennig and Hausdorf (2004) to simulate the observable q_d as a function of p_d by sampling from the null model with various values of p_d , to fit a linear regression explaining q_d from p_d , and then by estimating the real p_d by plugging the real

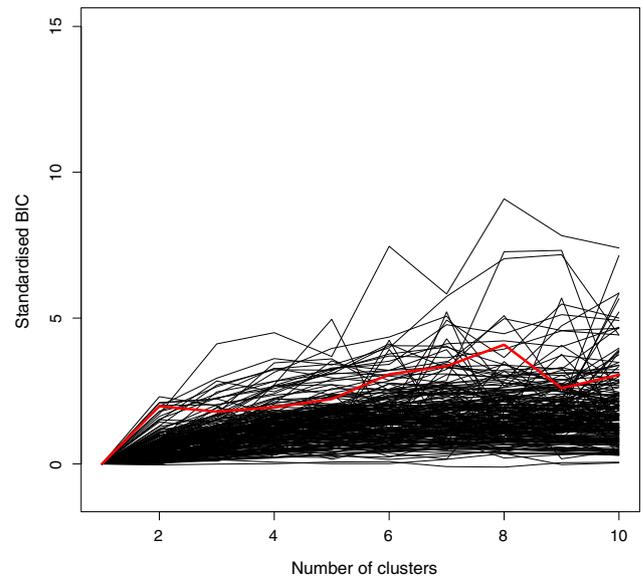


Fig. 6 ($k = 1$)-adjusted BIC values for Aegean islands snails data (red) and 500 bootstrap samples (black) drawn from null model with spatial autocorrelation for $k = 1, \dots, 10$. (Color figure online)

observed q_d into the regression equation, as implemented in the “prabclus”-package of R.

5.7 Parametric bootstrap

Repeat m times:

1. Estimate the null model parameters according to Sect. 5.6.
2. Generate n i.i.d. species in the bootstrap sample \mathbf{X}^* according to the algorithmic null model in Sect. 5.5.
3. Compute the Kulczynski dissimilarity matrix between the species.
4. Map them onto \mathbb{R}^4 by classical MDS.
5. Cluster them fitting a Gaussian mixture model with a uniform noise component for $k \in K$.
6. Let $V_{qk} = V(\mathbf{X}^*, C(\mathbf{X}^*, k))$, $q = 1, \dots, m$, be the resulting value of $\frac{\text{BIC}_q(k) - \text{BIC}_q(1)}{\text{BIC}_q(1)}$, where $\text{BIC}_q(k)$ is the BIC value for k mixture components.

The last step contains a subtle adjustment. Instead of the plain BIC, V_{qk} adjusts the BIC by the BIC value for $k = 1$. The reason for this is that the parametric bootstrap test compares a clustering alternative, i.e. $k > 1$, with a null model for a homogeneous population, i.e. $k = 1$. A dataset with large BIC(1) can be expected to have smaller variation (as expressed by covariance matrix eigenvalues) than a dataset with smaller BIC(1), and therefore also the former dataset’s BIC values for larger k will likely be larger. The adjustment corrects for this.

5.8 Results

The results are shown in Fig. 6. The real dataset does not really stick out, but its ($k = 1$)-adjusted BIC values (V) are certainly among the higher ones generated by the null model. Equation (2) yields $\hat{p} = 0.0697$, so that the BIC values of the snails data fail rather tightly to be significant evidence for real clustering at the 5%-level. Taking into account, the ad hoc nature of the model and the estimation, and the resulting anti-conservativity of the p values (see Sect. 2), the stronger statement that there is no evidence for clustering seems justified (by the way, using the raw BIC as V produces an even higher \hat{p}). Furthermore, different from what the raw BIC values indicated, Eq. (3) indicates $k = 2$ as clearly better, with calibrated V of 3.149, than $k = 8$ (second best with calibrated V of 1.971). Overall, taking into account the spatial autocorrelation of the unprocessed presence–absence data in this way changes the results quite a bit, compared to Sect. 5.3, and can particularly explain the extent to which clustering is observed through the BIC.

6 Concluding discussion

The present paper provides a general scheme by which cluster validation indexes can be used for testing homogeneity against a clustering alternative, and for calibrating the indexes so that their expected distribution can be taken into account for estimating the number of clusters. In the examples, it was shown that this approach can expose weaknesses of standard recommendations for how to use the validation indexes to estimate the number of clusters, and that it can detect that in some situations, in which researchers obtain a clustering, there is no evidence for any real clustering at all.

The scheme cannot be applied “straight out of the box”, but requires the researcher to make judgements on which structural features of the data cannot be taken as indicating a “real clustering”, and to use these for defining an appropriate null model. The examples shown here should illustrate how to go on about this task in a real situation.

The null models and parameter estimators used here have been set up in some kind of ad hoc-fashion, partly taking into account information from looking at the data. Also, parametric bootstrap will generally result in anti-conservative p values, because the p values are simulated from specific parameter values, and it cannot be ruled out that other parameter values can be found that are compatible with the data and yield still better values of the cluster validation index.

Both of these issues do not affect the interpretation of non-significant outcomes, because in any case a model with certain parameter values has been found that can explain the observed clustering. This means that despite the shortcom-

ings the interpretation is valid that there is no evidence for real clustering.

On the other hand, significant outcomes have to be interpreted with more care. They are certainly more convincing if the V -values for the real dataset look clearly different from those from all datasets generated by the null model in the bootstrap validity plot, rather than only just achieving $\hat{p} < 0.05$ or 0.01. Some sensitivity analysis, i.e. running the parametric bootstrap with slightly different parameters for the null model, may give the researcher a clearer impression of how stable the significance is.

The ad hoc-character of the null models and estimators presented here may not satisfy the theoretically oriented statistician, but it is meant to encourage the data analyst to set up such models where they could be helpful even if no worked out theory is available. Apart from the direct benefit of having a test of homogeneity and a calibration of indexes for estimating the number of clusters, it is also potentially instructive to think about clustering and non-clustering structural features of the data having the task of setting up such a null model in view. Carrying out this task may give researchers a clearer idea of what kind of clusters they are looking for, and what it means, in their field, to distinguish “clustered” from “homogeneous” data.

Further research is required regarding comparing different schemes for computing p values and estimating k as mentioned in Sect. 2. Another interesting issue is whether for estimating k the index values for k clusters should rather be compared with what is expected if there are $k - 1$ true clusters than with a homogeneous model with only one cluster. This, however, would require the construction of probability models for all numbers of clusters. In principle, one could think of fitting a mixture of null models with different parameters per cluster. In many cases, this will require considerable effort, and it can also not necessarily be taken for granted that a homogeneous null model for fitting the whole dataset is also suitable for fitting (and implicitly defining) a cluster. For example, one may want low within-cluster distances, but the null models used here are not constructed with having such an objective in mind.

Acknowledgments The research was supported by EPSRC Grant EP/K033972/1.

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

References

- Arbelaitz, O., Gurrutxaga, I., Muguerza, J., Perez, J.M., Perona, I.: An extensive comparative study of cluster validity indices. *Pattern Recognit.* **46**, 243–256 (2012)

- Bock, H.H.: Probabilistic models in cluster analysis. *Comput. Stat. Data Anal.* **23**, 5–28 (1996)
- Calinski, T., Harabasz, J.: A dendrite method for cluster analysis. *Commun. Stat. Theory Methods* **3**, 1–27 (1974)
- Cox, T.F., Cox, M.A.A.: *Multidimensional Scaling*, 2nd edn. Chapman and Hall/CRC, Boca Raton (2001)
- Drasgow, F.: Polychoric and polyserial correlations. In: Kotz, S., Johnson, N. (eds.) *The Encyclopedia of Statistics*, pp. 68–74. Wiley, New York (1986)
- Efron, B., Tibshirani, R.: *An Introduction to the Bootstrap*. CRC Press, Boca Raton (1994)
- Fraley, C., Raftery, A. E., Murphy, T. B., Scrucca L.: Mclust version 4 for R: Normal mixture modeling for model-based clustering, classification, and density estimation. Technical Report 597, Department of Statistics, University of Washington (2012)
- Fraley, C., Raftery, A.E.: How many clusters? Which clustering methods? Answers via model-based cluster analysis. *Comput. J.* **41**, 578–588 (1998)
- Fraley, C., Raftery, A.E.: Model-based clustering, discriminant analysis and density estimation. *J. Am. Stat. Assoc.* **97**, 611–631 (2002)
- Hausdorf, B., Hennig, C.: Biotic element analysis in biogeography. *Syst. Biol.* **52**, 712–723 (2003)
- Hausdorf, B., Hennig, C.: The influence of recent geography, palaeography and climate on the composition of the faune of the central aegean islands. *Biol. J. Linn. Soc.* **84**, 785–795 (2005)
- Hennig, C., Liao, T.F.: Comparing latent class and dissimilarity based clustering for mixed type variables with application to social stratification. *J. R. Stat. Soc. Ser. C* **62**, 309–369(2013)
- Hennig, Christian, Hausdorf, Bernhard: Distance-based parametric bootstrap tests for clustering of species ranges. *Comput. Stat. Data Anal.* **45**, 875–895 (2004)
- Jain, A.K., Dubes, R.C.: *Algorithms for Clustering Data*. Prentice Hall, Englewood Cliffs (1988)
- Kaufman, L., Rousseeuw, P.: *Finding Groups in Data*. Wiley, New York (1990)
- Keribin, C.: Consistent estimation of the order of mixture models. *Sankhya Ser. A* **62**(1), 49–66 (2000)
- Lin, Chien-Ju: A pattern-clustering method for longitudinal data—heroin users receiving methadone. PhD thesis, Department of Statistical Science, University College London, London (2014)
- Lin, C.-J., Hennig, C., Huang, C. L.: Clustering and a dissimilarity measure for methadone dosage time series. In: *Proceedings of ECDA-2014*, Bremen, Germany, page to appear. Springer, Berlin (2015)
- Milligan, G., Cooper, M.: An examination of procedures for determining the number of clusters in a data set. *Psychometrika* **50**(3), 159–179 (1985)
- Sugar, Catherine, James, Gareth: Finding the number of clusters in a dataset. *J. Am. Stat. Assoc.* **98**(463), 750–763 (2003)
- Tibshirani, R., Walther, G.: Cluster validation by prediction strength. *J. Comput. Graph. Stat.* **14**, 511–528 (2005)
- Xiong, H., Li, Z.: Clustering validation measures. In: Aggarwal, C.C., Reddy, C.K. (eds.) *Data Clustering: Algorithms and Applications*, pp. 571–606. CRC Press, Boca Raton (FL) (2014)