



Mixture polarization in inter-rater agreement analysis: a Bayesian nonparametric index

Giuseppe Mignemi¹ · Antonio Calcagni^{3,4} · Andrea Spoto¹ · Ioanna Manolopoulou²

Accepted: 4 December 2023
© The Author(s) 2024

Abstract

In several observational contexts where different raters evaluate a set of items, it is common to assume that all raters draw their scores from the same underlying distribution. However, a plenty of scientific works have evidenced the relevance of individual variability in different type of rating tasks. To address this issue the intra-class correlation coefficient (ICC) has been used as a measure of variability among raters within the Hierarchical Linear Models approach. A common distributional assumption in this setting is to specify hierarchical effects as independent and identically distributed from a normal with the mean parameter fixed to zero and unknown variance. The present work aims to overcome this strong assumption in the inter-rater agreement estimation by placing a Dirichlet Process Mixture over the hierarchical effects' prior distribution. A new nonparametric index λ is proposed to quantify raters polarization in presence of group heterogeneity. The model is applied on a set of simulated experiments and real world data. Possible future directions are discussed.

Keywords Bayesian nonparametrics · Inter-rater agreement · Dirichlet process mixture · Hierarchical Bayesian models

1 Introduction

In several contexts, decision-making relies heavily (or exclusively) on expert ratings, especially in situations where a direct quantification of quality of an object or a subject is either impossible or unavailable. Examples include applicant selection procedures, grading of student assignments in education, or risk evaluation in emergencies, all of which rely on observational ratings made by experts. For ease of exposition, throughout this paper we will refer to evaluation of students' work in an educational context as the primary example. To ensure consistency across different teachers, harmonization of marking criteria is often used to improve inter-rater

Extended author information available on the last page of the article

agreement and homogeneity (Gisev et al. 2013; Gwet 2008); however, discrepancies between grades assigned by different teachers may still persist (Bygren 2020; Barneron et al. 2019; Makransky et al. 2019; Zupanc and Štrumbelj 2018), reflecting each teacher's approach to evaluation. Therefore, statistical models that can capture inter-rater agreement (or disagreement) can shed light on heterogeneity between teachers and aid the mark moderation process (Bygren 2020; Barneron et al. 2019; Crimmins et al. 2016).

The specific context that we are considering in this work is the observational setting where a set of raters are evaluating different sets of items¹ out of a total population of items; these sets may be completely disjoint (i.e., each item is evaluated by exactly one rater). Each item is represented by a set of covariates, assumed to follow some distribution. Within a hierarchical statistical model, a common assumption is that raters (who may or may not include covariates) may each be characterized through a latent variable capturing e.g. whether an evaluator is generous or how they assess different aspects of the work. In the simplest setting, in an evaluation context where there is no space for subjectivity, these latent variables will be identical for all raters, in the sense that their view of the item is identical and as a result their evaluation style is assumed to be the same. However, it is well-known in many scientific fields, e.g., cognitive neuroscience (Barneron et al. 2019; Makransky et al. 2019; Briesch et al. 2014), statistics (Agresti 2015; Gelman et al. 2013) and psychometrics (Bartoš et al. 2020; Nelson and Edwards 2015; Hsiao et al. 2011), that individual variability in rating tasks (Wirtz 2020) needs to be accounted for when aggregating or interpreting individual raters' recommendations.

Existing works account for heterogeneity between raters through a latent variable within a mixed-effects model (Martinková et al. 2023; Bartoš et al. 2020; Nelson and Edwards 2015, 2008). In other words, a regression model is used where the rating is modelled conditionally on covariates with a random effect that varies across raters. However, the distribution of the latent variable is typically assumed to be unimodal, and cannot capture e.g. polarisation or clustering of rater types. The present work aims to extend these models to account for clustered variability between raters. Through a Bayesian approach, a Dirichlet process mixture prior is placed over the hierarchical rater effects in a linear model. This flexible prior naturally accommodates different clusters among raters (i.e., different distributions for the rater effects). A multiple-level model is specified in which observations (i.e., ratings) are nested within raters, and in turn these are nested within clusters. These clusters reflect distinct groups of raters in terms of their decision-making, and can be used to characterise the level of (dis)agreement. The level of multimodality (i.e., how separated the latent group densities are) quantifies the polarization of the latent groups. For instance, a large variance between teacher scores might be due to both the presence of two main divergent latent trends among them or to a high level of noise in their assessing (Koudenburg and Kashima 2022). It is important to differentiate the two cases and quantify the group polarization both for theoretical and practical purposes. Differentiate systematic differences of opinion against high level of noise might be needed (Koudenburg et al. 2021). They are two very different cases and much

¹ Commonly referred to as *subjects* in the rating context.

attention must be paid in distinguishing one another. The former is a case of high group polarization (Esteban and Ray 1994): two different teachers clusters emerge with a small within-cluster variance and a large variance between different clusters. In the second case only one cluster emerge with a large variance. It might be argue that in the first case, even that the overall agreement might be quite low since there are two main different trends among raters, there might be a high agreement within the same trend (Tang et al. 2022). Assuming the latent agreement among raters as the degree of latent similarity in rating, an index regarding the polarization of the different possible groups of raters might be informative (Koudenburg and Kashima 2022; Tang et al. 2022; Koudenburg et al. 2021). In this work we introduce a novel index to quantify the latent polarization among raters through the posterior distribution of the hierarchical effects (DiMaggio et al. 1996). It naturally derives from the nonparametric model and overcomes some strong assumptions (e.g., the number of latent groups, the ratings distribution) of the previous indices (Koudenburg et al. 2021; Esteban and Ray 1994). This nonparametric index, referred to as λ index, is based on the shape of the posterior distribution of the hierarchical effects. It connects two different research lines: it relates the works on distribution polarization of opinions (Koudenburg and Kashima 2022; Tang et al. 2022; Koudenburg et al. 2021) with those about the inter-rater agreement analysis (Martinková et al. 2023; Bartoš et al. 2020; Nelson and Edwards 2015; Gisev et al. 2013; Gwet 2008; Nelson and Edwards 2008).

The paper proceeds as follows: Sect. 2 is devoted to the general psychometric framework, the key concepts of inter-rater agreement, inter-rater reliability are introduced; the statistical model is specified in Sect. 3 and the adopted Gibbs sampler in Sect. 4; the novel rater similarity index is described in Sect. 5; simulation studies are reported in Sect. 6, as an illustrative example, a real data analysis is described in Sect. 8; it is followed by conclusion and future directions in Sect. 1.

2 Existing work in inter-rater agreement and hierarchical effects models

Several methods and statistical models that aim to account for inter-rater variability have appeared in the literature (Nelson and Edwards 2015; Gwet 2008; Cicchetti 1976). Models such as the Cultural Consensus Theory (Oravecz et al. 2014), which explores individuals' shared cultural knowledge, have been proposed to capture unobserved agreement and similar trends in groups of raters (Dressler et al. 2015). Two related but different concepts have been introduced: *inter-rater agreement* and *inter-rater reliability*. The former refers to the extent to which different raters' evaluations are concordant (i.e., they assign the same value to the same item), whereas the latter refers to the extent to which their evaluations consistently distinguish different items (Gisev et al. 2013). In other words, while the inter-rater agreement indices quantify the *observed* concordance, the inter-rater reliability indices aim to quantify the *consistency* of their evaluations (e.g., despite assigning different values, the distinction among the items is the same). The present work focuses on latent agreement

intended as homogeneity in the evaluators' point of view (Tang et al. 2022; Esteban and Ray 1994).

A number of methods are available to quantify both inter-rater agreement and inter-rater reliability. Indices for pairs (Nelson and Edwards 2008; McHugh 2012) or multiple raters (Jang et al. 2018), for binary (Gwet 2008), polytomous (Nelson and Edwards 2015) or continuous (Liljequist et al. 2019) ratings are commonly used in different contexts. Recent developments using the framework of Hierarchical Linear Models (i.e., HLMs) provide a more accurate estimation of inter-rater reliability accounting for different sources of variability (Martinková et al. 2023).

Despite the popularity of work on this issue, less attention has been paid to possible latent similarities of the raters (Wirtz 2020). From a psychometric point of view, it can be appealing to assess the extent to which different raters might be heterogeneous in their ratings (Martinková et al. 2023; Bartoš et al. 2020; Koudenburg et al. 2021; Casabianca et al. 2015; Nelson and Edwards 2015; Gisev et al. 2013; DeCarlo 2008; Gwet 2008; Nelson and Edwards 2008).

There are certain situations in which the subjective opinion of the raters is very informative; as a simple example, the type of teachers' training or experience can be thought of as latent states which affect a range of evaluations differently (Childs and Wooten 2023; Barneron et al. 2019; Bonefeld and Dickhäuser 2018; Dee 2005). Sometimes the major interest is not on the mere consistency between raters, but on their actual evaluation. For instance, in a selection process the actual students' scores are very relevant for their admission (Zupanc and Štrumbelj 2018). Even if a strict standardization of teachers evaluation is not feasible, some statistical methods can tackle these issues. In all these contexts the assessment of uniformity among raters could be useful and would provide further information about the rating process.

To this aim, existing work, e.g. Martinková et al. (2023); Nelson and Edwards (2015); Casabianca et al. (2015); Hsiao et al. (2011); Cao et al. (2010); DeCarlo (2008), adopts an hierarchical approach where correlations between ratings are naturally captured through an hierarchical Bayesian model. Each rater $i = 1, \dots, I$ is assumed to be rating a different set of items $\mathcal{J}_i \in \mathcal{J}$, $\mathcal{J}_i \cap \mathcal{J}_{i+1} = \emptyset$.² The rating y_{ij} of the item $j \in \mathcal{J}_i$ carried out by rater $i = 1, \dots, I$, is modelled as follows:

$$y_{ij} = \mathbf{x}'_{ij}\boldsymbol{\beta} + \mathbf{z}'_{ij}\mathbf{u}_i + \varepsilon_{ij}, \quad i = 1, \dots, I, j \in \mathcal{J}_i. \quad (1)$$

Here \mathbf{x}_{ij} and \mathbf{z}_{ij} are, respectively, $1 \times p$ and $1 \times q$ vectors of distinct explanatory variables of rating y_{ij} ; $\boldsymbol{\beta}$ is a $p \times 1$ vector of non varying effects and \mathbf{u}_i is a $q \times 1$ vector the hierarchical effects of rater i .

In the standard HLM formulation, the following distribution is specified for the rater effects:

$$\mathbf{u}_i \sim N_q(\mathbf{0}, \boldsymbol{\Sigma}), \quad i = 1, \dots, I.$$

Where $N_q(\cdot)$ stands for a q -variate normal distribution; Here $\mathbf{0}$ is a $q \times 1$ zero vector and $\boldsymbol{\Sigma}$ is a $q \times q$ positive semi-definite covariance matrix. For the hierarchical

² The multiple rating case (i.e., raters rate the same set of items, $\mathcal{J}_i = \mathcal{J}$, $i = 1, \dots, I$) is addressed in Appendix.

normal linear model $\varepsilon_{ij} \sim N(0, \sigma_\varepsilon^2)$, with \mathbf{u}_i and ε_{ij} typically assumed independent. The distribution of each vector-valued hierarchical effects \mathbf{u}_i is then assumed to follow some distribution and captures variability across different raters.

In the above mentioned example, y_{ij} is the score given to student $j \in \mathcal{J}_i$'s essay by teacher i . Since an observational approach is adopted (i.e., each raters rates a different set of items), the effect of the student is not identifiable (each student is rated only by one rater). Assuming that students effects are i.i.d., their variance is added to that of the residuals. In the univariate case (i.e., when $z_i = 1$, varying intercept model) the relevance of the raters effect $u_i \sim N(0, \sigma_u^2)$, where $\sigma_u^2 > 0$ is the variance, might be quantified through the *intraclass correlation coefficient* (i.e., ICC):

$$ICC = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_\varepsilon^2}.$$

It is the ratio between the variance of the raters effect and the total variability of the model, i.e., the proportion of variance of the score due to the teacher, which reflects the correlation of two ratings given by the same rater. Smaller values of ICC indicate a small effect of the rater on the student's score.

3 Dirichlet process mixture and hierarchical effects

The HLM assumption regarding the distribution of the hierarchical effects is crucial in characterising different possible clusters or latent patterns of heterogeneity among raters (Dorazio 2009). The common Gaussian assumption for the distribution of these effects may obscure skewness and multimodality present in the data. A more flexible specification of the hierarchical effects distribution can help capture more complex patterns of variability. Models that account for skew-normal (Lin and Lee 2008), skew-normal-cauchy (Kahrari et al. 2019), multivariate t (Wang and Lin 2014), extreme values (McCulloch and Neuhaus 2021) effects distributions have been proposed (Schielzeth et al. 2020). Nevertheless, they poorly account for the possible presence of multimodality in those distributions. In this regard, a mixture distribution has been proposed as a potential solution (Heinzl and Tutz 2013; Kyung et al. 2011; Kim et al. 2006). Each mode can then correspond to a cluster with a similar pattern (e.g., the same deviation from the population mean). Several works have explored this issue in the past two decades (Villarroel et al. 2009; Tutz and Oelker 2017). For instance, Verbeke and Lesaffre (Verbeke and Lesaffre 1996) proposed a standard normal mixture distributions for the hierarchical effects. James and Sugar (James and Sugar 2003) explored this approach in the context of functional data. De la Cruz-Mesía (De la Cruz-Mesía and Marshall 2006) proposed a mixture distribution for non-linear hierarchical effects in modelling continuous time autoregressive errors. A heteroscedastic normal mixture model in the hierarchical effects distribution was considered in linear (Komárek et al. 2010) and generalized hierarchical linear (Komárek and Komárková 2013) models. Despite the breadth of specifications for the mixture model, in all the aforementioned models, the number of mixture

components needs to be specified. Although this may not be a critical assumption in certain contexts, it may be questionable or detrimental in settings with a lack of a priori information on the level of multimodality, especially in cases where the characterisation of the multimodality is of direct interest.

When the number of components of the mixture is unknown, a Dirichlet Process Mixture (hereafter DPM) for the hierarchical effects is a natural extension (Gill and Casella 2009; Navarro et al. 2006; Verbeke and Lesaffre 1996). This nonparametric extension allows the model to capture an unknown marginal distribution of the hierarchical effects through the Dirichlet Process (Antoniak 1974; Ferguson 1973). Modeling the hierarchical effect \mathbf{u}_i as an infinite mixture of some distribution family (e.g., Normal) enables the model to account for possible multimodality without specifying the number of mixture components. Some existing works adopted this nonparametric approach and pose a DPM prior over the hierarchical effects (e.g., Heinzl and Tutz (2013); Heinzl et al. (2012); Kyung et al. (2011)).

The HLM of Eq. (1) is then specified in the same way as before through:

$$y_{ij} = \mathbf{x}'_{ij}\beta + \mathbf{z}'_i\mathbf{u}_i + \varepsilon_{ij}, \quad i = 1, \dots, I, \quad j \in \mathcal{J}_i.$$

The following hierarchical prior distribution is placed over the raters effects:

$$\begin{aligned} \mathbf{u}_i | \boldsymbol{\mu}_i, \mathbf{Q}_i &\sim N_q(\boldsymbol{\mu}_i, \mathbf{Q}_i) \\ \boldsymbol{\mu}_i, \mathbf{Q}_i | G &\stackrel{iid}{\sim} G \\ G &\sim DP(\alpha, G_0) \end{aligned}$$

where $\boldsymbol{\mu}_i$ and \mathbf{Q}_i are, respectively, the $q \times 1$ a location parameter vector and the $q \times q$ positive semi-definite covariance matrix for the hierarchical effects \mathbf{u}_i of rater $i = 1, \dots, I$. Here $\varepsilon_{ij} \sim N(0, \sigma_\varepsilon^2)$, $i = 1, \dots, I$, $j \in \mathcal{J}_i$; \mathbf{u}_i and ε_{ij} are assumed independent as before.

3.1 DPM as a generative process for the hierarchical effects

Here, $DP(\alpha, G_0)$ is a DPM with $\alpha > 0$ precision parameter and base measure G_0 . These specify the mixing distribution G (Heinzl and Tutz 2013), so that each realization of G is almost surely a discrete probability measure on the space (Ω, \mathcal{F}) (Blackwell 1973). Thus, since the DPM is a discrete generative process with non-zero probability of ties, some of the realizations might be identical to each other with probability determined by the precision parameter α . Therefore, specifying this hierarchical model on the components location parameters $\boldsymbol{\mu}$ induces a clustering in the hierarchical effects (i.e., the raters) (Kyung et al. 2011); hierarchical effects belonging to the same c -th cluster with location parameter $\boldsymbol{\mu}_c$ are then independent and identically distributed. In other words, in the context of the HLM, the DPM specifies the component-specific location parameter $\boldsymbol{\mu}_c$, so that each rater has each their own unique hierarchical effects value \mathbf{u}_i (Heinzl and Tutz 2013).

The DPM is a generative process commonly used in conjunction with a parametric family of distributions (e.g., Normal, Poisson), and the base measure parameter

G_0 denotes this specified distribution. Thus, for any element $A_n, n = 1, \dots, N$, of \mathcal{A} , a finite measurable partition of Ω ,

$$(G(A_1), G(A_2), \dots, G(A_N)) \sim Dir(\alpha G_0(A_1), \alpha G_0(A_2), \dots, \alpha G_0(A_N))$$

where $Dir(\cdot)$ stands for the Dirichlet distribution, and G_0 defines the expectation of G^3 , therefore they have the same support. The parameter α , a multiplicative constant of the vector-valued Dirichlet parameter, determines the probability of a new realization of the process to be different of the previous ones (Blackwell and MacQueen 1973). In other words, it governs the probability that the DPM generates a new cluster. Formally, the generative property of the DPM is that, for $i = 1, \dots, I$, with I being for instance the total number of raters:

$$G \sim DP(\alpha, G_0), \quad \mu_i | G \sim G$$

the probability that the new I -th realization μ_I of G assumes a different values than the previous ones is described by the well known *Pólya Urn Model*:

$$\mu_I | \mu_1, \mu_2, \dots, \mu_{I-1}, \alpha \sim \frac{\alpha}{\alpha + I - 1} G_0 + \frac{1}{\alpha + I - 1} \sum_{c=1}^C r_c$$

with $C \in \mathbb{N}$ being the number of already observed distinct clusters among the realizations of G (i.e., the number of the different values of μ already observed, in other words the number of clusters) and r_c counts the elements in the c -th cluster. Basically, since G is a discrete probability measure, the C clusters represent different point masses (or different sets of point masses in the multivariate case) and r_c is the frequency of each of them. Considering the conditional distribution of μ_I as a mixture distribution, the probability that μ_I is a new point mass sampled from G_0 is proportional to α , the probability that it is equal to the already observed c -th point mass is proportional to r_c . In this notation, the role of α in sampling a new (not already observed) value of μ_I (i.e., a new point mass, a new cluster) is interpretable.

To this regard, Sethuraman (1994) described a *stick-breaking construction* of the DP.⁴ In this formulation G is equivalent to:

$$G = \sum_{c=1}^{\infty} \pi_c \delta_{\mu_c}$$

where δ is the Dirac measure on μ_c and $\mu_c \stackrel{iid}{\sim} G_0$ is assumed. The weights $\{\pi_c\}_{c=1}^{\infty}$ of the infinite mixture result from the stick-breaking procedure as follows:

³ Considering the partition (A, A^c) of Ω and thus that $G(A) \sim Be(\alpha G_0(A), \alpha G_0(A^c))$ the expectation of $G(A)$ is defined as: $\mathbb{E}[G(A)] = \frac{\alpha G_0(A)}{\alpha G_0(A) + \alpha G_0(A^c)} = \frac{\alpha G_0(A)}{\alpha(G_0(A) + G_0(A^c))} = G_0(A)$.

⁴ Other stick-breaking representations might be used, e.g., Rigon and Durante (2021); Stefanucci and Canale (2021); Rodriguez and Dunson (2011).

$$\pi_c = v_c \prod_{l < c} (1 - v_l)$$

$$v_c \stackrel{iid}{\sim} Be(1, \alpha)$$

with $Be(\cdot)$ indicating the Beta distribution and $\{v_c\}_{c=1}^{\infty}$ being reparameterized weights. It is even more explicit in this construction that the random measure G is a mixture of point masses. The distribution of the random weights $\boldsymbol{\pi}$ (i.e., the probability of different allocation to the clusters) is governed through the stick-breaking process by the precision parameter α . Further details are given in the Appendix.

In practice, one of the established approximations to the stick-breaking process is to truncate the infinite number of components to a large, finite value:

$$G = \sum_{c=1}^R \pi_c \delta_{\boldsymbol{\mu}_c}$$

for large enough value of R (Tutz and Oelker 2017; Gelman et al. 2013).

In summary, the hierarchical effects distribution considering a stick breaking construction of the DPM might be then specified as follow:

$$\mathbf{u}_i | \boldsymbol{\mu}, \mathbf{Q}, \stackrel{iid}{\sim} \sum_{c=1}^R \pi_c N_q(\boldsymbol{\mu}_c, \mathbf{Q}_c), \quad i = 1, \dots, I$$

$$\boldsymbol{\mu}_c, \mathbf{Q}_c \stackrel{iid}{\sim} G_0$$

$$\pi_c = v_c \prod_{l < c} (1 - v_l), \quad \text{where}$$

$$v_c \stackrel{iid}{\sim} Be(1, \alpha), \quad c = 1 \dots, R$$

With this nonparametric model specification, latent common tendencies among raters might emerge through the components of the model (Heinzl and Tutz 2013; Heinzl et al. 2012; Kyung et al. 2011). The Bayesian approach allows us to characterize the shape of the distribution of the rater effects, as well as explore the effect of uncertainty on these (Gelman et al. 2013). For example, in an applied context, strict vs. accommodating are very common latent states that drive students' essays grading process (Zupanc and Štrumbelj 2018; Briesch et al. 2014; Dee 2005).

4 Prior distributions and estimation procedure

The DPM mixture model has been well studied in the literature in a variety of different settings, especially within Bayesian inference (Canale and Prünster 2017; Müller et al. 2015). Several sampling schemes have been proposed both in the Bayesian context (e.g., Canale and Dunson (2011); Dahlin et al. (2016); Kyung et al. (2011)) and in the frequentist one (e.g., Tutz and Oelker (2017)). Within the Bayesian framework, Gibbs sampling (Dahlin et al. 2016), slice sampler (Kyung et al. 2011; Walker 2007), Sequential Monte Carlo algorithms (Ulker et al. 2010), split-merge algorithms (Bouchard-Côté et al. 2017), have been proposed among others.

In this work, the model specification permits the use of conjugate priors, so that a blocked Gibbs sampling can be used (Heinzl and Tutz 2013; Heinzl et al. 2012; Kyung et al. 2011), with details shown below.

4.1 Prior specification

Several of the parameters in the model have conjugate prior distributions which allow easier computation.

- For the effects β the following hierarchical prior is assigned:

$$\begin{aligned} \beta | \mathbf{b}_\beta, \mathbf{B}_\beta &\sim N_p(\mathbf{b}_\beta, \mathbf{B}_\beta) \\ \mathbf{b}_\beta &\sim N_p(\mathbf{b}_0, \mathbf{S}_0) \\ \mathbf{B}_\beta &= \text{diag}(\sigma_{\beta_1}^2, \dots, \sigma_{\beta_p}^2) \\ \sigma_{\beta_m}^2 &\stackrel{iid}{\sim} IG(a_{\beta_0}, b_{\beta_0}) \end{aligned}$$

for $m = 1, \dots, p$, where p is the number of covariates associated to the effects β . Here, $IG(\cdot)$ stands for inverse-gamma with shape parameters $a_{\beta_0} > 0$ and rate parameters $b_{\beta_0} > 0$. Where \mathbf{b}_0 and \mathbf{S}_0 are, respectively, the $p \times 1$ vector of location parameters and the $p \times p$ positive semi-definite covariance matrix of \mathbf{b}_β (i.e., the location parameter vector of the non varying effect β); \mathbf{b}_β and \mathbf{S}_β are, respectively, the $p \times 1$ location parameter and the $p \times p$ positive semi-definite covariance matrix for β (i.e., the non varying effect). The set $\{\mathbf{b}_0, \mathbf{S}_0, a_{\beta_0}, b_{\beta_0}\}$ of the hyperparameters are specified by the user. A diagonal matrix is suggested for \mathbf{S}_0 as showed by Heinzl et al. (2012).

- A diagonal structure for the $q \times q$ prior covariance matrix \mathbf{Q}_r for the hierarchical effects is specified as follows for each mixture component $r = 1, \dots, R$ and each each related covariate $d = 1, \dots, q$:

$$\mathbf{Q}_r = \text{diag}(\sigma_{Q_{1r}}^2, \dots, \sigma_{Q_{qr}}^2)$$

For the base measure G_0 ⁵ and the precision parameter α of the DP mixture model the following priors are specified:

$$\begin{aligned} G_0 &= N_q(\boldsymbol{\mu}_0, \mathbf{D}_0) \times IG(a_{Q_0}, b_{Q_0})^q \\ \boldsymbol{\mu}_0 &\sim N_q(\mathbf{m}_0, \mathbf{W}_0) \\ \mathbf{D}_0 &= \text{diag}(\sigma_{D_{01}}^2, \dots, \sigma_{D_{0q}}^2) \\ \sigma_{D_{0d}}^2 &\sim IG(a_{D_0}, b_{D_0}) \\ \alpha &\sim Ga(c_0^\alpha, C_0^\alpha) \end{aligned}$$

⁵ Assuming independence between the location and the scale parameters of each mixture component, and between all the scale parameters for each covariate $d = 1, \dots, q$, G_0 is then the product of the q -variate normal and the q inverse gamma distributions.

for $d = 1, \dots, q$, where q is the number of covariates associated to the hierarchical effects \mathbf{u}_i , and for $a_{D_0}, a_{Q_0} > 0$ and $b_{Q_0}, b_{D_0} > 0$. Here $Ga(\cdot)$ stands for Gamma distribution with $c_0^\alpha > 0$ and $C_0^\alpha > 0$ respectively the shape and the rate parameters. Where \mathbf{m}_0 and \mathbf{W}_0 are, respectively, the $q \times 1$ location parameter vector and the $q \times q$ positive semi-definite covariance matrix of $\boldsymbol{\mu}_0$ (i.e. the location parameter of the base measure G_0); $\boldsymbol{\mu}_0$ and \mathbf{D}_0 are, respectively, the $q \times 1$ location parameter vector and the $q \times q$ positive semi-definite covariance matrix of the base measure G_0 . The set $\{a_{Q_0}, b_{Q_0}, \mathbf{m}_0, \mathbf{W}_0, a_{D_0}, b_{D_0}, c_0^\alpha, C_0^\alpha\}$ of the hyperparameters need to be fixed. A diagonal structure is suggested for \mathbf{W}_0 as above.

- The following prior is assigned to the noise variance:

$$\sigma_\varepsilon \sim Ga(a_\varepsilon, b_\varepsilon)$$

with $a_\varepsilon > 0$ and $b_\varepsilon > 0$ hyperparameters fixed by the user as well.

4.2 Posterior sampling

Since most of the parameters in the model have conjugate prior distributions, a blocked Gibbs sampling algorithm was used for the posterior sampling (Ishwaran and James 2001).

The parameter vector for the model is $\theta = \{\beta, \mathbf{b}_\beta, \mathbf{B}_\beta, \boldsymbol{\mu}_0, \mathbf{D}_0, \mathbf{Q}, \sigma_\varepsilon, \alpha, \boldsymbol{\pi}, \mathbf{c}\}$ which is updated at each state of the Markov chain of the Gibbs sampling. Here $\mathbf{c} = (c_1, \dots, c_I)$ is the allocation parameter of the raters to the clusters and. Further details on the following sampling are given in the Appendix. The closed-form posteriors are as follows.

1. Update parameters referring to effects β :

$$\beta | \mathbf{b}_\beta, \mathbf{B}_\beta, \mathbf{u}, \sigma_\varepsilon, \mathbf{y} \sim N_p(\mathbf{b}_\beta^*, \mathbf{B}_\beta^*)$$

For each covariate $m = 1, \dots, p$ associated with a non varying effect β_m ,

$$b_{\beta_m} | \sigma_{\beta_m}^2, \beta_m \sim N \left(\left(\frac{1}{\sigma_{\beta_m}^2} + \frac{1}{s_{0_m}^2} \right)^{-1} \left(\frac{\beta_m}{\sigma_{\beta_m}^2} + \frac{m_{\beta_m}}{s_{0_m}^2} \right), \left(\frac{1}{\sigma_{\beta_m}^2} + \frac{1}{s_{0_m}^2} \right)^{-1} \right)$$

$$\sigma_{\beta_m}^2 | b_{\beta_m}, \beta_m \sim IG \left(a_{b_0} + \frac{1}{2}, \frac{1}{2} (\beta_m - b_{\beta_m}) \right)$$

2. Update parameters referring to hierarchical effects:

- For each rater $i = 1, \dots, I$:

$$\mathbf{u}_i | \boldsymbol{\mu}_{c_i}, \boldsymbol{\mu}_0, \mathbf{Q}_0, \beta, \sigma_\varepsilon, \mathbf{y}_i \sim N_q(\boldsymbol{\mu}_{c_i}^*, \mathbf{Q}_{c_i}^*),$$

where $\boldsymbol{\mu}_{c_i}$ is the location parameter vector of the cluster where the i -th rater is allocated.

- For each component $r = 1, \dots, R$ of the truncated mixture: - If $\nexists i : c_i = r$ (if no rater are currently allocated into cluster r), for each covariate $d = 1, \dots, q$ associated to an hierarchical effect (independently):

$$\begin{aligned} \boldsymbol{\mu}_r | \boldsymbol{\mu}_0, \mathbf{D}_0 &\sim N_q(\boldsymbol{\mu}_0, \mathbf{D}_0) \\ \sigma_{Q_{dr}}^2 &\sim IG(a_{Q_0}, b_{Q_0}) \end{aligned}$$

- If $\exists i : c_i = r$ (if at least one rater assigned to component r), for each covariate $d = 1, \dots, q$ associated to an hierarchical effect (independently):

$$\begin{aligned} \mu_{r_d} | \sigma_{Q_m}^2, \mu_0, \sigma_{D_{0_m}}^2, \mathbf{u}, \mathbf{c} &\sim N(\mu_0, \sigma_{D_{0_m}}^{2*}) \\ \sigma_{Q_{dr}}^2 | \boldsymbol{\mu}, \mathbf{u} &\sim IG(a_{Q_0}^*, b_{Q_0}^*) \end{aligned}$$

Essentially, at each iteration t , if the r -th cluster is empty the component location parameters μ_r are sampled from the prior as suggested by (Gelman et al. 2013), otherwise they are drawn from the above mentioned closed-form posterior.

- Each rater $i = 1, \dots, I$ is re-allocated into a cluster:

$$c_i | \boldsymbol{\pi}, \boldsymbol{\mu}, \mathbf{Q}, \mathbf{u}_i \sim \text{Cat}(\boldsymbol{\omega}_i^*)$$

where $\text{Cat}(\cdot)$ stands for Categorical distribution, and $\boldsymbol{\omega}_i^*$ is reported in the Appendix. A truncated approximation for the DPM mixture model was used (Gelman et al. 2013; Heinzl et al. 2012) for a large value of R . The stick-breaking construction was used to generate the mixture weights $\pi_{1:R}$.

- For each component $r = 1, \dots, R - 1$:

$$v_r | \boldsymbol{\Phi}, \alpha \sim \text{Be} \left(1 + c_r, \alpha + \sum_{l=c+1}^R r_l \right)$$

and $v_R = 1$ for the last cluster. Here c_r is the number of raters assigned to the cluster r , and r_l is the number of raters assigned to the cluster l .

- The precision parameter is updated as follows:

$$\alpha | v_1, \dots, v_{R-1} \sim \text{Ga} \left(R - 1 + a_\alpha, b_\alpha - \sum_{c=1}^{R-1} \ln(1 - v_r) \right)$$

- For each covariate $d = 1, \dots, q$ associated with an hierarchical effect the base measure parameters are updated:

$$\mu_{0_d} | \sigma_{D_{0_d}}^2, \boldsymbol{\mu} \sim N \left(\left(\frac{I}{\sigma_{D_{0_d}}^2} + \frac{1}{\sigma_{W_{0_d}}^2} \right)^{-1} \left(\frac{I}{\sigma_{D_{0_d}}^2} \bar{\mu}_d + \frac{m_{0_d}}{\sigma_{W_{0_d}}^2} \right), \left(\frac{I}{\sigma_{D_{0_d}}^2} + \frac{1}{\sigma_{W_{0_d}}^2} \right)^{-1} \right)$$

where $\bar{\mu}_d$ is the mean of the location parameters related to the d -th covariate over all the clusters.

$$\sigma_{D_{0m}}^2 | \mu_{0m}, \mathbf{u} \sim IG \left(a_{D_0} + \frac{I}{2}, b_{D_0} + \frac{1}{2} \sum_{i=1}^I (\mu_{i_m} - \mu_{0m})^2 \right)$$

3. Update the error variance:

$$\sigma_\epsilon^2 | \beta, \mathbf{u}, \mathbf{y} \sim IG \left(a_\epsilon + \frac{1}{2} I | \mathcal{J} |, b_\epsilon + \frac{1}{2} \sum_{i=1}^I \sum_{j \in \mathcal{J}_i} (y_{ij} - \mathbf{x}'_{ij} \beta - \mathbf{z}'_{ij} \mathbf{u}_i)^2 \right).$$

Here $|\mathcal{J}|$ is the cardinality of the set of all the rated items \mathcal{J} , it equals the number of observations.

5 The nonparametric λ index

The marginal posterior distribution of the hierarchical effects in the model outlined above captures information about the polarization or disagreement among raters (on the assumption that the model captures the data adequately). The ICC (i.e., intra-class correlation coefficient, (Martinková et al. 2023; Bartoš et al. 2020; Agresti 2015; Gelman et al. 2013)) might adequately quantify inter-rater variability if the normal distributional assumption of the rater hierarchical effect holds. Two assumptions are made computing the standard ICC considering a normal distributed hierarchical effect. Firstly, that the raters are sampled from the same population. Secondly, that possible different latent trends among raters are not interesting or eventually regarded as disagreement ratings. This might be a good first approximation of the rating process. Nevertheless, when more detailed considerations are needed, or subtle heterogeneity among raters is expected, the standard ICC might be less informative and inaccurate. Besides the latter issue, further information about the shape of the posterior might be quantified. For instance, in presence of a bimodal hierarchical effects distribution with two very distant modes (for example, when opinions are polarised), considering the posterior distribution of σ_u as an index of variability among raters might be misleading.

Several indexes have been proposed to quantify group opinion polarization (e.g., (Tang et al. 2022; Koudenburg and Kashima 2022; Koudenburg et al. 2021; Esteban and Ray 1994)) and to measure distribution bimodality (e.g., the Ashman's D (Forchheimer et al. 2015) or the bimodal separation index (Zhang et al. 2003)). The strong assumptions behind their use limit them to be valid options only in the parametric context or when the number of clusters is known. A model based nonparametric index is here proposed to overcome these limitations.

To this end the full estimated distribution of \mathbf{u} resulting from the model might be useful. At each iteration t , the density of \mathbf{u} is given by the corresponding mixture model given the parameters at iteration t . Following the formulation of (Gelman et al. 2013),

the set of modes and antimodes (i.e., the lowest frequent value between two modes) is identified. When the distribution of \mathbf{u} is multimodal, the latent polarization (disagreement) λ is then defined as the log ratio between the mean density of the modes and the that of the anti-modes, it is zero when it is unimodal:

$$\lambda = \begin{cases} \log \left(\frac{\frac{1}{M} \sum_{m=1}^M f_u(\gamma_m)}{\frac{1}{M-1} \sum_{m=1}^{M-1} f_u(\zeta_m)} \right), & \text{if } M > 1 \\ 0, & \text{otherwise.} \end{cases}$$

Where M is the number of modes γ_m , $m = 1, \dots, M$ and the number of antimodes ζ_m , $m = 1, \dots, M - 1$ of the density of \mathbf{u} ; $f_u(\cdot)$ denotes the density at a specific point. Larger values of λ indicate strongly multimodal distribution of the hierarchical effects, whereas smaller values are evidence of weak multimodality, thus the estimated hierarchical effects are less concentrated.

As it is shown in Fig. 1 larger values of λ indicate distribution polarization, whereas smaller values indicate a less concentrated and more spread density distribution. The λ index is strongly affected by both location and scale parameters of the mixture components. For this reason it might be very informative in presence of multimodal distributions. Assuming such a raters' group polarization as a result of low latent agreement among raters, the λ index might be a useful diagnostic tool.

6 Simulation studies

The following simulations aim to evidence how the values of λ varying across different polarization settings. The first simulation investigates the role of the precision parameter α and the variance of the mixture components in determining the values of λ . The second one shows the complementary role of λ in the inter-rater agreement analysis and how this index varies across different settings.

6.1 Simulation 1: DPM and λ

6.1.1 Simulation setting

The first simulation study explores the role that the precision parameter of the Dirichlet Process and the variance of the components have in determining the values of the log-density index λ . For simplicity purpose the mixture components are assumed to have the same variance Q in this simulation, so the component subscription will be omitted. The objective is to study the effect of α and Q , on λ conditional on all the other variables. Since the former has a crucial role in the determination on the point masses of G , and thus the concentration of its realizations, an inverse relation between α and λ is expected if Q is fixed. Likewise, an inverse relation between Q and λ is expected if α is fixed. It is interpretable as an index of the sharpness of the modes. For this reason both the precision parameter of the DPM and the variance

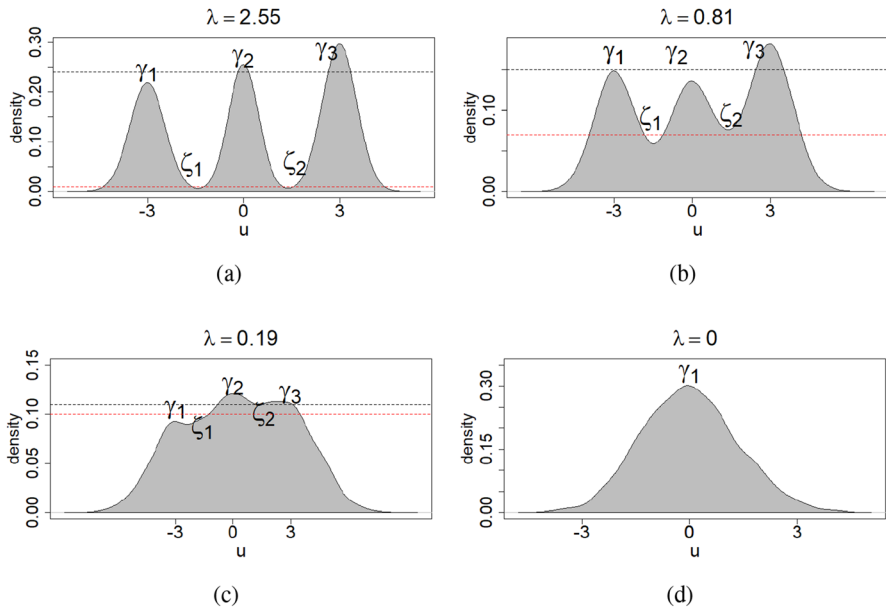


Fig. 1 Different values of λ indicate different polarization levels. Three different values of λ were computed for three different mixture distributions, respectively. The realizations of these distribution are here referred to as u . Black dotted lines indicate the mean mode density, red dotted lines indicate the mean antimode density. **a** High polarization: the mixture components are highly and clearly separate, the mean density values of the modes is far larger then the mean density value of the antimodes; the log-density ratio between these two quantities is $\lambda = 2.55$. **b** Medium polarization: the mixture components are clearly separated, but the mean density values of the modes is closer to the mean density value of the antimodes; the log-density ratio between these two quantities is $\lambda = 0.81$. **c** Low polarization: the mixture components are not clearly separated, the mean density values of the modes is very close to the mean density value of the antimodes; the log-density ratio between these two quantities is $\lambda = 0.19$. **d** No polarization: the mixture distribution has only one mode (i.e., γ_1) and $\lambda = 0$ since the number of mode is not greater then one

of its components are expected to have an effect on λ . Controlling for Q (i.e., keeping it fixed), the expected relation is: the smaller α , i.e. the precision of the DPM mixture, the larger λ , i.e. the relative density around the modes; controlling for α (i.e., keeping it fixed), the expected relation is: the smaller Q , i.e. the variance of the components of the DPM, the larger λ . The parameters of the base measures G_0 have a non-negligible role in determining λ , so in this section focus is devoted to the relation between the precision parameter α , the mixture components variance Q and the index λ . Indeed, in all the study simulations the values of the other parameters involved in the DPM have been kept fixed across the scenarios.

Data generating process

The experimental design is as follows. For 4 different values of $\alpha = (0.1, 1, 5, 20)$ and 2 different values of $Q = (0.1, 1.5)$ a set of independent observations $u = 1, \dots, n$ are drawn from the following DPM:

$$\begin{aligned}
 u_i | \mu_c, Q, &\sim \sum_{c=1}^R \pi_c N(\mu_c, Q), \quad i = 1, \dots, n \\
 \mu_c &\stackrel{iid}{\sim} G_0 \\
 \pi_c &= v_c \prod_{l < c} (1 - v_l), \quad \text{where} \\
 v_c &\stackrel{iid}{\sim} Be(1, \alpha), \quad c = 1 \dots, R.
 \end{aligned}$$

Where μ_c and π_c are the location parameter and the mixing proportion of the component c , respectively; G_0 is the base measure; and v_c is the parameter of the stick-breaking. Following the above mentioned truncated stick-breaking construction, here R is the maximum number of observable cluster. Across the eighth scenarios the following quantities are assigned: the number of observations $n = 500$, the maximum number of clusters $R = 50$, the base measure $G_0 : U(-6, 6)$. Here, $U(\cdot)$ stands for uniform distribution. The use of these distributions in the present experimental context aims to highlight the effect of different values of α and Q on λ in a more evident and interpretable manner.

6.1.2 Results

As shown in Tables 1 and 2 as α increases, and so the number of point masses of G increases as well, λ decreases. The density of the observations $u = 1, \dots, n$ is concentrated around few point masses (few modes) for lower value of α and is spread out the larger. Note also the change of density of the antimodes. As expected, it is proportional to the precision parameter in a positive fashion. As the observations are more spread as α increases, there are fewer intervals in the support with relative small density: λ index decreases at larger values of α (column-wise Tables 1 and 2). A similar proportional relation is observed between the variance of the mixture components Q and λ when α is kept fixed (row-wise Tables 1 and 2). Smaller values of both α and Q result in a high polarized distribution of $u = 1, \dots, n$ and correspond to larger values of λ . Whereas larger values of both α and Q result in a low polarized distribution and correspond to smaller values of λ . It is an index of how spread the density is over the support of the hierarchical effects.

From an interpretative point of view, λ indicates the degree of overlap between the infinitely many clusters. It might be informative of the separation between them. Since this quantification is based on a non-parametric density, λ is not directly related to the number of the group, or to the cluster location. It indicates the degree to which the independent observations drawn from a DPM overlap; the variance of the cluster Q also plays a crucial role. The index thus quantifies the combined effect of the parameters to assess the extent to which possible different opinions (i.e., the modes) might be strongly shared among the raters (i.e., the modes are sharp pick of density). To this regard, λ is a polarization index in presence of heterogeneity. The higher the polarization levels, the larger the values of the index. The practical interpretation and the operational decisions must be guided by the field of application.

6.2 Simulation 2: inter-rater agreement and λ

6.2.1 Simulation setting

The following simulation study aims to highlight the complementary role of λ as an additional summary metric in inter-rater agreement analysis. The varying intercept parametrization is hereafter adopted as univariate case for the raters effects. To this aim the standard modelling approach (i.e., the normal distributed varying intercept and the resulting ICC) is compared with the nonparametric proposed above (i.e., the DPM prior over the varying intercept and λ). The experiment evaluates the performance of both standard ICC and λ in the presence of heterogeneity between raters' evaluations due to a multimodal distribution of the hierarchical effects.

Data generating process

Three experimental scenarios were planned, in which a different clustering on the raters' intercept parameter was specified in the generative model. In each scenario the rater's intercept u_i was generated from a bimodal Gaussian mixture. The location parameters of the mixture components were fixed across the scenarios, $\mu_1 = -3$ and $\mu_2 = 3$; whereas decreasing values (1, 0.5, 0.1) were assigned to the components scale parameters Q_1 and Q_2 (see Table 3). This resulted in different polarization scenarios. The mixture components were kept equiprobable ($\pi_c = 0.5$), $c \in \{1, 2\}$ throughout. The number of raters $I = 100$ and the number of items $J = 250$ were fixed across the scenarios. One continuous covariate x_{ij} with an effect $\beta = 2$ was used and it was the same across the scenarios.

Standard model approach

The following priors were specified for the standard hierarchical effect model (i.e., the varying intercepts are assumed to be i.i.d. normal distributed):

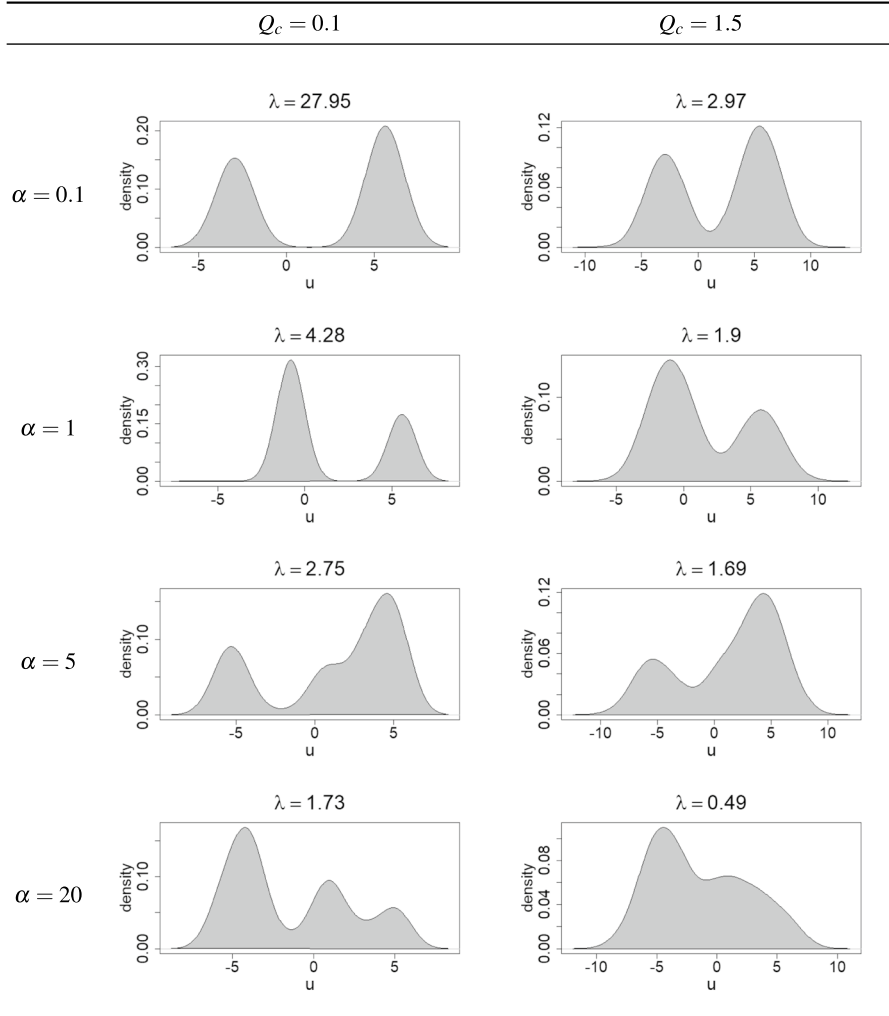
$$\begin{aligned}\beta &\sim N(0, 5), \\ \sigma_\varepsilon &\sim \text{Exp}(0.2), \\ \sigma_u &\sim \text{Exp}(0.2), \\ u_i &\stackrel{iid}{\sim} N(0, \sigma_u),\end{aligned}$$

for $i = 1, \dots, I$; $\text{Exp}(\cdot)$ stands for the exponential distribution and β is the non-hierarchical effect, σ_u and σ_ε are the hierarchical effect and the noise variances parameters, respectively. A logic of complexity penalization was used in the choice of the above mentioned priors distributions (Simpson et al. 2017). The posterior of each standard hierarchical effect model were sampled using NUTS-Hamiltonian MCMC in Stan language (Stan Development Team 2022).

Nonparametric model approach

The set of priors introduced in Sect. 4 were elicited for the DPM models with the following hyperparameters as suggested by (Heinzl et al. 2012): $\mathbf{b}_0 = \mathbf{0}$, $\mathbf{S}_0 = 1000\mathbf{I}_p$, $a_{\beta_0} = 0.005$, $b_{\beta_0} = 0.005$, $\mathbf{m}_0 = \mathbf{0}$, $\mathbf{W}_0 = 100\mathbf{I}_q$, $a_{D_0} = 0.5$, $b_{D_0} = 0.5$, $a_{Q_0} = 0.001$, $b_{Q_0} = 0.001$, $a_\alpha = 2$, $b_\alpha = 2$, $a_\varepsilon = 0.005$, $b_\varepsilon = 0.005$. As result of some preliminary analysis, a dense grid of 481 equally-spaced values from -12 to 12 (i.e., with a fixed interval of 0.05) was used to monitoring the mixture density of the

Table 1 The eight scenarios correspond to a DPM with different values of the precision parameter α and the components variance Q . Each scenario correspond to a specific combination of these two parameters. All the other quantities are fixed across the scenarios. The realizations of the DPM are here indicated as $u = 1, \dots, n$. Different combinations of α and Q result in different values of λ . For fixed values of Q (column-wise), a proportional relation is shown between α and λ : when the first increases, the second decreased. Similarly, for fixed values of α (row-wise), a proportional relation is shown between Q and λ : when the first increases, the second decreased. Smaller values of both α and Q result in a high polarized distribution of $u = 1, \dots, n$ and correspond to larger values of λ . Whereas larger values of both α and Q result in a low polarized distribution and correspond to smaller values of λ



nonparametric varying intercept u_i at each iteration. The posterior distribution of the nonparametric hierarchical effect is obtained as the set of the mean density of each point of the grid over the iterations (Gelman et al. 201320132013).

Table 2 Parameters values at each scenarios. Each of them correspond to a DPM with different values of the precision parameter α and the components variance Q . Both α and Q have an effect on the distribution polarization of the realizations of the DPM. As a result different values of λ are observed

	α	Q	λ
Scenario 1	0.1	0.1	27.95
Scenario 2	0.1	1.5	2.97
Scenario 3	1	0.1	4.28
Scenario 4	1	1.5	1.9
Scenario 5	5	0.1	2.75
Scenario 6	5	1.5	1.69
Scenario 7	20	0.1	1.73
Scenario 8	20	1.5	0.49

Table 3 True raters' hierarchical effects distribution across different scenarios. A Gaussian mixture is specified as distribution of the hierarchical effects $u_i = 1, \dots, I$. The location parameters of two components of the mixture are kept fixed across the scenarios and decreasing values were assigned to the respective scale parameters

Scenario 1	$u_i \stackrel{iid}{\sim} 0.5 \cdot N(-3, 1) + 0.5 \cdot N(3, 1)$
Scenario 2	$u_i \stackrel{iid}{\sim} 0.5 \cdot N(-3, 0.5) + 0.5 \cdot N(3, 0.5)$
Scenario 3	$u_i \stackrel{iid}{\sim} 0.5 \cdot N(-3, 0.1) + 0.5 \cdot N(3, 0.1)$

Table 4 95% HPD intervals of the hierarchical effects variance σ_u from the standard models (i.i.d. normal distributed varying intercepts) and of the Grid density of the hierarchical effect in DPM models

	σ_u	Grid density
Scenario 1	(2.80, 3.75)	$(-4.60, -1.60) \cup (1.40, 4.60)$
Scenario 2	(2.65, 3.55)	$(-3.95, -1.85) \cup (1.75, 4.35)$
Scenario 3	(2.57, 3.46)	$(-3.70, -2.30) \cup (2.30, 3.60)$

In all the computations for both the models 55,000 iteration with 5000 burn-in were used, the Markov chains were thinned the by a factor of 50, resulting in samples of size 1000 (Heinzel et al. 2012).

6.2.2 Results

As shown in Table 4 the standard model (i.e., that in which hierarchical raters intercepts are assumed to be i.i.d. normally distributed) due to the rigid distributional assumption of the hierarchical parameters is not able to capture the possible multimodal distribution and it resulted in a large value of the hierarchical effect variance σ_u (see Table 4 and Fig. 2). As a result, the ICC didn't capture almost any difference among the three different scenarios (see Table 4 and Fig. 3). On the contrary, the DPM model, due to the flexible nonparametric specification of the intercepts prior,

Table 5 95% HPD intervals of the ICC from the standard models (i.i.d. normal distributed varying intercepts) and of λ from the DPM models

	ICC	λ
Scenario 1	(0.952, 0.973)	(1.22, 6.33)
Scenario 2	(0.948, 0.970)	(1.14, 10.69)
Scenario 3	(0.945, 0.969)	(0.05, 31.94)

Table 6 95% credible intervals of β , DPM and residuals related parameters. Here β is the non varying effect, b_β and σ_β are, respectively, the related location and scale hyperparameters; μ_0 and σ_{D_0} are the location and scale parameters of the base measure G_0 , respectively. The precision parameter α and the residuals standard deviation σ_ϵ are also reported

	Scenario 1	Scenario 2	Scenario 3
β	(1.87, 2.12)	(1.87, 2.13)	(1.87, 2.12)
b_β	(-7.15, 11.03)	(-7.26, 10.58)	(-7.44, 10.97)
σ_β	(0.11, 91.40)	(0.11, 88.28)	(0.10, 87.69)
μ_0	(-0.16, 1.16)	(0.12, 1.33)	(-0.49, 0.70)
σ_{D_0}	(6.63, 12.63)	(6.09, 11.30)	(6.52, 11.68)
σ_ϵ	(0.43, 0.46)	(0.42, 0.46)	(0.42, 0.46)
α	(8.05, 18.06)	(8.15, 18.09)	(8.12, 18.03)




Fig. 2 95% HPD intervals of σ_u from the standard models (i.i.d. normal distributed varying intercepts): **a** Scenario 1: lower polarization. **b** Scenario 2: medium polarization. **c** Scenario 3: higher polarization. Scenario 3: higher polarization. These models, due to their rigid distributional assumption, poorly differentiate the three different polarization scenarios

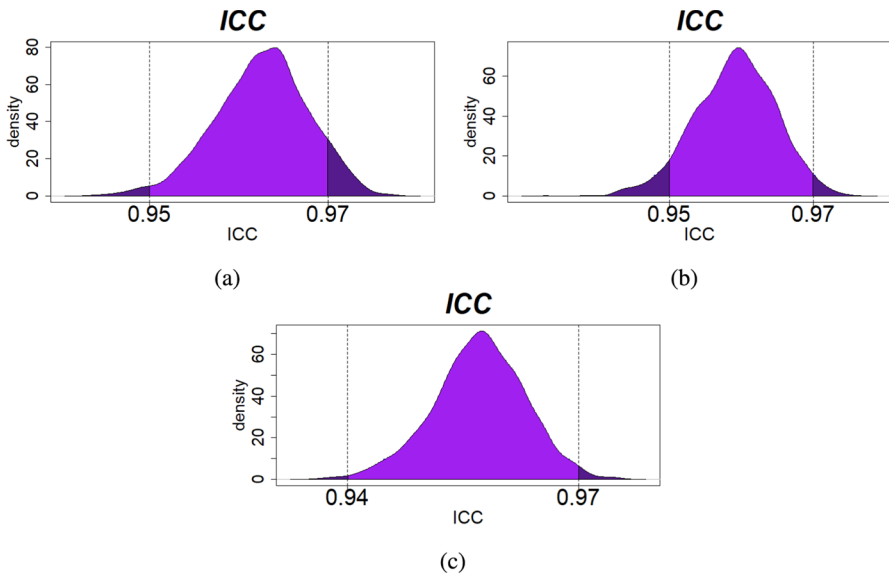


Fig. 3 95% HPD intervals of ICC from the the standard models (i.i.d. normal distributed varying intercepts): **a** Scenario 1: lower polarization. **b** Scenario 2: medium polarization. **c** Scenario 3: higher polarization

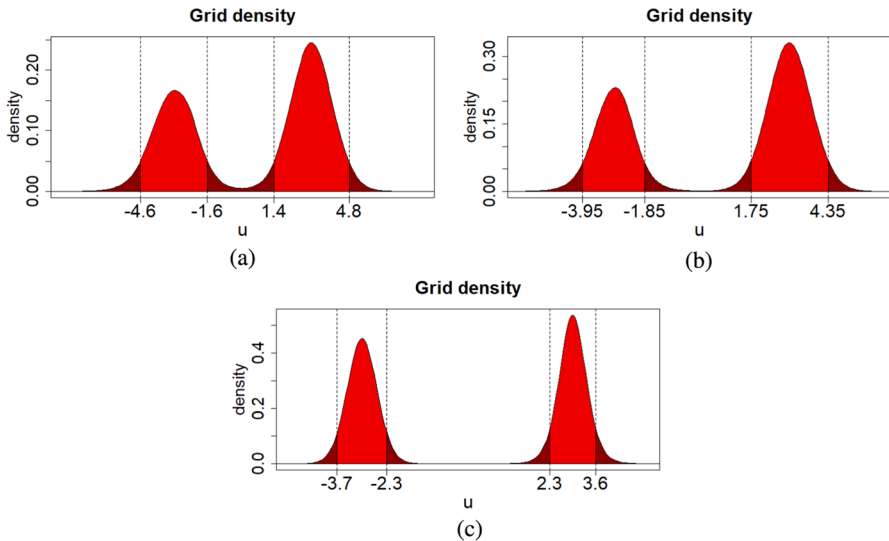


Fig. 4 95% HPD intervals of the hierarchical effect from the DPM. The different mixture used to generate the data emerged clearly from the posterior of the grid adopted to monitoring u . **a** Scenario 1. **b** Scenario 2. **c** Scenario 3

showed a good performance. As evident from Fig. 4 and Table 4 the DPM model was far more able to reproduce the data generating process. The different mixture

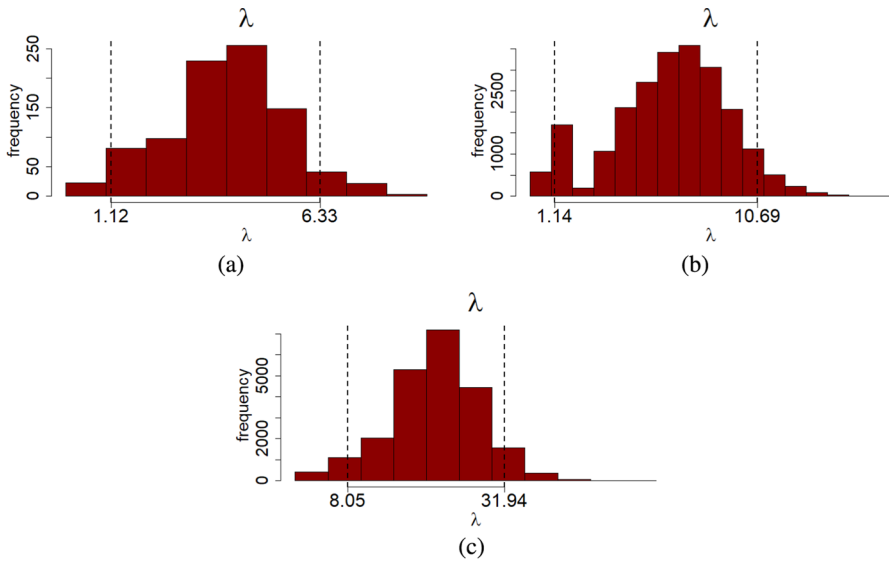


Fig. 5 Posterior distribution of λ : **a** Scenario 1: lower polarization. **b** Scenario 2: medium polarization. **c** Scenario 3: higher polarization. The black dotted lines stands for 95% credible intervals

Table 7 95% HPD intervals of the other parameters of the standard models (i.i.d. normal distributed varying intercepts)

	Scenario 1	Scenario 2	Scenario 3
β	(1.99, 2.00)	(1.98, 2.01)	(1.99, 2.01)
σ_ϵ	(0.60, 0.64)	(0.62, 0.63)	(0.62, 0.63)

used to generate the data emerged clearly from the posterior of the grid adopted to monitoring \mathbf{u} (Fig. 4). Since the DPM model properly learn the multi-modalities of the raters intercepts density, the index λ , being based on the ratio between the mean density of the modes and that of the antimodes present in the grid at each iteration, showed to be able to differentiate the three different polarization scenarios. It gives some interesting information regarding the shape of the non-parametric mixture distribution. The 95% credible interval of λ (see Table 5 and Fig. 5) as estimated in the three different scenarios highlighted different degrees of amplitude and separation (i.e., different degrees of polarization) along them. The index λ is computed as a logarithm of the ratio of the average mode density against the density of the antimodes at each iteration t of the posterior sampler. So, in the first scenario, the values of the 95% credible interval are smaller, indicating that in most of the iterations the difference between the mean density at the modes and that of the antimodes was very small. In terms of the third scenario, λ assumed rather larger values along the iterations as evidence that the mode density is far larger than that of the antimodes. In other words, the rater clusters were separated and clearly distinct. The parameters of

the DPM are the most influential with regard to λ . Specifically, the location parameters μ_c , $c = 1, \dots, R$, and the scale parameters Q_c , $c = 1, \dots, R$, showed to have a combined effect of the proposed index. The 95% HDP intervals of the parameters of both the DPM prior model and that with normal distributional assumption are reported in Tables 6 and 7, respectively.

7 Large scale performance assessment

The evaluation heterogeneity of teachers is a long-standing issue in psychometrics (Uto 2022; Shirazi 2019; Bonefeld and Dickhäuser 2018; Casabianca et al. 2015; DeCarlo 2008). Highly biased scored might have a detrimental effect on students proficiency and education (Chin et al. 2020; Paredes 2014; Cooper 2003). The proposed nonparametric model and the index λ might be valuable tools to address this issue. They might help to shed light on very biased assessment contexts and to provide fairer scores. The estimated hierarchical effect of each teacher (which may be interpreted as the teacher's bias) might be used to adjust the observed score. The index λ might quantify teachers polarization in their grading.

The Matura data set

As an illustrative real data application, a large scale performance assessment data set was analysed (Zupanc and Štrumbelj 2018). The DPM-model was applied to a large-scale essay assessment data obtained during the nation-wide external examination conducted by the National Examination Centre in upper secondary schools in Slovenia also known as *Matura* and analyzed in Zupanc and Štrumbelj (2018). These data were related to the spring term argumentative essays for years between 2010 and 2014. Particular attention is devoted to the distinction between two main aspects of essay writing: the language correctness (i.e., the presence of grammatical or syntactic errors) and the the good argumentation of the content (i.e., a good and clear presentation of all the arguments). Regarding the data structure, students are nested within the teachers. So that each student's essay is evaluated by one trained teacher, who is asked to grade it concerning two different rubrics. An essay can receive a score between 0 and 20 for the language-related rubric and between 0 and 30 for the content-related one. Prior analysis of these data (Zupanc and Štrumbelj 2018) revealed that heterogeneity among teachers was broadly down to two types: strict and lenient. The two different trends might be captured by the model and their polarization quantified by the λ index.

For this reason $N=2616$ students' essays, each scored by one of $I=18$ different teachers, were considered for the analysis.⁶ The objective of this application is to analyze teachers' individual differences in scoring the essay content, controlling for its language correctness. How lenient or strict they are in scoring the quality of an essay content, without the effect of the language correctness. The content score is commonly ways more susceptible to idiosyncrasies or biases of the teacher than the language-related score, which is generally more objective (Childs and Wooten 2023; Zhu et al. 2021; Shirazi 2019). Accordingly, the content-related score was specified

⁶ For illustrative purposes, only the variables related to the first teachers were considered.

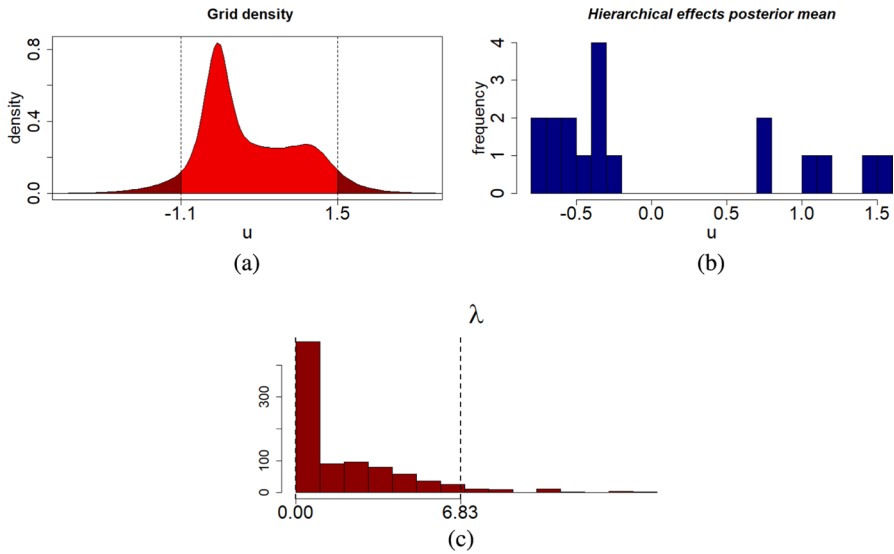


Fig. 6 **a** 95% HPD of the monitoring grid for the teachers' hierarchical effects $u = 1, \dots, I$. **b** Posterior mean of the hierarchical effect of each teacher. Two different clusters emerged from the analysis, as expected: the more lenient (to the right-hand side) and the stricter (to the left-hand side). **c** λ ' 95% credible intervals. It indicate a moderate polarized posterior distribution of the posterior hierarchical effects

as outcome variable and the language-related score as covariate with a non varying effect. A DPM hierarchical prior was specified over the teachers' intercepts. All the scores were re-scaled for this analysis to get a easier parameters value interpretation⁷ (Gelman et al. 2013).

7.1 Results

The language-related score showed a posterior mean effect of 0.27 on the content-related score, with a (0.16, 0.38) 95% credible interval. The language correctness of the essay writing had moderate role in predicting the evaluation of the its content. As shown by Fig. 6a the DPM-model learned the presence of two main trends from the data. The bimodal non-parametric distribution over the grid suggested that the teachers were rather heterogeneous in the essay scoring process. More precisely, they seemed to be slightly polarized around two main tendencies. Some teacher showed a slightly more lenient or stricter than the others (i.e., who had a larger or smaller hierarchical effect posterior mean, respectively), see Fig. 6. The λ index showed a posterior mean of 1.87 which suggested a low polarization. The λ 95%

⁷ The following transformation was applied to standardize the score: $f(x) = \frac{x - \bar{x}}{\hat{\sigma}_x}$, where $\bar{x} = \frac{1}{N} \sum_{n=1}^N x_n$ was the sample mean and $\hat{\sigma}_x = \sqrt{\frac{1}{N} \sum_{n=1}^{N-1} (x_n - \bar{x})^2}$ the sample standard deviation.

Table 8 95% credible intervals of β , DPM and residuals related parameters. Here β is the non varying effect, b_β and σ_β are, respectively, the related location and scale hyperparameters; μ_0 and σ_{D_0} are the location and scale parameters of the base measure G_0 , respectively. The precision parameter α and the residuals standard deviation σ_ε are also reported

β	(0.16, 0.38)
b_β	(−14.80, 20.32)
σ_β	(0.11, 98.33)
μ_0	(−0.33, 0.34)
σ_{D_0}	(0.19, 1.03)
σ_ε	(0.13, 0.15)
α	(3.00, 4.17)

HPD interval was (0.0, 6.83) which indicated a non negligible occurrence of quite high values of λ . All the other parameters credible intervals are reported in Table 8

Assuming this latent group polarization as a low latent agreement among raters, the λ index might be used in a diagnostic manner. Considering the present application, some solutions might be suggested for a fairer assessment process. Firstly, assuming a very negligible noise term, the teacher's estimated bias might be removed from the actual score. Another practical solution might be the implementation *ad hoc* training aimed to a much more shared point of view in essay scoring.

8 Conclusions

Most of the statistical models commonly used to analyze data from such observational contexts haven't shown to be very flexible to certain types of heterogeneity among raters. The common HLMs with a normal (or unimodal) distributional assumption for the hierarchical effects cannot capture any possible latent clusters, i.e. any multimodality. Indeed, the residual covariance modelled through the hierarchical effects might be informative about different latent similarities among raters. In this regard, incorporating a DPM in the prior of the hierarchical effects distribution is a flexible choice to address this issue.

Consequently, the estimation of the agreement among the raters should take into account the possible multimodal distribution of the hierarchical effects. Interest might not be exclusively on the proportion of variance attributable to the hierarchical effects over the total variance (i.e., the main interpretation of the ICC); instead, it might be more appealing to explore the entire multimodal density. Since the DPM naturally accommodates clusters among hierarchical effects (i.e., among raters), it is natural to consider the extent to which the mixture components are separated. Since λ is based on the density approximated through the grid approach $f(\mathbf{u})$, it reflects both the clustering induced by the Dirichlet process and the variance of the mixture components. Due to the particular information carried by λ it might be more informative about the latent

agreement among raters than the solely ICC. The latter is very useful when the normal distributional assumption of the hierarchical effects holds. However, in the presence of multimodality the estimate of the variance of the hierarchical effect σ_u is not accurate (it might be over-estimated) and the related ICC might be non-informative.

In contexts in which strong beliefs about the exact number of cluster are present or it is supported by some sort of evidence, an hierarchical model with a prior finite mixture distribution over the hierarchical effects is expected to have comparably good performance as well. The parametric variance of a mixture might be take into account in the ICC formula in these cases. For the above mentioned reasons, added flexibility and the shrinkage property the DPM was here preferred.

Many other studies are needed to fully understand the performance of λ across different combinations of the Dirichlet process parameters. Future works might highlight the role of λ when the rating is either expressed on a dichotomous or on a polytomous scale. Further studies might highlight the computation of λ when multivariate hierarchical effects are specified. Comparisons between this index and the others widely used in these cases (Tang et al. 2022; Forchheimer et al. 2015; Zhang et al. 2003) might be a focus of future studies. Further application of λ in a non-parametric context might be studied (Canale and Prünster 2017).

Appendix

Remarks for multiple ratings

When raters rate the same set of items $\mathcal{J}_i = \mathcal{J}$, $i = 1, \dots, I$ a varying intercept can be identified for each item (Martinková et al. 2023; Agresti 2015; Nelson and Edwards 2015). These term might be added to Eq. (1) (which is the same in both the standard and nonparametric formulation):

$$y_{ij} = \mathbf{x}'_{ij}\beta + \mathbf{z}'_i\mathbf{u}_i + \delta_i + \varepsilon_{ij}, \quad i = 1, \dots, I, j \in \mathcal{J}. \quad (2)$$

In both the standard HLM (i.e., assuming a multivariate normal distributed hierarchical rater effect) and the nonparametric HLM (i.e., specifying a DPM over the rater effect) the following distribution might be specified:

$$\delta_j \sim N(0, \sigma_\delta^2) \quad j = 1, \dots, J.$$

where $\sigma_\delta > 0$ is the scale parameter of δ and $J = |\mathcal{J}|$. See Sects. 2 and 3 for the other quantities and their distribution assumption. Specifying a conjugate prior for σ_δ additional steps might be added to the Gibbs sampling for the nonparametric HLM.

The main results of the present work and the interpretation of λ (see Sect. 5) still hold for this model specification.

Details on Dirichlet process mixture

As noticed above, α is proportional to the *concentration* of the realizations of G in point masses. Indeed, considering the partition (A, A^c) of Ω , the variance of $G(A)$ is defined as

$$\text{Var}[G(A)] = \frac{G_0(A)(1 - G_0(A))}{\alpha + 1}$$

Thus, larger values of α , conditioning on the number of raters I , reduce the variability of the DP, i.e. the process samples most of the time from G_0 , G tends to be an infinite number of point masses: the empirical distribution of G tends to become a discrete approximation of the parametric G_0 . In this case there is no a strong clustering since the probability of ties is very low. On the contrary, smaller values of α induce a strong clustering, the random weights distribution concentrate the probability mass to few points of the support of G and the probability of ties is higher. Which in the present model means that several u_i will be independent and identically distributed from a normal distribution indexed by the same parameters. Moreover, Antoniak (Antoniak 1974) demonstrated that

$$\mathbb{E}[C|I, \alpha] \approx \alpha \ln\left(\frac{I + \alpha}{\alpha}\right)$$

where C is the number of clusters. Thus, the expected number of point masses of G is proportional to both the α and the number of raters I . Every consideration regarding the role of the precision parameter on the distribution of G should be conditioned to I .

Details on the Gibbs sampling

Further details regarding some parameters of the posterior sampling are showed as follow.

The following matrix notation is here adopted: $\mathbf{X}_i = (\mathbf{x}'_{i1}, \dots, \mathbf{x}'_{i|\mathcal{J}_i|})$, $\mathbf{Z}_i = (\mathbf{z}'_{i1}, \dots, \mathbf{z}'_{i|\mathcal{J}_i|})$, are the design matrices for each rater $i = 1, \dots, I$; and $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_I)$ and $\mathbf{Z} = \text{diag}(\mathbf{Z}_1, \dots, \mathbf{Z}_I)$ are the full design matrices.

1. Referring to the non varying effects:

$$\mathbf{b}_\beta^* = \left(\mathbf{B}_\beta^{-1} + \frac{1}{\sigma_\epsilon^2} \mathbf{X}'\mathbf{X} \right)^{-1} \left(\mathbf{B}_\beta^{-1} \mathbf{b}_\beta + \frac{1}{\sigma_\epsilon^2} \mathbf{X}'(\mathbf{y} - \mathbf{Z}\mathbf{u}) \right)$$

$$\mathbf{B}_\beta^* = \left(\mathbf{B}_\beta^{-1} + \frac{1}{\sigma_\epsilon^2} \mathbf{X}'\mathbf{X} \right)^{-1}$$

2. Referring to hierarchical effects:

- For each rater $i = 1, \dots, I$:

$$\boldsymbol{\mu}_{c_i}^* = \left(\mathbf{D}_0^{-1} + \frac{1}{\sigma_\varepsilon^2} \mathbf{Z}'_i \mathbf{Z}_i \right)^{-1} \left(\mathbf{D}_0^{-1} \boldsymbol{\mu}_{c_i} + \frac{1}{\sigma_\varepsilon^2} \mathbf{Z}'_i (\mathbf{y}_i - \mathbf{X}_i \boldsymbol{\beta}) \right)$$

$$\mathbf{Q}_{c_i}^* = \left(\mathbf{D}_0^{-1} + \frac{1}{\sigma_\varepsilon^2} \mathbf{Z}'_i \mathbf{Z}_i \right)^{-1}$$

Here $\boldsymbol{\mu}_{c_i}$ is the location parameter vector of the cluster where the rater i is allocated.

- For each component $r = 1, \dots, R$ and each variable $d = 1, \dots, q$, associated with an hierarchical effect:

$$\mu_{0_r}^* = \left(\frac{c_r}{\sigma_{Q_d}^2} + \frac{1}{\sigma_{D_{0_d}}^2} \right)^{-1} \left(\frac{c_r}{\sigma_{Q_d}^2} \bar{u}_{d,r} + \frac{\mu_{0_r}}{\sigma_{D_{0_d}}^2} \right)$$

$$\sigma_{D_{0_d}}^{2*} = \left(\frac{c_r}{\sigma_{Q_d}^2} + \frac{1}{\sigma_{D_{0_d}}^2} \right)^{-1}$$

$$\sigma_{Q_{dr}}^2 | \boldsymbol{\mu}, \mathbf{u} \sim IG \left(a_{Q_0} + \frac{r_c}{2}, b_{Q_0} + \frac{1}{2} \sum_{i=1}^{r_c} (u_{i_d} - \mu_{dr})^2 \right)$$

Here \bar{u}_{dr} is the mean of the d -th hierarchical effect in the cluster r .

- For rater $i = 1, \dots, I$ and each component $r = 1, \dots, R$:




Fig. 7 Trace Plots of some parameters from the second scenario. As it is shown they all converge properly

$$\omega_{i_r}^* = \frac{\pi_r N_q(\mathbf{u}_i | \boldsymbol{\mu}_r, \mathbf{Q}_r)}{\sum_{r=1}^R \pi_r N_q(\mathbf{u}_i | \boldsymbol{\mu}_r, \mathbf{Q}_r)}.$$

Some trace plots

See Fig. 7.

Funding Open access funding provided by Università degli Studi di Padova. The authors have no relevant financial or non-financial interests to disclose.

Declarations

Conflict of interest The authors have no conflicts of interest to declare that are relevant to the content of this article. All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no financial or proprietary interests in any material discussed in this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Agresti A (2015) Foundations of linear and generalized linear models
- Antoniak CE (1974) Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *Ann Stat* 2(6):1152–1174
- Barneron M, Allalouf A, Yaniv I (2019) Rate it again: using the wisdom of many to improve performance evaluations. *J Behav Dec Mak* 32(4):485–492
- Bartoš F, Martinkova P, Brabec M (2020) Testing heterogeneity in inter-rater reliability, pp 347–364
- Blackwell D (1973) Discreteness of Ferguson selections. *Ann Stat* 1(2):356–358
- Blackwell D, MacQueen JB (1973) Ferguson distributions via Polya urn schemes. *Ann Stat* 1(2):353–355
- Bonefeld M, Dickhäuser O (2018) (Biased) grading of students' performance: students' names, performance level, and implicit attitudes. *Front Psychol* 9:481
- Bouchard-Côté A, Doucet A, Roth A (2017) Particle Gibbs split-merge sampling for Bayesian inference in mixture models. *J Mach Learn Res* 18:1–39
- Briesch A, Hemphill E, Volpe R, Daniels B (2014) An evaluation of observational methods for measuring response to Classwide intervention. In: *School psychology quarterly: the official journal of the Division of School Psychology, American Psychological Association*, vol 30
- Bygren M (2020) Biased grades? Changes in grading after a blinding of examinations reform. *Assess Eval Higher Educ* 45(2):292–303
- Canale A, Dunson DB (2011) Bayesian kernel mixtures for counts. *J Am Stat Assoc* 106(496):1528–1539 (PMID: 22523437)

- Canale A, Prünster I (2017) Robustifying Bayesian nonparametric mixtures for count data. *Biometrics* 73(1):174–184
- Cao J, Stokes SL, Zhang S (2010) A Bayesian approach to ranking and rater evaluation: an application to grant reviews. *J Educ Behav Stat* 35(2):194–214
- Casabianca JM, Lockwood JR, Mccaffrey DF (2015) Trends in classroom observation scores. *Educ Psychol Meas* 75:311–337
- Childs TM, Wooten NR (2023) Teacher bias matters: an integrative review of correlates, mechanisms, and consequences. *Race Ethn Educ* 26(3):368–397
- Chin MJ, Quinn DM, Dhaliwal TK, Lovison VS (2020) Bias in the air: a nationwide exploration of teachers' implicit racial attitudes, aggregate bias, and student outcomes. *Educ Res* 49(8):566–578
- Cicchetti DV (1976) Assessing inter-rater reliability for rating scales: resolving some basic issues. *Br J Psychiatry* 129(5):452–456
- Cooper CW (2003) The detrimental impact of teacher bias: lessons learned from the standpoint of African American mothers. *Teach Educ Quart* 30(2):101–116
- Crimmins G, Nash G, Oprescu F, Alla K, Brock G, Hickson-Jamieson B, Noakes C (2016) Can a systematic assessment moderation process assure the quality and integrity of assessment practice while supporting the professional development of casual academics? *Assess Evaluat Higher Educ* 41(3):427–441
- Dahlin J, Kohn R, Schön TB (2016) Bayesian inference for mixed effects models with heterogeneity
- De la Cruz-Mesia R, Marshall G (2006) Non-linear random effects models with continuous time autoregressive errors: a Bayesian approach. *Stat Med* 25(9):1471–1484
- DeCarlo LT (2008) Studies of a latent-class signal-detection model for constructed-response scoring. *ETS Res Rep Ser*. <https://doi.org/10.1002/j.2333-8504.2008.tb02149.x>
- Dee TS (2005) A teacher like me: does race, ethnicity, or gender matter? *Am Econ Rev* 95(2):158–165
- DiMaggio P, Evans J, Bryson B (1996) Have American's social attitudes become more polarized? *Am J Sociol* 102(3):690–755
- Dorazio RM (2009) On selecting a prior for the precision parameter of Dirichlet process mixture models. *J Stat Plann Inference* 139(9):3384–3390
- Dressler WW, Balieiro MC, dos Santos JE (2015) Finding culture change in the second factor: stability and change in cultural consensus and residual agreement. *Field Methods* 27(1):22–38
- Esteban J-M, Ray D (1994) On the measurement of polarization. *Econometrica* 62(4):819–851
- Ferguson TS (1973) A Bayesian analysis of some nonparametric problems. *Ann Stat* 1(2):209–230
- Forchheimer D, Forchheimer R, Haviland D (2015) Improving image contrast and material discrimination with nonlinear response in bimodal atomic force microscopy. *Nat Commun* 6:6270
- Gelman A, Carlin J, Stern H, Dunson D, Vehtari A, Rubin D (2013) *Bayesian data analysis*. Chapman and Hall/CRC, Boca Raton
- Gill J, Casella G (2009) Nonparametric priors for ordinal Bayesian social science models: specification and estimation. *J Am Stat Assoc* 104(486):453–454
- Gisev N, Bell JS, Chen TF (2013) Interrater agreement and interrater reliability: key concepts, approaches, and applications. *Res Soc Adm Pharm* 9(3):330–338
- Gwet KL (2008) Computing inter-rater reliability and its variance in the presence of high agreement. *Br J Math Stat Psychol* 61(1):29–48
- Heinzl F, Tutz G (2013) Clustering in linear mixed models with approximate Dirichlet process mixtures using EM algorithm. *Stat Model* 13(1):41–67
- Heinzl F, Kneib T, Fahrmeir L (2012) Additive mixed models with Dirichlet process mixture and p-spline priors. *AStA Adv Stat Anal* 96:47–68
- Hsiao CK, Chen P-C, Kao W-H (2011) Bayesian random effects for interrater and test-retest reliability with nested clinical observations. *J Clin Epidemiol* 64(7):808–814
- Ishwaran H, James L (2001) Gibbs sampling methods for stick-breaking priors. *J Am Stat Assoc* 96:161–173
- James GM, Sugar CA (2003) Clustering for sparsely sampled functional data. *J Am Stat Assoc* 98(462):397–408
- Jang JH, Manatunga AK, Taylor AT, Long Q (2018) Overall indices for assessing agreement among multiple raters. *Stat Med* 37(28):4200–4215
- Kahrari F, Ferreira CS, Arellano-Valle RB (2019) Skew-normal-Cauchy linear mixed models. *Sankhya B Indian J Stat* 81(2):185–202

- Kim S, Tadesse MG, Vannucci M (2006) Variable selection in clustering via Dirichlet process mixture models. *Biometrika* 93(4):877–893
- Komárek A, Komárková L (2013) Clustering for multivariate continuous and discrete longitudinal data. *Ann Appl Stat* 7(1):177–200
- Komárek A, Hansen BE, Kuiper EMM, van Buuren HR, Lesaffre E (2010) Discriminant analysis using a multivariate linear mixed model with a normal mixture in the random effects distribution. *Stat Med* 29(30):3267–3283
- Koudenburg N, Kashima Y (2022) A polarized discourse: effects of opinion differentiation and structural differentiation on communication. *Personal Soc Psychol Bull* 48(7):1068–1086
- Koudenburg N, Kiers HAL, Kashima Y (2021) A new opinion polarization index developed by integrating expert judgments. *Front Psychol* 12:738258
- Kyung M, Gill J, Casella G (2011) New findings from terrorism data: Dirichlet process random-effects models for latent groups. *J R Stat Soc Ser C (Appl Stat)* 60(5):701–721
- Liljequist D, Elfving B, Skavberg Roaldsen K (2019) Intraclass correlation—a discussion and demonstration of basic features. *PLoS ONE* 14(7):1–35
- Lin TI, Lee JC (2008) Estimation and prediction in linear mixed models with skew-normal random effects for longitudinal data. *Stat Med* 27(9):1490–1507
- Makransky G, Terkildsen T, Mayer R (2019) Role of subjective and objective measures of cognitive processing during learning in explaining the spatial contiguity effect. *Learn Instr* 61:23–34
- Martinková P, Bartoš F, Brabec M (2023) Assessing inter-rater reliability with heterogeneous variance components models: flexible approach accounting for contextual variables. *J Educ Behav Stat* 48(3):349–383
- McCulloch CE, Neuhaus JM (2021) Improving predictions when interest focuses on extreme random effects. *J Am Stat Assoc* 118(541):504–513
- McHugh M (2012) Interrater reliability: the kappa statistic. *Biochemia medica : časopis Hrvatskoga društva medicinskih biokemičara /HDMB* 22:276–82
- Müller P, Quintana FA, Jara A, Hanson T (2015) Bayesian nonparametric data analysis, vol 1. Springer, Berlin
- Navarro DJ, Griffiths TL, Steyvers M, Lee MD (2006) Modeling individual differences using Dirichlet processes. *J Math Psychol*, 50(2):101–122. Special Issue on Model Selection: Theoretical Developments and Applications
- Nelson KP, Edwards D (2008) On population-based measures of agreement for binary classifications. *Canad J Stat* 36:411–426
- Nelson K, Edwards D (2015) Measures of agreement between many raters for ordinal classifications. *Stat Med* 34:3116–3132
- Oravecz Z, Vandekerckhove J, Batchelder WH (2014) Bayesian cultural consensus theory. *Field Methods* 26(3):207–222
- Paredes V (2014) A teacher like me or a student like me? Role model versus teacher bias effect. *Econ Educ Rev* 39:38–49
- Rigon T, Durante D (2021) Tractable Bayesian density regression via logit stick-breaking priors. *J Stat Plann Inference* 211:131–142
- Rodríguez A, Dunson D (2011) Nonparametric Bayesian models through probit stick-breaking processes. *Bayesian Anal* 6:145–178
- Schielzeth H, Dingemanse NJ, Nakagawa S, Westneat DF, Allogue H, Teplitsky C, Réale D, Dochtermann NA, Garamszegi LZ, Araya-Ajoy YG (2020) Robustness of linear mixed-effects models to violations of distributional assumptions. *Methods Ecol Evolut* 11(9):1141–1152
- Sethuraman J (1994) A constructive definition of Dirichlet priors. *Stat Sin* 4(2):639–650
- Shirazi MA (2019) For a greater good: Bias analysis in writing assessment. *SAGE Open* 9(1):2158244018822377
- Simpson D, Rue H, Riebler A, Martins TG, Sørbye SH (2017) Penalising model component complexity: a principled, practical approach to constructing priors. *Stat Sci* 32(1):1–28
- Stan Development Team (2022) RStan: the R interface to Stan. R package version 2.21.7
- Stefanucci M, Canale A (2021) Multiscale stick-breaking mixture models. *Stat Comput* 31:13
- Tang T, Ghorbani A, Squazzoni F, Chorus CG (2022) Together alone: a group-based polarization measurement. *Qual Quant* 56:3587–3619
- Tutz G, Oelker M-R (2017) Modelling clustered heterogeneity: fixed effects, random effects and mixtures. *Int Stat Rev* 85(2):204–227

- Ulker Y, Günsel B, Cemgil T (2010) Sequential Monte Carlo samplers for Dirichlet process mixtures. In: Teh YW, Titterton M (eds), Proceedings of the 13th international conference on artificial intelligence and statistics, volume 9 of proceedings of machine learning research, pp 876–883, Chia Laguna Resort, Sardinia, Italy. PMLR
- Uto M (2022) A Bayesian many-facet Rasch model with Markov modeling for rater severity drift. *Behav Res Methods* 55(7):3910–3928
- Verbeke G, Lesaffre E (1996) A linear mixed-effects model with heterogeneity in the random-effects population. *J Am Stat Assoc* 91(433):217–221
- Villarroel L, Marshall G, Barón AE (2009) Cluster analysis using multivariate mixed effects models. *Stat Med* 28(20):2552–2565
- Walker SG (2007) Sampling the Dirichlet mixture model with slices. *Commun Stat Simul Comput* 36(1):45–54
- Wang W-L, Lin T-I (2014) Multivariate t nonlinear mixed-effects models for multi-outcome longitudinal data with missing values. *Stat Med* 33(17):3029–3046
- Wirtz MA (2020) Interrater reliability. Springer, Cham, pp 2396–2399
- Zhang C, Mapes BE, Soden BJ (2003) Part a no. 594 q. *J R Meteorol Soc* 129:2847–2866
- Zhu Y, Fung AS-L, Yang L (2021) A methodologically improved study on raters' personality and rating severity in writing assessment. *SAGE Open* 11(2):21582440211009476
- Zupanc K, Štrumbelj E (2018) A Bayesian hierarchical latent trait model for estimating rater bias and reliability in large-scale performance assessment. *PLoS ONE* 13(4):1–16

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Authors and Affiliations

Giuseppe Mignemi¹  · Antonio Calcagni^{3,4} · Andrea Spoto¹ · Ioanna Manolopoulou²

✉ Giuseppe Mignemi
giuseppe.mignemi@phd.unipd.it
<https://github.com/MignemiG>

Antonio Calcagni
antonio.calcagni@unipd.it

Andrea Spoto
andrea.spoto@unipd.it

Ioanna Manolopoulou
i.manolopoulou@ucl.ac.uk

¹ University of Padova, Padua, Italy

² University College London, London, UK

³ DPSS, University of Padova, Padua, Italy

⁴ GNCS Research Group, National Institute of Advanced Mathematics (INdAM), Rome, Italy