

Unimodal Likelihood Models for Ordinal Data

Anonymous authors

Paper under double-blind review

Abstract

Ordinal regression (OR) is the classification of ordinal data, in which the underlying target variable is categorical and considered to have a natural ordinal relation for the explanatory variable. We in this study suppose the unimodality of conditional probability distributions as a natural ordinal relation of the ordinal data. Under this supposition, unimodal likelihood models are expected promising for improving the generalization performance in OR tasks. Demonstrating that previous unimodal likelihood models have a weak representation ability, we thus develop more representable unimodal models including the most representable one. Our took OR experiments show that the developed more representable unimodal models yielded better generalization performance for real-world ordinal data compared with previous unimodal models and popular statistical OR models having no unimodality guarantee.

1 Introduction

Ordinal regression (OR, also called ordinal classification) is the classification of ordinal data, in which the underlying target variable is categorical and considered to have a natural ordinal relation for the explanatory variable; see Section 2 for a detailed formulation. Typical examples of the target label set of ordinal data are grouped continuous variables like age groups {‘0 to 9 years old’, ‘10 to 19 years old’, ..., ‘90 to 99 years old’, ‘over 100 years old’} or assessed ordered categorical variables like human’s rating {‘strongly agree’, ‘agree’, ‘neutral’, ‘disagree’, ‘strongly disagree’} (Anderson, 1984), and various practical tasks have been tackled within the OR framework: for example, face-age estimation (Niu et al., 2016; Cao et al., 2019; Anonymous, 2022), information retrieval (Liu, 2011), credit or movie rating (Kim & Ahn, 2012; Yu et al., 2006), and questionnaire survey in social research (Chen et al., 1995; Bürkner & Vuorre, 2019).

Consider an example of the questionnaire survey about support for a certain idea that requires subjects to respond from {‘strongly agree’, ‘agree’, ...}. Here, it would seem possible that subjects, who have features specific to those who typically respond ‘agree’, respond ‘neutral’, but unlikely that they respond ‘disagree’. Such a phenomenon can be rephrased as the unimodality of the conditional probability distribution (CPD) of the underlying events. The hypothesis “many statisticians or practitioners often judge that the data have a natural ordinal relation and decide to treat them within the OR framework, with unconsciously expecting their unimodality” may be convincing, as we confirmed experimentally that many ordinal data, treated in previous OR studies that do not consider the unimodality, tend to have a unimodal CPD (see Section 6.2).

Commonly, the generalization performance of a classifier (OR method) based on the statistical modeling depends on the underlying data distribution and the representation ability of the model: A model that is too strongly restricted or unnecessarily flexible to represent the data may degrade its generalization performance. Assuming the unimodality of the CPD as a natural ordinal relation of ordinal data, unimodal likelihood models are considered promising for improving the generalization performance in OR tasks.

Existing studies (da Costa et al., 2008; Iannario & Piccolo, 2011; Beckham & Pal, 2017) were inspired by the shape of the probability mass function (PMF) of elementary categorical probability distributions such as binomial, Poisson, and uniform distributions, and developed unimodal likelihood models. In this paper, we introduce several notions that characterize the shape of PMFs in Section 2.1, and, on the basis of them, show that their models have a weak representation ability in Section 3. We thus propose more representable unimodal models in Sections 4 and 5; In particular, a model described in Section 4.2 is the most strongly representable among the class of unimodal likelihood models.

We took experimental comparisons between 2 previous unimodal models, 2 popular statistical OR models without the unimodality guarantee, and 8 proposed unimodal models. As we will show the experimental results and considerations in Section 6.3 and the Appendix, we confirmed that our proposed more representable unimodal models were effective to improve the generalization performances for the conditional probability estimation and OR tasks for many data that have been treated in previous OR studies as ordinal data.

2 Preliminaries

2.1 Ordinal Regression Tasks and Ordinal Data

The OR task is a classification task. Denoting a categorical target variable underlying the data as $Y \in [K] := \{1, \dots, K\}$ and an explanatory variable underlying the data as $\mathbf{X} \in \mathbb{R}^d$, we formulate the OR task as searching for a classifier $f : \mathbb{R}^d \rightarrow [K]$ that is good in the sense that the task risk $\mathbb{E}[\ell(f(\mathbf{X}), Y)]$ becomes small for a specified task loss $\ell : [K]^2 \rightarrow [0, +\infty)$, where the expectation $\mathbb{E}[\cdot]$ is taken for all included random variables (here \mathbf{X} and Y). Popular task losses in OR tasks include not only the zero-one loss $\ell_{zo}(j, k) := \mathbb{1}\{j \neq k\}$, where $\mathbb{1}\{c\}$ values 1 if a condition c is true and 0 otherwise, but also V-shaped losses reflecting one’s preference of smaller prediction errors over larger ones such as the absolute loss $\ell_{abs}(j, k) := |j - k|$ and the squared loss $\ell_{sq}(j, k) := (j - k)^2$.

In the OR framework, it is supposed that an underlying categorical target variable Y of the data is equipped with an ordinal relation naturally interpretable in the relationship with its underlying explanatory variable \mathbf{X} , as examples described in the head of Section 1. We here assume that the target labels are encoded to $1, \dots, K$ in an order-preserving manner, like from ‘strongly agree’, ‘agree’, ‘neutral’, ‘disagree’, ‘strongly disagree’ to $1, \dots, 5$ or $5, \dots, 1$. The OR framework considers the classification of such ordinal data. Note that, like most previous OR studies have discussed the OR without formal common understanding of what constitutes ordinal data, it is difficult to define ordinal data any more rigorously, and we in this paper refer to the data discussed in previous OR studies as ordinal data.

As we declared in Section 1, this study basically assumes, as a natural ordinal relation of ordinal data, the unimodality in the theoretical discussion or the almost-unimodality to real-world ordinal data, precisely defined in the following:

Definition 1. For a vector $\mathbf{p} = (p_k)_{k \in [K]} \in \mathbb{R}^K$, we define $M(\mathbf{p}) := \min(\arg \max_k (p_k)_{k \in [K]})^1$, and say that \mathbf{p} is unimodal if it satisfies

$$p_1 \leq \dots \leq p_{M(\mathbf{p})} \text{ and } p_{M(\mathbf{p})} \geq \dots \geq p_K. \quad (1)$$

Also, we call $M(\mathbf{p})$ the mode of \mathbf{p} if \mathbf{p} is a PMF satisfying $\mathbf{p} \in \Delta_{K-1}$, where Δ_{K-1} is the $(K-1)$ -dimensional probability simplex $\{(p_k)_{k \in [K]} \mid \sum_{k=1}^K p_k = 1, p_k \in [0, 1] \text{ for } k = 1, \dots, K\}$. Moreover, if the CPD $(\Pr(Y = y | \mathbf{X} = \mathbf{x}))_{y \in [K]}$ is unimodal at any \mathbf{x} in whole the domain \mathbb{R}^d or in its sub-domain $\mathcal{X} \subseteq \mathbb{R}^d$ with a large probability $\Pr(\mathbf{X} \in \mathcal{X})$, we say that the data is unimodal or almost-unimodal.

This paper presents a qualitative discussion on the representation ability of various statistical OR models: It aims to clarify whether each model is suitable for representing data that follow a distribution with various structural properties that we consider to be expectable in unimodal ordinal data. As a preparation for that discussion, we further introduce below the notions of decay rates, scale, homoscedasticity, heteroscedasticity, and skewness that are related to structural properties of the data distribution:

Definition 2. For a PMF $\mathbf{p} = (p_k)_{k \in [K]} \in \Delta_{K-1}$ having a mode $m \in [K]$, we call $\frac{p_k}{p_{k+1}}$ for $k = 1, \dots, m-1$ (if $m \neq 1$) and $\frac{p_k}{p_{k-1}}$ for $k = m+1, \dots, K$ (if $m \neq K$) the decay rates (DRs) of \mathbf{p} .

Definition 3. For a PMF $\mathbf{p} = (p_k)_{k \in [K]} \in \Delta_{K-1}$ having a mode $m \in [K]$, we call $\sum_{k=1}^{m-1} p_k + \sum_{k=m+1}^K p_k = 1 - p_m$ the scale of \mathbf{p} , where we define the summation such that $\sum_{k=i}^j f(k)$ is zero irrelevant to the function f as far as $i > j$. If the scale of the CPD $\mathbf{p}(\mathbf{x}) = (\Pr(Y = y | \mathbf{X} = \mathbf{x}))_{y \in [K]}$ is similar or dissimilar over whole the domain \mathbb{R}^d of the explanatory variable \mathbf{X} , we say that the data is homoscedastic or heteroscedastic. In particular, we refer to heteroscedastic data as mode-wise heteroscedastic or overall heteroscedastic data, if

¹arg min and arg max may return a set, and min is applied to convert it to a point statistic and simplify the discussion.




Figure 1: Instances of the unimodal 10-dimensional PMFs $(p_k)_{k \in [10]}$, whose mode is 5, in black solid line, and their DRs $\{\frac{p_k}{p_{k+1}}\}_{k=1}^4$ and $\{\frac{p_k}{p_{k-1}}\}_{k=6}^{10}$ in red dotted line.

the scale of their $\mathbf{p}(\mathbf{x})$ is similar or can be dissimilar in a domain where the conditional mode $M(\mathbf{p}(\mathbf{x}))$ is the same.

Definition 4. For a PMF $\mathbf{p} = (p_k)_{k \in [K]} \in \Delta_{K-1}$ having a mode $m \in [K]$, we call $\sum_{k=1}^{m-1} p_k - \sum_{k=m+1}^K p_k$ the skewness of \mathbf{p} , and we say that \mathbf{p} is skew or less skew when the absolute value of its skewness (absolute skewness) is large or small. If the absolute skewness of the CPD $\mathbf{p}(\mathbf{x}) = (\Pr(Y = y | \mathbf{X} = \mathbf{x}))_{y \in [K]}$ is large or small over whole the domain \mathbb{R}^d of the explanatory variable \mathbf{X} , we say that the data is skew or less skew. In particular, we refer to skew data as mode-wise skew or overall skew data, if the skewness of their $\mathbf{p}(\mathbf{x})$ is similar or can be dissimilar in a domain where the conditional mode $M(\mathbf{p}(\mathbf{x}))$ is the same.

For a unimodal PMF, its DRs are smaller than or equal to 1. The scale and skewness of PMFs are typically treated as qualitative notions, but we measure them numerically for the sake of clarity (see $\sum_{k=1}^{m-1} p_k + \sum_{k=m+1}^K p_k$ and $\sum_{k=1}^{m-1} p_k - \sum_{k=m+1}^K p_k$ in Definitions 3 and 4). The numerical measure in each definition is intended just to further clarify our qualitative discussion, not that we see any particular importance in that choice of the measure (other measures can be used instead). The readers should recognize our treatment of them with reference to the description of Definitions 2, 3, and 4 and Figure 1. The figure displays 3 instances of the unimodal PMFs and their DRs: we say that the right PMF has a larger scale than the left PMF and that the center PMF is less skew and the right PMF is skew.

Additionally, we introduce the degree of freedom:

Definition 5. The degree of freedom (DF) of a vector-valued function \mathbf{g} defined on \mathbb{R}^d refers to the maximum value over all $\mathbf{x} \in \mathbb{R}^d$ of the dimension at the point $\mathbf{g}(\mathbf{x})$ of the manifold $\{\mathbf{g}(\mathbf{x}) : \mathbf{x} \in \mathbb{R}^d\}$.

General CPDs $(\Pr(Y = y | \mathbf{X} = \mathbf{x}))_{y \in [K]}$ including those underlying unimodal data have up to $(K - 1)$ -DF (or called full-DF) (it is not K since $\sum_{y=1}^K \Pr(Y = y | \mathbf{X} = \mathbf{x}) = 1$). We use the DF as one indicator of the representation ability of statistical likelihood models defined below; we consider that the larger its DF, the richer the representation ability of the statistical model tends.

2.2 Likelihood Models and Ordinal Regression Methods

Suppose that one has a set of observations (ordinal data) $\{(\mathbf{x}_i, y_i)\}_{i=1}^n$, each of which is drawn independently from an identical distribution of (\mathbf{X}, Y) . Every OR method covered in this paper assumes, as a model of the conditional probability $\Pr(Y = y | \mathbf{X} = \mathbf{x})$, a certain likelihood model $\hat{\Pr}(Y = y | \mathbf{X} = \mathbf{x}) = P(y; \mathbf{g}(\mathbf{x}))$ with a fixed part P (say, the softmax function in multinomial logistic regression model) and learnable part $\mathbf{g} \in \mathcal{G}$ (say, a certain neural network model), where we call P the link function and \mathbf{g} the learner model in a model class \mathcal{G} . Note that the distinction between P and \mathbf{g} is to aid in understanding relation between multiple models, and this paper does not emphasize a strict mathematical distinction (for example, (P, \mathcal{G}) and $(2P, \{\mathbf{g}/2 \mid \mathbf{g} \in \mathcal{G}\})$ yield an equivalent likelihood model).

The key notion in the discussion of this paper is the representation ability of the likelihood model. Strictly speaking, the representation ability we discuss is defined by the following relationship:

Definition 6. We say that likelihood models based on (P_1, \mathcal{G}_1) has a stronger representation ability (or is more representable) than those based on (P_2, \mathcal{G}_2) , if there exists $\mathbf{g}_1 \in \mathcal{G}_1$ such that $P_1(y; \mathbf{g}_1(\mathbf{x})) = P_2(y; \mathbf{g}_2(\mathbf{x}))$ for any $\mathbf{x} \in \mathbb{R}^d$ and $y \in [K]$, for any $\mathbf{g}_2 \in \mathcal{G}_2$.

The representation abilities of two different likelihood models may not be strictly comparable according to Definition 6, but even in such cases we in this paper discuss the relationship between the representation abilities from a qualitative understanding of those likelihood models as much as possible.

In this paper, we set up an OR method as learning a model \mathbf{g} from the class \mathcal{G} through the maximum likelihood estimation $\max_{\mathbf{g} \in \mathcal{G}} \frac{1}{n} \sum_{i=1}^n \log P(y_i; \mathbf{g}(\mathbf{x}_i))$ (which we call a conditional probability estimation task), and then constructing a classifier under the task with the task loss ℓ as $f(\mathbf{x}) = f_\ell(\mathbf{x}; (P(y; \hat{\mathbf{g}}(\mathbf{x})))_{y \in [K]})$ with $f_\ell(\mathbf{x}; (p_k)_{k \in [K]}) := \min(\arg \min_j (\sum_{k=1}^K p_k \ell(j, k)))_{j \in [K]}$ and an obtained model $\hat{\mathbf{g}}$.

Therefore, the only difference of our-treated OR methods appears in the link function P and model class \mathcal{G} of the likelihood model. Under these settings, generic principles suggest that a compact likelihood model that can adequately represent the data is promising for an OR method with a better generalization performance. On the other hand, we believe that our subsequent discussions hold as well even if changing the loss function used in the parameter fitting procedure and the decision function: As options for the parameter fitting other than the maximum likelihood estimation, we can apply robust alternatives (Bianco & Yohai, 1996; Croux et al., 2013) to every our-treated likelihood model. Anonymous (2022) has developed a decision function that allows misspecification of the likelihood model. As another decision function, it would also be possible to assume a certain prior on the model \mathbf{g} , take the Bayesian discussion further, and use a decision function based on the posterior predictive probability distribution.

3 Existing Unimodal Likelihood Models

3.1 Binomial Models

The (shifted) binomial distribution $(P_b(k; p))_{k \in [K]}$ is unimodal at any $p \in [0, 1]$, where

$$P_b(k; p) := \binom{K-1}{k-1} p^{k-1} (1-p)^{K-k} \text{ for } k \in [K], p \in [0, 1], \quad (2)$$

and where $\binom{k}{l} := \frac{k!}{l!(k-l)!}$ is the binomial coefficient. Inspired by the shape of the PMF of the binomial distribution, da Costa et al. (2008) considered unimodal likelihood models based on the link function

$$P_{\text{bin}}(y; u) := \binom{K-1}{y-1} \left(\frac{1}{1+e^{-u}} \right)^{y-1} \left(\frac{1}{1+e^u} \right)^{K-y} \text{ for } y \in [K], u \in \mathbb{R} \quad (3)$$

and an \mathbb{R} -valued learner model $\mathbf{g}(\mathbf{x})$ (applied to \mathbf{u} in (3)). Thereafter, Beckham & Pal (2017) introduced the scaling factor: in other words, they proposed the link function

$$P_{\text{bin}}(y; u, s) := \frac{e^{\log(P_{\text{bin}}(y; u))/s}}{\sum_{k=1}^K e^{\log(P_{\text{bin}}(k; u))/s}} \text{ for } y \in [K], u \in \mathbb{R}, s \in \mathbb{R}_+, \quad (4)$$

where $\mathbb{R}_+ := \{v \mid 0 < v < \infty\}$. They applied a \mathbf{x} -dependent model $\mathbf{g}(\mathbf{x})$ to \mathbf{u} and a \mathbf{x} -independent positive parameter to s and learned them; we call models consisting of the link function P_{bin} and learner models in $\mathcal{G} = \{(g(\cdot), s) \mid g : \mathbb{R}^d \rightarrow \mathbb{R}_+, s \in \mathbb{R}_+\}$ the Binomial (BIN) models.


The BIN models $(P_{\text{bin}}(y; \mathbf{g}(\mathbf{x}), s))_{y \in [K]}$ is parametrically constrained with the 1-DF. One would know that the mode of the binomial distribution is

$$M((P_b(k; p))_{k \in [K]}) = \min(\{\lceil Kp \rceil, \lfloor Kp \rfloor + 1\} \cap [K]), \quad (5)$$

where $\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ are the ceil and floor functions, and hence $\max\{1, Kp\} \leq M((P_b(k; p))_{k \in [K]}) \leq \min\{Kp+1, K\}$. These known properties and simple calculations can show the mode, unimodality, and DRs of the BIN model:

Theorem 1. *It holds that*

$$(i) \quad \frac{P_{\text{bin}}(y+1; u, s)}{P_{\text{bin}}(y; u, s)} = \left(\frac{(K-y)p}{y(1-p)} \right)^{1/s} \text{ with } p = \frac{1}{1+e^{-u}} \text{ for all } y \in [K-1], u \in \mathbb{R}, s \in \mathbb{R}_+.$$

Figure 2: Instances of the BIN models (3) with $K = 10$.²

Then, for any $u \in \mathbb{R}$ and $s \in \mathbb{R}_+$, it holds that, for $m = M((P_{\text{bin}}(y; u, s))_{y \in [K]})$ that is (5) with $p = \frac{1}{1+e^{-u}}$,

(ii) $P_{\text{bin}}(1; u, s) \leq \dots \leq P_{\text{bin}}(m; u, s)$ if $m \neq 1$, and $P_{\text{bin}}(m; u, s) \geq \dots \geq P_{\text{bin}}(K; u, s)$ if $m \neq K$,

(iii) $\frac{P_{\text{bin}}(1; u, s)}{P_{\text{bin}}(2; u, s)} \leq \dots \leq \frac{P_{\text{bin}}(m-1; u, s)}{P_{\text{bin}}(m; u, s)} (\leq 1)$ if $m \neq 1$, and $(1 \geq) \frac{P_{\text{bin}}(m+1; u, s)}{P_{\text{bin}}(m; u, s)} \geq \dots \geq \frac{P_{\text{bin}}(K; u, s)}{P_{\text{bin}}(K-1; u, s)}$ if $m \neq K$.

We call the constraint on the sequence of DRs like Theorem 1, (iii) as the DRs' unimodality (constraint), considering that $(\frac{p_1}{p_2}, \dots, \frac{p_{m-1}}{p_m}, \frac{p_{m+1}}{p_m}, \dots, \frac{p_K}{p_{K-1}})$ is unimodal if a PMF $(p_k)_{k \in [K]}$ having a mode m satisfies that constraint. One can find that, owing to the DRs' unimodality constraint, the BIN models cannot exactly represent unimodal MPFs in Figure 1. In addition to the DRs' unimodality constraint, the BIN model $(P_{\text{bin}}(y; g(\mathbf{x}), s))_{y \in [K]}$ is always less skew especially when its mode is close to the labels' intermediate value $\frac{1+K}{2}$ (recall that the mode (5) and median, $\lceil (K-1)p \rceil + 1$ or $\lfloor (K-1)p \rfloor + 1$, of the binomial distribution $(P_b(k; p))_{k \in [K]}$ are close (Kaas & Buhrman, 1980)), and tends homoscedastic (see Figure 2; there, the scale $1 - \max_y (P_{\text{bin}}(y; u, s))_{y \in [10]}$ for each s is similar for many u except at both ends). We will relax in Section 5.2 the restriction of representability regarding the scale of BIN models, but they cannot avoid the restriction regarding the DRs and skewness.

3.2 Poisson Models

The (shifted) Poisson distribution $(P_p(k; \lambda))_{k \in \mathbb{N}}$ is unimodal at any $\lambda > 0$, where

$$P_p(k; \lambda) := \frac{\lambda^{k-1} e^{-\lambda}}{(k-1)!} \text{ for } k \in \mathbb{N}, \lambda > \mathbb{R}_+. \quad (6)$$

da Costa et al. (2008) truncated (6) within $[K]$ and standardize it to develop likelihood models: $\hat{P}r(Y = y | \mathbf{X} = \mathbf{x}) = P_p(y; g(\mathbf{x})) / \sum_{k=1}^K P_p(k; g(\mathbf{x}))$ with an \mathbb{R}_+ -valued learner model $g(\mathbf{x})$. As in the case for BIN models, Beckham & Pal (2017) developed a scaled link function


$$P_{\text{poi}}(y; u, s) := \frac{e^{v_y/s}}{\sum_{k=1}^K e^{v_k/s}} \text{ for } y \in [K], u, s \in \mathbb{R}_+, \text{ with } v_k = (k-1) \log(u) - \log((k-1)!) \text{ for } k \in [K], \quad (7)$$

which is a generalization of $P_{\text{poi}}(y; u, 1)$ of (da Costa et al., 2008).³ They applied a \mathbf{x} -dependent \mathbb{R}_+ -valued learner model $g(\mathbf{x})$ to u and a \mathbf{x} -independent positive parameter to s (i.e., $\mathcal{G} = \{(g(\cdot), s) \mid g : \mathbb{R}^d \rightarrow \mathbb{R}_+, s \in \mathbb{R}_+\}$); We call these models the Poisson (POI) Models.

The POI model $(P_{\text{poi}}(y; g(\mathbf{x}), s))_{y \in [K]}$ is also parametrically constrained with the 1-DF. The POI model $(P_{\text{poi}}(y; g(\mathbf{x}), s))_{y \in [K]}$ has a similar DRs' unimodality constraint to the BIN model, and always has mode-wise heteroscedasticity especially when s is small and K is large: its scale tends small and large respectively when $M((P_{\text{poi}}(y; g(\mathbf{x}), s))_{y \in [K]})$ is close to 1, K , or smaller side of $\{2, \dots, K-1\}$ and larger side; See Figure 3 with $s = 0.2$. They, therefore, cannot represent homoscedastic or overall heteroscedastic data. Furthermore, we find that the POI models are special instances of PO-ACL (and PO-ORD-ACL) models discussed later and have a weaker representation ability; Refer to also Sections 4.1, 5.1, and 5.2 and Theorem 3.

²The display style in Figure 1 is like showing a PMF created by cutting the curves in Figure 2 out at each u .

³In (Beckham & Pal, 2017), v_k in (7) is defined with minus u , but which is meaningless after the softmax transformation.

Figure 3: Instances of the POI models (7) with $K = 10$.

4 Novel Unimodal Likelihood Models

4.1 Ordered Adjacent Categories Logit Models

The BIN and POI models are unimodal, but have just 1-DF and may be too weakly representable for some data. We thus developed more strongly representable unimodal likelihood models that have up to full-DF, i.e., $(K - 1)$ -DF. This section describes ordered adjacent categories logit (ORD-ACL) models developed by modifying existing ACL models that do not have the unimodality guarantee.

The (naïve) ACL models (see (Simon, 1974; Andrich, 1978; Goodman, 1979; Masters, 1982), and (Agresti, 2010, Section 4.1)) are designed to model the relationship between the underlying conditional probabilities of the adjacent categories, $\frac{\Pr(Y=y|\mathbf{X}=\mathbf{x})}{\Pr(Y \in \{y, y+1\}|\mathbf{X}=\mathbf{x})}$, as

$$\frac{\hat{\Pr}(Y = y|\mathbf{X} = \mathbf{x})}{\hat{\Pr}(Y \in \{y, y+1\}|\mathbf{X} = \mathbf{x})} = \frac{1}{1 + e^{-g_y(\mathbf{x})}} \text{ for } y \in [K-1] \quad (8)$$

with an $\mathbb{R}^{(K-1)}$ -valued learner model \mathbf{g} . Considering the normalization condition $\sum_{y=1}^K \hat{\Pr}(Y = y|\mathbf{X} = \mathbf{x}) = 1$, it can be found that ACL models depend on the ACL link function

$$P_{\text{acl}}(y; \mathbf{u}) := \frac{\prod_{l=1}^{y-1} e^{-u_l}}{\sum_{k=1}^K \prod_{l=1}^{k-1} e^{-u_l}} \text{ for } y \in [K], \mathbf{u} \in \mathbb{R}^{K-1}, \quad (9)$$

together with the model class $\mathcal{G} = \{\mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^{K-1}\}$, where we define the product such that $\prod_{k=i}^j f(k)$ is one irrelevant to the function f as far as $i > j$.

We have analyzed properties of the ACL link function and then found a simple condition so that the ACL models turn to get the unimodality guarantee:

Theorem 2. *It holds that*

$$(i) \quad \frac{P_{\text{acl}}(y+1; \mathbf{u})}{P_{\text{acl}}(y; \mathbf{u})} = e^{-u_y} \text{ for all } y \in [K-1], \mathbf{u} \in \mathbb{R}^{K-1}.$$

Assume that $u_0(= -\infty) \leq \dots \leq u_{m-1} \leq 0 \leq u_m \leq \dots \leq u_K(= +\infty)$ for some $m \in [K]$. Then, it holds that

$$(ii) \quad P_{\text{acl}}(1; \mathbf{u}) \leq \dots \leq P_{\text{acl}}(m; \mathbf{u}) \text{ if } m \neq 1, \text{ and } P_{\text{acl}}(m; \mathbf{u}) \geq \dots \geq P_{\text{acl}}(K; \mathbf{u}) \text{ if } m \neq K,$$

$$(iii) \quad \frac{P_{\text{acl}}(1; \mathbf{u})}{P_{\text{acl}}(2; \mathbf{u})} \leq \dots \leq \frac{P_{\text{acl}}(m-1; \mathbf{u})}{P_{\text{acl}}(m; \mathbf{u})} (\leq 1) \text{ if } m \neq 1, \text{ and } (1 \geq) \frac{P_{\text{acl}}(m+1; \mathbf{u})}{P_{\text{acl}}(m; \mathbf{u})} \geq \dots \geq \frac{P_{\text{acl}}(K; \mathbf{u})}{P_{\text{acl}}(K-1; \mathbf{u})} \text{ if } m \neq K.$$

On the basis of the unimodality guarantee stated in Theorem 2, (ii), we propose ORD-ACL models $(P_{\text{acl}}(y; \hat{\mathbf{g}}(\mathbf{x})))_{y \in [K]}$ applying the ACL link function (9) and ordered learner model $\hat{\mathbf{g}}(\mathbf{x})$. Here, the ordered learner model $\hat{\mathbf{g}}$ can be implemented through the re-parameterization as

$$\hat{g}_k(\mathbf{x}) = \begin{cases} g_1(\mathbf{x}), & \text{for } k = 1, \\ \hat{g}_{k-1}(\mathbf{x}) + \rho(g_k(\mathbf{x})), & \text{for } k = 2, \dots, K-1, \end{cases} \quad (10)$$

with another model $\mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^{K-1}$ and a fixed non-negative function $\rho : \mathbb{R} \rightarrow [0, \infty)$ such as $\rho_{\text{exp}}(u) := e^u$, $\rho_{\text{sq}}(u) := u^2$ (we denote the procedure (10) as $\check{\mathbf{g}} = \rho[\mathbf{g}]$; i.e., the model class $\mathcal{G} = \{\rho[\mathbf{g}] \mid \mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^{K-1}\}$).

It is a good aspect in the modeling of the ordinal data assumed to be unimodal that the likelihood models are guaranteed to be unimodal. In contrast, the ACL models also incidentally impose a constraint on the DRs (see Theorem 2, (i) and (iii)): The conditional model at \mathbf{x} (let m) satisfies $0 \leq \dot{g}_m(\mathbf{x})$ if $m = 1$, $\dot{g}_{m-1}(\mathbf{x}) < 0$ if $m = K$ or $\dot{g}_{m-1}(\mathbf{x}) < 0 \leq \dot{g}_m(\mathbf{x})$ otherwise, $\frac{P_{\text{acl}}(y; \check{\mathbf{g}}(\mathbf{x}))}{P_{\text{acl}}(y+1; \check{\mathbf{g}}(\mathbf{x}))} = e^{\dot{g}_y(\mathbf{x})}$ is non-decreasing in y for $y = 1, \dots, m-1$, and $\frac{P_{\text{acl}}(y; \check{\mathbf{g}}(\mathbf{x}))}{P_{\text{acl}}(y-1; \check{\mathbf{g}}(\mathbf{x}))} = e^{-\dot{g}_{y-1}(\mathbf{x})}$ is non-increasing in y for $y = m+1, \dots, K$. Nevertheless, it would be important that although a similar DRs' unimodality constraint was identified in the BIN and POI models, the ORD-ACL models are so representable that they fully encompass the BIN and POI models.

4.2 V-Shaped Stereotype Logit Models

We next describe V-shaped stereotype logit (VS-SL) models, which are developed by modifying SL models, unimodal, full-DF, and more flexible than ORD-ACL models.

The SL models (refer to (Anderson, 1984) and (Agresti, 2010, Section 4.3)) attempt to model the conditional probabilities of multiple categories paired with a certain fixed (stereotype) category, $\frac{\Pr(Y=1|\mathbf{X}=\mathbf{x})}{\Pr(Y \in \{1, y\}|\mathbf{X}=\mathbf{x})}$, as

$$\frac{\hat{\Pr}(Y = 1|\mathbf{X} = \mathbf{x})}{\hat{\Pr}(Y \in \{1, y\}|\mathbf{X} = \mathbf{x})} = \frac{1}{1 + e^{-g_y(\mathbf{x})}} \text{ for } y \in [K] \quad (11)$$

with an \mathbb{R}^K -valued learner model \mathbf{g} such that $g_1(\mathbf{x}) = 0$ for any $\mathbf{x} \in \mathbb{R}^d$ (i.e., the model class $\mathcal{G} = \{\mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^K \mid g_1(\cdot) = 0\}$ for the link function P_{sl}). Considering the normalization condition, we introduce the SL link function as

$$P_{\text{sl}}(y; \mathbf{u}) := \frac{e^{-u_y}}{\sum_{k=1}^K e^{-u_k}} \text{ for } y \in [K], \mathbf{u} \in \mathbb{R}^K. \quad (12)$$


With the SL link function, our proposed VS-SL models take the format $(P_{\text{sl}}(y; \check{\mathbf{g}}(\mathbf{x})))_{y \in [K]}$ with a V-shaped learner model $\check{\mathbf{g}} = \tau[\mathbf{g}]$ (described below) for $\mathbf{g} = \rho[\mathbf{g}]$ and another \mathbb{R}^K -valued learner model \mathbf{g} (i.e., the model class $\mathcal{G} = \{\tau[\rho[\mathbf{g}]] \mid \mathbf{g} : \mathbb{R}^d \rightarrow \mathbb{R}^K\}$ for the link function P_{sl}). First, ρ transforms an arbitrary model \mathbf{g} to an ordered model $\check{\mathbf{g}}$ (so $\dot{g}_1(\mathbf{x}) \leq \dots \leq \dot{g}_K(\mathbf{x})$ for any $\mathbf{x} \in \mathbb{R}^d$). Next, τ transforms an ordered model $\check{\mathbf{g}}$ to a V-shaped model $\tau[\check{\mathbf{g}}]$, where the notation $\tau[\mathbf{u}]$ for an \mathbb{R}^K -valued object \mathbf{u} implies the element-wise application $(\tau \circ u_k)_{k \in [K]}$. The function $\tau(u)$ for $u \in \mathbb{R}$ is supposed to be non-negative, non-increasing in $u < 0$, and non-decreasing in $u > 0$ such as $\tau_{\text{abs}}(u) := |u|$ and $\tau_{\text{sq}}(u) := u^2$. It holds that $\check{g}_1(\mathbf{x}) \geq \dots \geq \check{g}_m(\mathbf{x})$ and $\check{g}_m(\mathbf{x}) \leq \dots \leq \check{g}_K(\mathbf{x})$ with $m = \min(\arg \min_k (|\dot{g}_k(\mathbf{x})|)_{k \in [K]})$ if $\tau = \tau_{\text{abs}}, \tau_{\text{sq}}$. Finally, the SL link function transforms $\check{\mathbf{g}}$ to a unimodal likelihood model that is normalized and remains non-negative.

Most importantly, the VS-SL models $(P_{\text{sl}}(y; \check{\mathbf{g}}(\mathbf{x})))_{y \in [K]}$ is ensured to be unimodal and further can represent arbitrary unimodal data if ignoring the restriction owing to the structure of an implemented learner model \mathbf{g} in $\tau[\rho[\mathbf{g}]]$ (or equivalently, if using an extremely flexible \mathbf{g}). Therefore, the VS-SL models can be viewed the most representable ones among the class of the unimodal likelihood models.

5 Variant Models

5.1 Proportional-Odds Models

The ACL models $(P_{\text{acl}}(y; \mathbf{g}(\mathbf{x})))_{y \in [K]}$, ORD-ACL models $(P_{\text{acl}}(y; \rho[\mathbf{g}](\mathbf{x})))_{y \in [K]}$, SL models $(P_{\text{sl}}(y; \mathbf{g}(\mathbf{x})))_{y \in [K]}$, and VS-SL models $(P_{\text{sl}}(y; \tau[\rho[\mathbf{g}]](\mathbf{x})))_{y \in [K]}$ have full-DF and may be too flexible and difficult to learn for some data especially when the sample size is not large enough. In the statistical research (e.g., see (McCullagh, 1980)), the proportional-odds (PO) constraint that constrains the DF of the learner model to 1 is often applied to avoid such a trouble. We call $(P_{\text{acl}}(y; \mathbf{b} - \mathbf{g}(\mathbf{x}) \cdot \mathbf{1}))_{y \in [K]}$ the PO-ACL model, $(P_{\text{acl}}(y; \rho[\mathbf{b}] - \mathbf{g}(\mathbf{x}) \cdot \mathbf{1}))_{y \in [K]}$ the PO-ORD-ACL model, and $(P_{\text{sl}}(y; \tau[\rho[\mathbf{b}] - \mathbf{g}(\mathbf{x}) \cdot \mathbf{1}]))_{y \in [K]}$ the PO-VS-SL model, where \mathbf{b} and $\mathbf{1}$ are learnable

Figure 4: Instances of the PO-ACL models with $K = 10$.

parameter and all-1 vector with appropriate dimension ($(K - 1)$ for ACL and ORD-ACL models, or K for VS-SL models) and g is an \mathbb{R} -valued learner model.⁴

The PO-ORD-ACL and PO-VS-SL models are novel. The significance of these models we expect lies in the theoretical guarantee of the unimodality, not in the improvement of probability prediction and classification performances (e.g., in comparison between PO-ORD-ACL and ORD-ACL models): For example, the PO-ACL model $(P_{\text{acl}}(y; \mathbf{b} - g(\mathbf{x}) \cdot \mathbf{1}))_{y \in [K]}$ may result in ordered parameter \mathbf{b} depending on the overall tendency of the training data even if not explicitly imposing the ordering constraint during the parameter fitting (but not always).

Additionally, recalling the POI model reviewed in Section 3.2, it would be clear that the POI model is a special instance of the PO-ACL models:

Theorem 3. *For any $y \in [K]$, $u \in \mathbb{R}$, $s \in \mathbb{R}_+$, it holds that $P_{\text{poi}}(y; u, s) = P_{\text{acl}}(y; \mathbf{b} - v \cdot \mathbf{1})$ with $\mathbf{b} = (b_k)_{k \in [K-1]}$ satisfying $b_k = \log(k)/s$ for $k = 1, \dots, K - 1$ and $v = \log(u)/s$.*

The PO-ACL model can further represent homoscedastic data and more various mode-wise heteroscedastic data because it can adjust its parameter \mathbf{b} according to the model fit to the data, compared with the POI model; Compare Figure 3 for the POI models and Figure 4 for the PO-ACL models.

5.2 Overall Heteroscedastic Models

The BIN, POI, and PO models have just 1-DF. The strong constraint of these models can be interpreted as incidentally assuming the homoscedasticity or mode-wise heteroscedasticity. We in this section describes their overall homoscedastic (OH) extension, OH models.

McCullagh (1980) studied an OH extension of a certain PO model. According to his idea, we propose the OH-ACL model $(P_{\text{acl}}(y; \{\mathbf{b} - g(\mathbf{x}) \cdot \mathbf{1}\}/s(\mathbf{x})))_{y \in [K]}$, OH-ORD-ACL model $(P_{\text{acl}}(y; \{\rho[\mathbf{b}] - g(\mathbf{x}) \cdot \mathbf{1}\}/s(\mathbf{x})))_{y \in [K]}$, and OH-VS-SL model $(P_{\text{sl}}(y; \tau[\{\rho[\mathbf{b}] - g(\mathbf{x}) \cdot \mathbf{1}\}/s(\mathbf{x})]))_{y \in [K]}$, where \mathbf{b} and $\mathbf{1}$ are learnable parameter and all-1 vector with appropriate dimension, and g and s are \mathbb{R} - and \mathbb{R}_+ -valued learner models. The \mathbb{R}_+ -valued scale model s can be implemented as $s(\mathbf{x}) = r + \rho(t(\mathbf{x}))$ with a constant $r > 0$ for avoiding zero-division, fixed non-negative function ρ , and another \mathbb{R} -valued learner t .

The significance of OH models is that their scaling factor can vary depending on \mathbf{x} . Extending previous BIN and POI models by (Beckham & Pal, 2017), we also introduce the OH-BIN model $(P_{\text{bin}}(y; g(\mathbf{x}), s(\mathbf{x})))_{y \in [K]}$ and OH-POI model $(P_{\text{poi}}(y; g(\mathbf{x}), s(\mathbf{x})))_{y \in [K]}$. For the implementation $s(\mathbf{x}) = r + \rho(t(\mathbf{x}))$ of the scaling factor for these models, the use of a small $r > 0$ enables to represent a wider class distribution.

These OH models can represent overall heteroscedastic data, but their representation ability is parametrically constrained with the 2-DF; Compare Figures 2, 3, 4, and 5.

⁴The PO-SL model $P_{\text{sl}}(y; \mathbf{b} - g(\mathbf{x}) \cdot \mathbf{1}) = e^{-by} / \sum_{k=1}^K e^{-bk}$ is constant regarding \mathbf{x} and may be less representable. Therefore, Anderson (1984) considered to use a learner model $g(\mathbf{x}) = \mathbf{a} - g(\mathbf{x}) \cdot \mathbf{c}$ with learnable parameters \mathbf{a}, \mathbf{c} ($\in \mathbb{R}^K$ here) s.t. $a_1 = c_1 = 0$ and \mathbb{R} -valued learner model $g : \mathbb{R}^d \rightarrow \mathbb{R}$ (we call g a rotatable-odds (RO) learner model) for the SL link function; we call this model the RO-SL model. We can also consider novel RO-ACL models: they are more representable than the PO-ACL (including PO-ORD-ACL) models, but have no unimodality guarantee since the RO learner model $g(\mathbf{x}) = \mathbf{a} - g(\mathbf{x}) \cdot \mathbf{c}$ does not necessarily satisfy the ordering condition $g_1(\mathbf{x}) \leq \dots \leq g_{K-1}(\mathbf{x})$ (depending on \mathbf{c}).

Figure 5: Instances of the OH-ACL models with $K = 10$.

5.3 Combination-with-Uniform Models

When a CPD $(p_k(\mathbf{x}))_{k \in [K]}$ is unimodal, the combination-with-uniform (CU) CPD, $q(\mathbf{x}) \cdot (\frac{1}{K})_{k \in [K]} + \{1 - q(\mathbf{x})\} \cdot (p_k(\mathbf{x}))_{k \in [K]}$ with e.g., $q(\mathbf{x}) = \frac{1}{1 + e^{-t(\mathbf{x})}}$ for an \mathbb{R} -valued model t , is also unimodal. For example, Piccolo (2003); Iannario & Piccolo (2011) studied a combination of uniform and binomial model. The CU extension of low-DF (like PO and OH) models can increase their DF by 1.

6 Numerical Experiments

6.1 Experimental Purposes

From generic principles, it can be expected that compact likelihoods model that can adequately represent the data will yield better generalization performance in the conditional probability estimation task, and accordingly that OR methods based on such likelihood models will yield better generalization performance in the OR task. We took numerical experiments in order to verify whether our proposed likelihood models, developed based on this working hypothesis, provide better performances than previous unimodal likelihood models with weak representation ability and popular statistical OR models with no unimodality guarantee.

6.2 Experimental Settings

We selected 21 real-world datasets of those used in experiments by the previous OR study (Gutierrez et al., 2015) with the total sample size n_{tot} that is 1000 or more, and used them for numerical experiments.⁵ AB5, ..., CE5' (resp AB10, ..., CE10') are datasets generated by discretizing a real-valued target of datasets, which are often used to benchmark regression methods, by 5 (resp. 10) different bins with equal proportions. SW, ..., CA originally have a categorical target, and the authors of (Gutierrez et al., 2015) judged that their targets have a natural ordinal relation.

Table 1 shows the dataset name, dataset properties, n_{tot} , d , and K , and mean and standard deviation (STD) of 100 test MUs and DRs' MU: The mean unimodality (MU) is a numerical criterion to evaluate the unimodality of the conditional probability distribution of the data, and it is defined, for a likelihood model $\hat{\Pr}(Y = \cdot | \mathbf{X} = \cdot)$ and n used data points, as $\frac{1}{n} \sum_{i=1}^n \mathbb{1}\{(\hat{\Pr}(Y = y | \mathbf{X} = \mathbf{x}_i))_{y \in [K]} \text{ is unimodal.}\}$. As that likelihood model, we used the test SL model under the task-P (described below) that can represent any conditional probability distribution. We trained a likelihood model with a training sample of size $n_{\text{tra}} = 800$, and evaluated the MU with an obtained likelihood model and a remaining test sample of size $n_{\text{tes}} = n_{\text{tot}} - n_{\text{tra}}$. We repeated this procedure 100 trials with a randomly-set different sample setting and initial parameters of the likelihood model to obtain 100 test MUs. We also give a relative reference value in Table 2: It shows a ratio that 10^6 samples of the PMF uniformly and randomly drawn from the probability simplex in \mathbb{R}^K was unimodal for $K = 3, \dots, 10$. Comparing the test MU in Table 1 and the relative reference value of the same K in Table 2, one can find that many real-world data treated as ordinal data by existing OR research tend strongly unimodal. Additionally, introducing the DRs' MU $\frac{1}{n} \sum_{i=1}^n \mathbb{1}\{(\hat{\Pr}(Y = y | \mathbf{X} = \mathbf{x}_i))_{y \in [K]} \text{ satisfied the DRs' unimodality.}\}$, we took a similar procedure; see Tables 1

⁵One can get the datasets from a researchers' site (<http://www.uco.es/grupos/ayrna/orreview>) of (Gutierrez et al., 2015), or our GitHub repository (<https://github.com/Anonymous>) together with our used program codes.

Table 1: It shows the dataset name, dataset properties n_{tot} , d , and K , and mean and STD (as ‘mean \pm STD’) of test MUs and test DRs’ MUs.

dataset name	n_{tot}	d	K	MU	DRs’ MU
AB5 (abalon5)	4177	11	5	.8915 \pm .0662	.4042 \pm .0743
BA5 (bank5)	8192	8	5	.9944 \pm .0387	.4293 \pm .1389
BA5’ (bank5’)	8192	32	5	.9887 \pm .0167	.7041 \pm .1210
CO5 (computer5)	8192	12	5	.9956 \pm .0137	.5623 \pm .1041
CO5’ (computer5’)	8192	21	5	1.0000 \pm .0001	.5851 \pm .1170
CH5 (cal.housing5)	20640	8	5	.9065 \pm .1027	.3993 \pm .1206
CE5 (census5)	22784	8	5	.7887 \pm .0929	.3221 \pm .1050
CE5’ (census5’)	22784	16	5	.8332 \pm .0736	.3827 \pm .0978
AB10 (abalon10)	4177	11	10	.3218 \pm .1315	.0076 \pm .0164
BA10 (bank10)	8192	8	10	.8923 \pm .1463	.0225 \pm .0432
BA10’ (bank10’)	8192	32	10	.5101 \pm .2261	.0019 \pm .0054

dataset name	n_{tot}	d	K	MU	DRs’ MU
CO10 (computer10)	8192	12	10	.8006 \pm .1431	.0935 \pm .0883
CO10’ (computer10’)	8192	21	10	.8189 \pm .1617	.0517 \pm .0737
CH10 (cal.housing10)	20640	8	10	.3311 \pm .1713	.0051 \pm .0095
CE10 (census10)	22784	8	10	.2042 \pm .1017	.0008 \pm .0018
CE10’ (census10’)	22784	16	10	.3151 \pm .1153	.0027 \pm .0053
SW (SWD)	1000	10	4	.9993 \pm .0027	.9853 \pm .0204
LE (LEV)	1000	4	5	.9547 \pm .0666	.2208 \pm .1555
ER (ERA)	1000	4	9	.7909 \pm .1045	.2126 \pm .1403
WR (winequality-red)	1599	11	6	.9894 \pm .0457	.1415 \pm .1259
CA (car)	1728	21	4	.8735 \pm .2037	.0276 \pm .0223

Table 2: It shows a ratio that 10^6 samples of the PMF uniformly and randomly drawn from Δ^{K-1} satisfied the unimodality (resp. the DRs’ unimodality) in the row “MU” (resp. in the row “DRs’ MU”).

K	3	4	5	6	7	8	9	10
MU	.6666	.3330	.1329	.0444	.0125	.0032	.0007	.0001
DRs’ MU	.6666	.3330	.1329	.0276	.0055	.0009	.0001	.0000

and 2. The DRs’ MU for ordinal data were larger than those for uniform random PMF, but their difference appears to be relatively smaller than that for MU.

We considered 4 tasks: a conditional probability estimation task (task-P), and 3 OR tasks with the zero-one task loss ℓ_{zo} (task-Z), absolute task loss ℓ_{abs} (task-A), and squared task loss ℓ_{sq} (task-S). Results for the task-P, -Z, -A, and -S are respectively evaluated based on the negative log likelihood (NLL), mean zero-one error (MZE), mean absolute error (MAE), and mean squared error (MSE). Note that, for a likelihood model $\hat{\text{Pr}}(Y = \cdot | \mathbf{X} = \cdot)$ and n used data points, the NLL is defined as $-\frac{1}{n} \sum_{i=1}^n \log \hat{\text{Pr}}(Y = y_i | \mathbf{X} = \mathbf{x}_i)$, and the MZE, MAE, and MSE are respectively defined as $\frac{1}{n} \sum_{i=1}^n \ell(f(\mathbf{x}_i), y_i)$ with $f(\mathbf{x}) = f_{\ell}(\mathbf{x}; (\hat{\text{Pr}}(Y = y | \mathbf{X} = \mathbf{x}))_{y \in [K]})$ for $\ell = \ell_{\text{zo}}, \ell_{\text{abs}}, \ell_{\text{sq}}$.

We tried 12 likelihood models: previous 1-DF BIN and POI models; proposed 1-DF PO-ORD-ACL and PO-VS-SL models; proposed 2-DF OH-BIN, OH-POI, OH-ORD-ACL, and OH-VS-SL models; previous full-DF ACL and SL models; proposed full-DF ORD-ACL and VS-SL models. Here, we used link functions $\rho = \rho_{\text{exp}}$ to ensure the non-negativity and for a positive learner model for POI and OH-POI models and $\tau = \tau_{\text{sq}}$ to ensure the V-shape, and a positive constant $r = 0.01$ in the scaling model $s(\mathbf{x}) = r + \rho(t(\mathbf{x}))$ for OH models. We implemented all learner models with a 4-layer fully-connected neural network model that shares weights in except for the final layer and has 100 nodes activated with the sigmoid function in addition to bias nodes in every hidden layer. Note that, for the OH models, we implemented their 1DT and scaling models (g and t in the notation in Section 5.2) with two isolated weight-shared networks (each of which is described above), because their performance was significantly degraded when the 1DT and scaling models were implemented with a single weight-shared network. We trained a model with a training sample and Adam optimization for 1000 epochs according to the maximum likelihood estimation, and evaluated the NLL, MZE, MAE, and MSE with a remaining test sample at the end of each epoch. Then, for the task-P, -Z, -A, or -S, we adopted a model at the timing when the test NLL, MZE, MAE, or MSE got minimum as the test model under the corresponding task.

We experimented with 6 training sample size settings $n_{\text{tra}} = 25, 50, 100, 200, 400, 800$, to see the dependence of behaviors of each method on the training sample size.

For all combinations of 21 datasets, 4 tasks, 12 likelihood models, and 6 training sample size settings, we repeated above-described procedure 100 trials with a randomly-set different training sample and initial parameters, to obtain 100 test errors under the corresponding task.

6.3 Experimental Results

First, we provide experimental considerations about comparisons within each set of methods that have interesting differences while maintaining commonality; See Table 3.

Table 3: A cell for an error and training sample size n_{tra} in a sub-table for a group of specified methods shows the total (over the 21 datasets) number of times that each method won in the Mann-Whitney U-test with p -value 0.05 regarding the error for each pair of two methods within the group. In each cell in a sub-table, methods that won the most times are highlighted in the corresponding color.

POI v.s. PO-ORD-ACL							PO-ORD-ACL v.s. OH-ORD-ACL v.s. ORD-ACL						
	$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$		$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$
NLL	12,5	12,5	13,4	14,5	11,8	7,7	NLL	7,15,28	6,13,28	1,14,26	3,16,22	4,14,30	1,14,33
MZE	6,5	5,7	3,11	2,10	3,11	1,11	MZE	6,2,17	6,2,22	3,7,28	2,10,34	3,13,27	3,16,27
MAE	5,1	4,2	1,2	1,4	2,6	2,5	MAE	7,2,16	9,1,22	8,2,26	9,1,27	7,5,22	8,7,26
MSE	7,1	6,1	3,2	1,2	2,2	1,4	MSE	2,0,13	5,0,15	3,1,23	9,1,24	8,4,18	12,4,19
BIN v.s. POI v.s. PO-ORD-ACL v.s. PO-VS-SL							PO-VS-SL v.s. OH-VS-SL v.s. VS-SL						
	$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$		$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$
NLL	37,36,17,20	37,36,14,23	33,34,13,20	32,33,13,20	29,30,16,20	23,24,16,25	NLL	11,22,17	12,24,19	11,26,16	3,20,21	5,15,29	4,19,29
MZE	16,20,6,10	16,14,9,16	23,7,15,20	25,6,13,23	24,7,18,26	22,6,19,26	MZE	9,5,10	8,5,10	4,9,17	2,14,25	0,18,28	1,21,24
MAE	11,13,2,5	7,8,4,7	6,3,5,11	12,2,5,12	13,4,9,11	15,5,9,10	MAE	19,1,9	21,0,11	14,1,17	11,3,22	7,5,26	10,5,25
MSE	15,15,2,4	8,13,2,5	8,8,4,6	10,2,4,5	5,4,5,5	5,4,9,9	MSE	13,1,10	13,0,12	11,3,13	13,3,15	8,3,17	12,5,16
BIN v.s. OH-BIN							ACL v.s. ORD-ACL						
	$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$		$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$
NLL	3,13	4,14	4,11	4,7	7,9	6,11	NLL	5,9	3,8	5,8	9,5	7,2	7,3
MZE	0,6	0,6	0,11	1,15	1,17	0,16	MZE	0,16	0,17	0,11	2,8	3,5	6,5
MAE	1,1	2,5	3,4	5,5	6,6	7,3	MAE	0,15	0,14	1,12	1,8	1,8	2,6
MSE	0,1	1,2	2,3	8,2	14,1	15,0	MSE	0,14	0,14	0,10	1,7	2,8	2,5
POI v.s. OH-POI							SL v.s. VS-SL						
	$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$		$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$
NLL	2,7	4,10	5,10	5,11	6,11	4,14	NLL	2,14	3,15	5,13	5,13	3,12	4,12
MZE	1,8	0,12	0,16	0,18	1,17	1,18	MZE	0,21	0,19	0,18	2,17	2,14	2,10
MAE	1,6	0,7	4,9	5,8	5,8	8,6	MAE	0,21	0,20	0,19	0,17	1,15	2,12
MSE	0,3	2,6	4,6	6,3	9,3	13,1	MSE	0,20	0,20	0,19	0,15	1,15	1,13
OH-BIN v.s. OH-POI v.s. OH-ORD-ACL v.s. OH-VS-SL							ACL v.s. ORD-ACL v.s. SL v.s. VS-SL						
	$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$		$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$
NLL	28,37,8,10	32,35,9,13	23,28,13,13	14,21,17,19	15,15,24,30	16,14,23,36	NLL	40,44,4,16	36,46,3,16	38,41,8,14	36,24,9,16	30,14,8,26	26,14,10,31
MZE	22,30,4,4	23,24,5,9	29,27,6,16	36,18,6,16	34,12,12,24	26,12,15,25	MZE	21,43,0,37	20,44,0,33	18,40,0,32	18,30,4,29	20,22,4,27	21,21,5,22
MAE	19,22,3,2	22,26,3,2	19,26,4,4	25,20,4,6	19,12,12,11	17,14,17,13	MAE	20,42,0,29	19,42,0,29	19,37,0,29	19,29,0,28	15,26,2,26	18,25,3,22
MSE	23,29,2,1	20,23,3,2	17,22,2,2	13,14,7,7	7,12,13,11	9,10,24,18	MSE	19,40,0,28	19,36,0,30	14,37,0,28	18,29,0,21	14,27,2,25	14,24,3,24

POI v.s. PO-ORD-ACL The POI model is a special instance of the PO-ORD-ACL models in which the mode-wise heteroscedasticity is fixed, and has a weaker representation ability. This relation might make the POI model work better when n_{tra} was small and the PO-ORD-ACL model work better when n_{tra} was large, in every task.

BIN v.s. POI v.s. PO-ORD-ACL v.s. PO-VS-SL (1-DF) The PO-ORD-ACL and PO-VS-SL models can adjust the mode-wise heteroscedasticity, while the BIN and POI models cannot. In this respect, the PO-ORD-ACL and PO-VS-SL models tend more representable than the BIN and POI models. For this reason, the BIN and POI models worked better when n_{tra} was small, and the PO-ORD-ACL and PO-VS-SL models worked better when n_{tra} was large.

BIN v.s. OH-BIN, and POI v.s. OH-POI The OH-BIN and OH-POI models are respectively OH-generalizations of the BIN and POI models, and have a stronger representation ability. Thus, OH models were better regarding the NLL that was directly optimized. However, OH models were bad for some tasks (especially, task-A and -S) when n_{tra} was large, which is a counter-intuitive result and requires further analysis but may be because the difference of the restricted structures of the two likelihood models makes a difference in the compatibility between the model and task.

OH-BIN v.s. OH-POI v.s. OH-ORD-ACL v.s. OH-VS-SL (2-DF) The situation is similar to that for the 1-DF models, BIN v.s. POI v.s. PO-ORD-ACL v.s. PO-VS-SL.

PO-ORD-ACL v.s. OH-ORD-ACL v.s. ORD-ACL, and PO-VS-SL v.s. OH-VS-SL v.s. VS-SL The VS-SL, OH-VS-SL, and PO-VS-SL models (and the ORD-ACL, OH-ORD-ACL, and PO-ORD-ACL models) are more representable in that order. Presumably for this reason, the PO- and OH-VS-SL models worked better with small n_{tra} , and the VS-SL model worked better with large n_{tra} .

Table 4: A cell for a method and training sample size n_{tra} in a sub-table for an error shows the total (over the 21 datasets) numbers of times that the method was top 1, 2, and 3 in the rank regarding the mean of the test errors within those for 12 methods. In each block for n_{tra} in a sub-table, the 1st, 2nd, and 3rd best results are respectively highlighted in the red, green, and blue colors.

NLL							MAE						
	$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$		$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$
BIN	0,4,5	0,3,5	0,1,5	0,3,5	2,2,2	0,0,0	BIN	2,4,6	1,6,9	1,4,7	1,2,6	1,4,5	0,0,0
POI	4,9,11	5,7,10	3,5,7	2,2,3	1,1,1	0,0,0	POI	1,10,12	1,4,6	0,2,2	0,0,1	0,0,0	0,1,1
PO-ORD-ACL	0,0,2	0,1,1	1,1,1	1,1,1	0,1,2	0,0,0	PO-ORD-ACL	0,1,2	1,3,4	1,3,6	0,4,4	0,0,2	0,0,0
PO-VS-SL	3,3,4	1,3,4	0,1,2	0,0,1	0,1,2	0,2,2	PO-VS-SL	4,5,9	5,5,8	5,6,11	2,4,7	1,2,2	0,2,4
OH-BIN	0,6,8	1,7,10	5,7,9	2,3,4	1,1,2	1,1,3	OH-BIN	2,4,7	2,3,10	1,5,6	4,4,6	0,1,3	1,1,2
OH-POI	6,9,12	6,8,11	3,7,9	3,4,7	0,2,4	0,1,1	OH-POI	7,10,12	3,9,9	4,5,8	2,4,4	1,4,4	0,2,2
OH-ORD-ACL	0,0,1	0,0,0	1,1,3	1,2,3	0,0,1	0,0,3	OH-ORD-ACL	0,1,1	0,0,1	0,0,0	0,1,1	0,0,1	0,2,3
OH-VS-SL	0,0,3	1,1,2	0,2,3	0,1,2	3,4,5	5,6,8	OH-VS-SL	1,1,2	1,2,2	0,2,2	0,2,2	2,3,5	2,2,2
ACL	4,5,8	3,5,8	2,6,11	4,10,15	6,10,15	4,9,16	ACL	0,0,1	0,0,1	0,1,2	2,4,8	2,7,12	5,9,15
ORD-ACL	4,6,9	4,7,10	6,10,10	7,10,11	0,9,12	0,7,11	ORD-ACL	4,5,10	7,9,10	8,10,13	7,11,13	7,10,13	5,10,15
SL	0,0,0	0,0,0	0,1,2	1,3,3	2,4,6	2,4,6	SL	0,0,0	0,0,0	0,0,0	0,0,1	1,1,2	2,3,7
VS-SL	0,0,0	0,0,2	0,0,1	0,3,8	6,7,11	9,12,13	VS-SL	0,1,1	0,1,3	1,4,6	3,6,10	6,10,14	6,10,12

MZE							MSE						
	$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$		$n_{\text{tra}}=25$	$n_{\text{tra}}=50$	$n_{\text{tra}}=100$	$n_{\text{tra}}=200$	$n_{\text{tra}}=400$	$n_{\text{tra}}=800$
BIN	2,2,5	0,2,3	1,1,2	0,1,2	0,0,0	0,0,0	BIN	1,2,7	3,5,8	2,7,11	1,5,9	2,6,9	1,3,6
POI	1,7,8	1,3,5	0,0,1	0,0,0	0,1,1	0,0,1	POI	4,9,12	4,6,12	1,3,6	0,1,5	0,1,3	0,1,4
PO-ORD-ACL	0,1,6	1,2,3	0,1,1	0,0,0	0,0,1	0,0,1	PO-ORD-ACL	1,1,1	0,1,2	0,2,4	0,2,3	1,1,3	0,3,6
PO-VS-SL	1,4,7	2,2,4	1,3,3	1,2,3	0,0,0	0,1,1	PO-VS-SL	2,3,5	2,2,4	4,4,5	3,6,6	0,5,5	2,5,8
OH-BIN	4,6,10	3,9,11	5,11,13	9,13,13	8,12,14	5,7,7	OH-BIN	2,6,10	2,5,11	2,4,5	0,1,4	0,0,1	0,1,1
OH-POI	9,11,12	5,8,12	3,5,11	2,4,7	0,3,5	0,3,5	OH-POI	8,10,13	4,9,11	2,7,10	2,2,5	0,2,3	1,1,2
OH-ORD-ACL	1,1,1	0,1,2	0,1,4	0,2,3	0,0,1	1,3,5	OH-ORD-ACL	0,1,1	0,0,0	0,0,0	0,0,1	0,1,3	0,2,2
OH-VS-SL	1,2,3	2,3,4	1,4,5	1,3,4	3,6,10	3,3,8	OH-VS-SL	0,1,2	0,1,1	0,1,1	1,1,1	1,2,3	2,2,3
ACL	0,1,1	0,1,2	0,1,3	1,4,4	4,7,8	5,7,9	ACL	1,1,2	1,1,1	0,0,1	1,4,6	3,5,6	1,5,8
ORD-ACL	2,6,6	6,9,13	9,11,12	5,8,10	2,7,8	2,6,12	ORD-ACL	1,6,7	4,11,11	8,10,14	8,11,12	7,9,12	7,9,11
SL	0,0,0	0,0,0	0,0,0	1,1,3	1,1,3	0,3,4	SL	0,0,0	0,0,0	0,0,0	0,2,2	1,2,4	1,3,5
VS-SL	0,1,4	1,2,4	1,4,8	1,4,14	3,5,12	5,9,10	VS-SL	1,2,3	1,1,2	2,4,6	5,7,9	6,8,11	6,7,7

ACL v.s. ORD-ACL The ACL model can represent any data, and the ORD-ACL model can represent unimodal data with the DRs’ unimodality constraint. The unimodality of the ORD-ACL model might improve the generalization performance when n_{tra} was small, but the ORD-ACL model was not good when n_{tra} was large perhaps because its DRs’ unimodality constraint restricted its representation ability too much for our tried data; see again DRs’ MU in Tables 1 and 2.

SL v.s. VS-SL The SL model can represent any data, and the VS-SL model can represent any unimodal data. Therefore, the VS-SL model has more restricted representation ability than the SL model. Unlike the comparison between ACL and ORD-ACL models, the VS-SL model gave better results than the SL model, probably thanks to the validity of the unimodal hypothesis.

ACL v.s. ORD-ACL v.s. SL v.s. VS-SL (full-DF) The ORD-ACL model was the best in many cases, and the VS-SL model was better when n_{tra} was large. This is thought due to their unimodality and the DRs’ unimodality constraint of the ORD-ACL model.

In summary, POI, PO-VS-SL, OH and ORD-ACL models were especially good when n_{tra} was small, ORD-ACL and VS-SL models were especially good when n_{tra} was large, as shown in Table 4. These results imply the significance of our proposed unimodal likelihood models. For further details of the experimental results, refer to the supplement.

7 Conclusion

In this paper, we pointed out that previous unimodal BIN and POI models have a weak representation ability from the perspective of the DRs’ unimodality constraint, scale, skewness, and DF, and then developed more representable unimodal ORD-ACL and VS-SL models as well as their PO-constrained version and its OH-extension. In our experiments, 1-DF or 2-DF OH models worked better when the training sample size was small, and full-DF ORD-ACL and VS-SL models worked better (than low-DF models or full-DF non-unimodal models) when the training sample size was large due to their stronger representation ability and unimodality.

References

- Alan Agresti. *Analysis of Ordinal Categorical Data*, volume 656. John Wiley & Sons, 2010.
- John A Anderson. Regression and ordered categorical variables. *Journal of the Royal Statistical Society: Series B (Methodological)*, 46(1):1–22, 1984.
- David Andrich. A rating formulation for ordered response categories. *Psychometrika*, 43(4):561–573, 1978.
- Anonymous. Modified threshold method for ordinal regression. *Submitted to Transactions on Machine Learning Research*, 2022. URL <https://openreview.net/forum?id=PInXz6Gasv>. Under review.
- Christopher Beckham and Christopher Pal. Unimodal probability distributions for deep ordinal classification. In *International Conference on Machine Learning*, pp. 411–419, 2017.
- Ana M Bianco and Victor J Yohai. Robust estimation in the logistic regression model. In *Robust Statistics, Data Analysis, and Computer Intensive Methods*, pp. 17–34. 1996.
- Paul-Christian Bürkner and Matti Vuorre. Ordinal regression models in psychology: A tutorial. *Advances in Methods and Practices in Psychological Science*, 2(1):77–101, 2019.
- Wenzhi Cao, Vahid Mirjalili, and Sebastian Raschka. Consistent rank logits for ordinal regression with convolutional neural networks. *arXiv preprint arXiv:1901.07884*, 6, 2019.
- Chuansheng Chen, Shin-Ying Lee, and Harold W Stevenson. Response style and cross-cultural comparisons of rating scales among east asian and north american students. *Psychological Science*, 6(3):170–175, 1995.
- Christophe Croux, Gentiane Haesbroeck, and Christel Ruwet. Robust estimation for ordinal regression. *Journal of Statistical Planning and Inference*, 143(9):1486–1499, 2013.
- Joaquim F Pinto da Costa, Hugo Alonso, and Jaime S Cardoso. The unimodal model for the classification of ordinal data. *Neural Networks*, 21(1):78–91, 2008.
- Leo A Goodman. Simple models for the analysis of association in cross-classifications having ordered categories. *Journal of the American Statistical Association*, 74(367):537–552, 1979.
- Pedro Antonio Gutierrez, Maria Perez-Ortiz, Javier Sanchez-Monedero, Francisco Fernandez-Navarro, and Cesar Hervas-Martinez. Ordinal regression methods: survey and experimental study. *IEEE Transactions on Knowledge and Data Engineering*, 28(1):127–146, 2015.
- Maria Iannario and Domenico Piccolo. Cub models: Statistical methods and empirical evidence. *Modern Analysis of Customer Surveys: with Applications using R*, pp. 231–258, 2011.
- Rob Kaas and Jan M Buhrman. Mean, median and mode in binomial distributions. *Statistica Neerlandica*, 34(1):13–18, 1980.
- Kyoung-Jae Kim and Hyunchul Ahn. A corporate credit rating model using multi-class support vector machines with an ordinal pairwise partitioning approach. *Computers & Operations Research*, 39(8):1800–1811, 2012.
- Tie-Yan Liu. *Learning to Rank for Information Retrieval*. Springer Science & Business Media, 2011.
- Geoff N Masters. A rasch model for partial credit scoring. *Psychometrika*, 47(2):149–174, 1982.
- Peter McCullagh. Regression models for ordinal data. *Journal of the Royal Statistical Society: Series B (Methodological)*, 42(2):109–127, 1980.
- Zhenxing Niu, Mo Zhou, Le Wang, Xinbo Gao, and Gang Hua. Ordinal regression with multiple output cnn for age estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4920–4928, 2016.

Domenico Piccolo. On the moments of a mixture of uniform and shifted binomial random variables. *Quaderni di Statistica*, 5(1):85–104, 2003.

Gary Simon. Alternative analyses for the singly-ordered contingency table. *Journal of the American Statistical Association*, 69(348):971–976, 1974.

Shipeng Yu, Kai Yu, Volker Tresp, and Hans-Peter Kriegel. Collaborative ordinal regression. In *Proceedings of the International Conference on Machine Learning*, pp. 1089–1096, 2006.

A Supplemental Experimental Results

We here describe details of the experimental results.

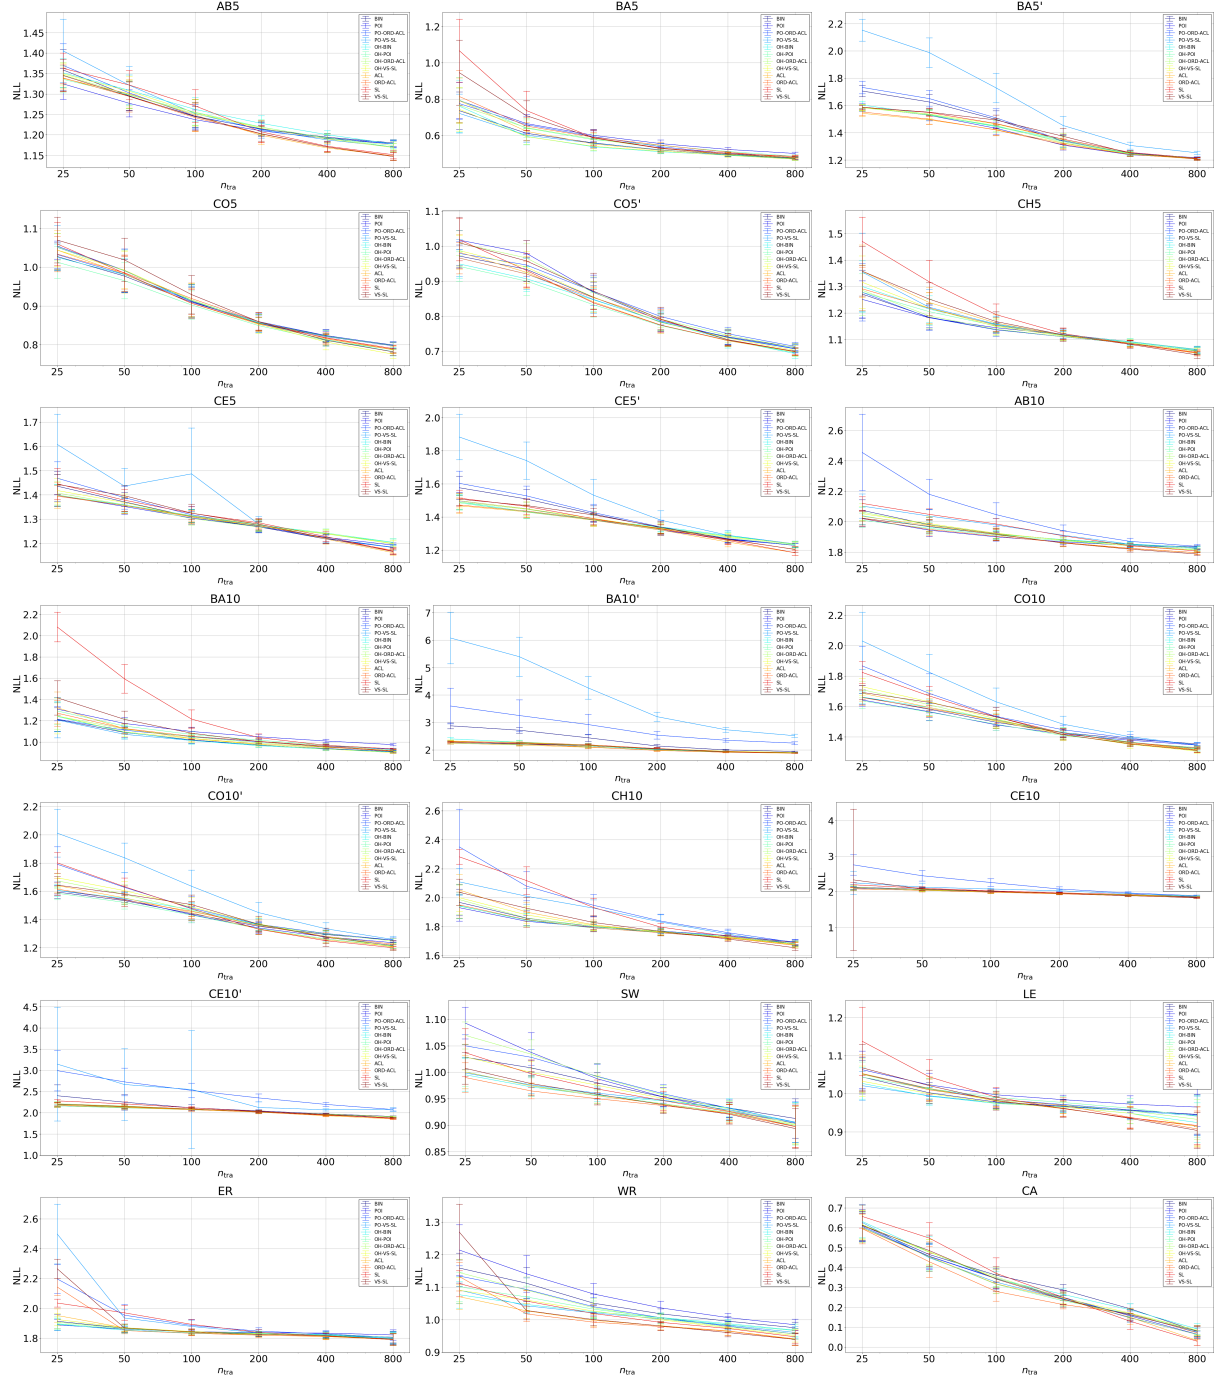


Figure 6: It shows errorbar-plots of mean and STD of the test NLLs of all methods for each dataset.

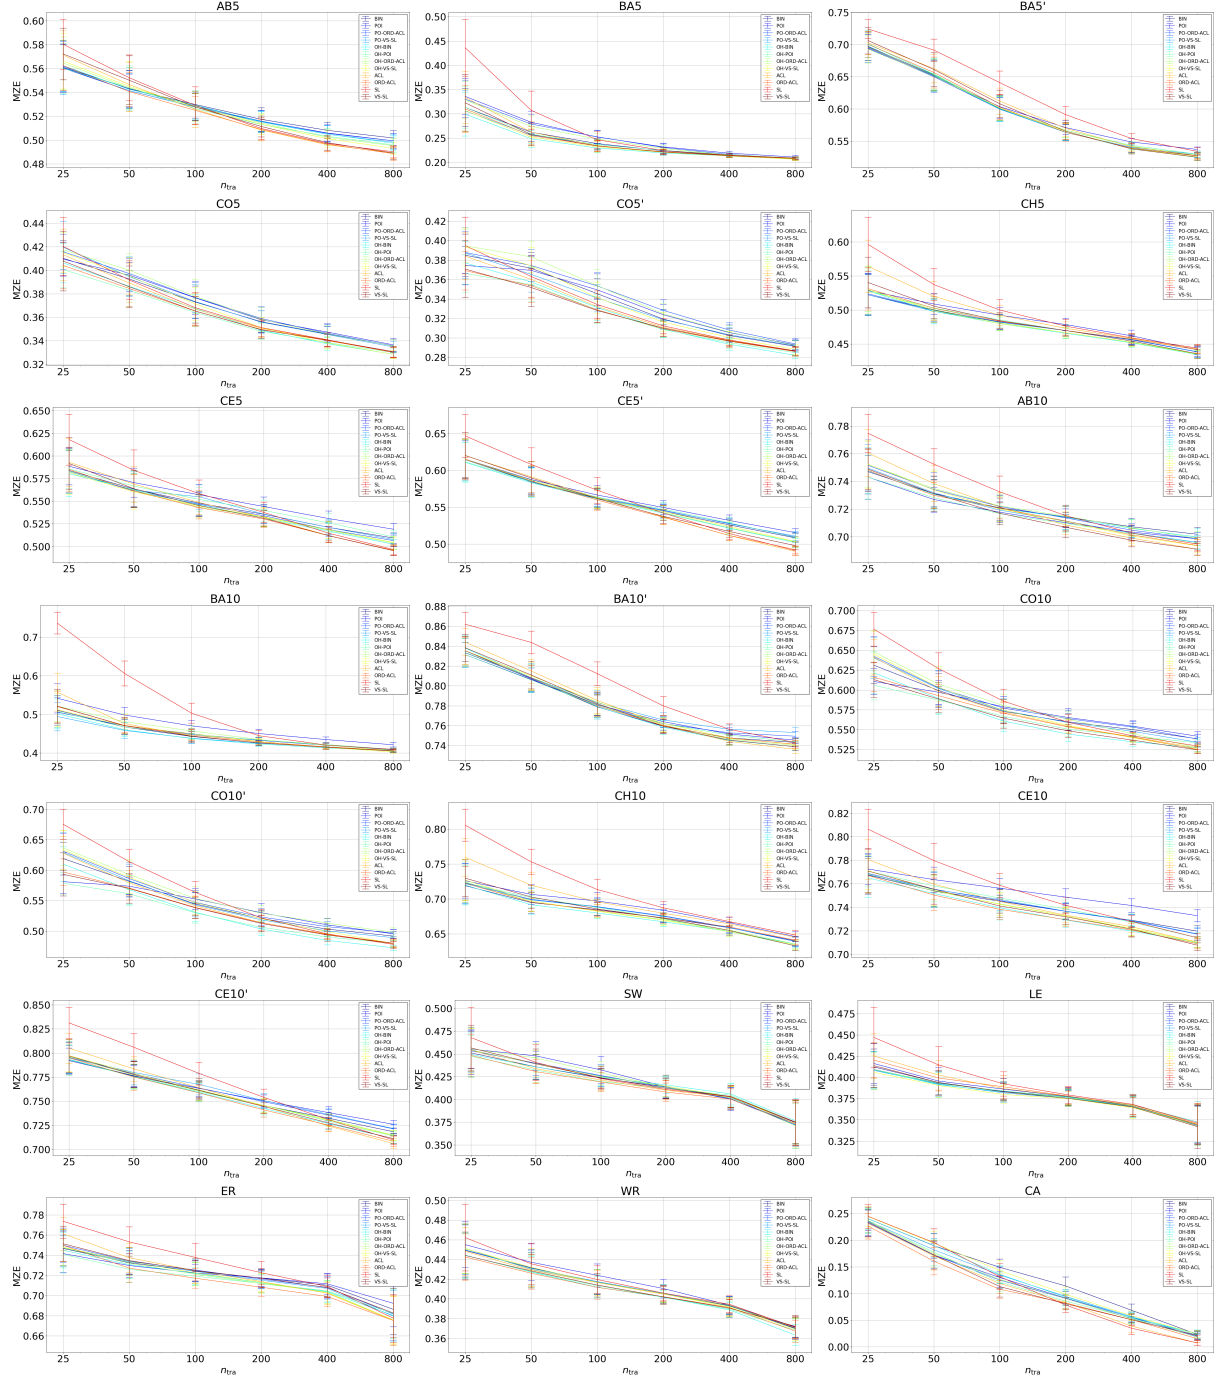


Figure 7: It shows errorbar-plots of mean and STD of the test MZEs of all methods for each dataset.

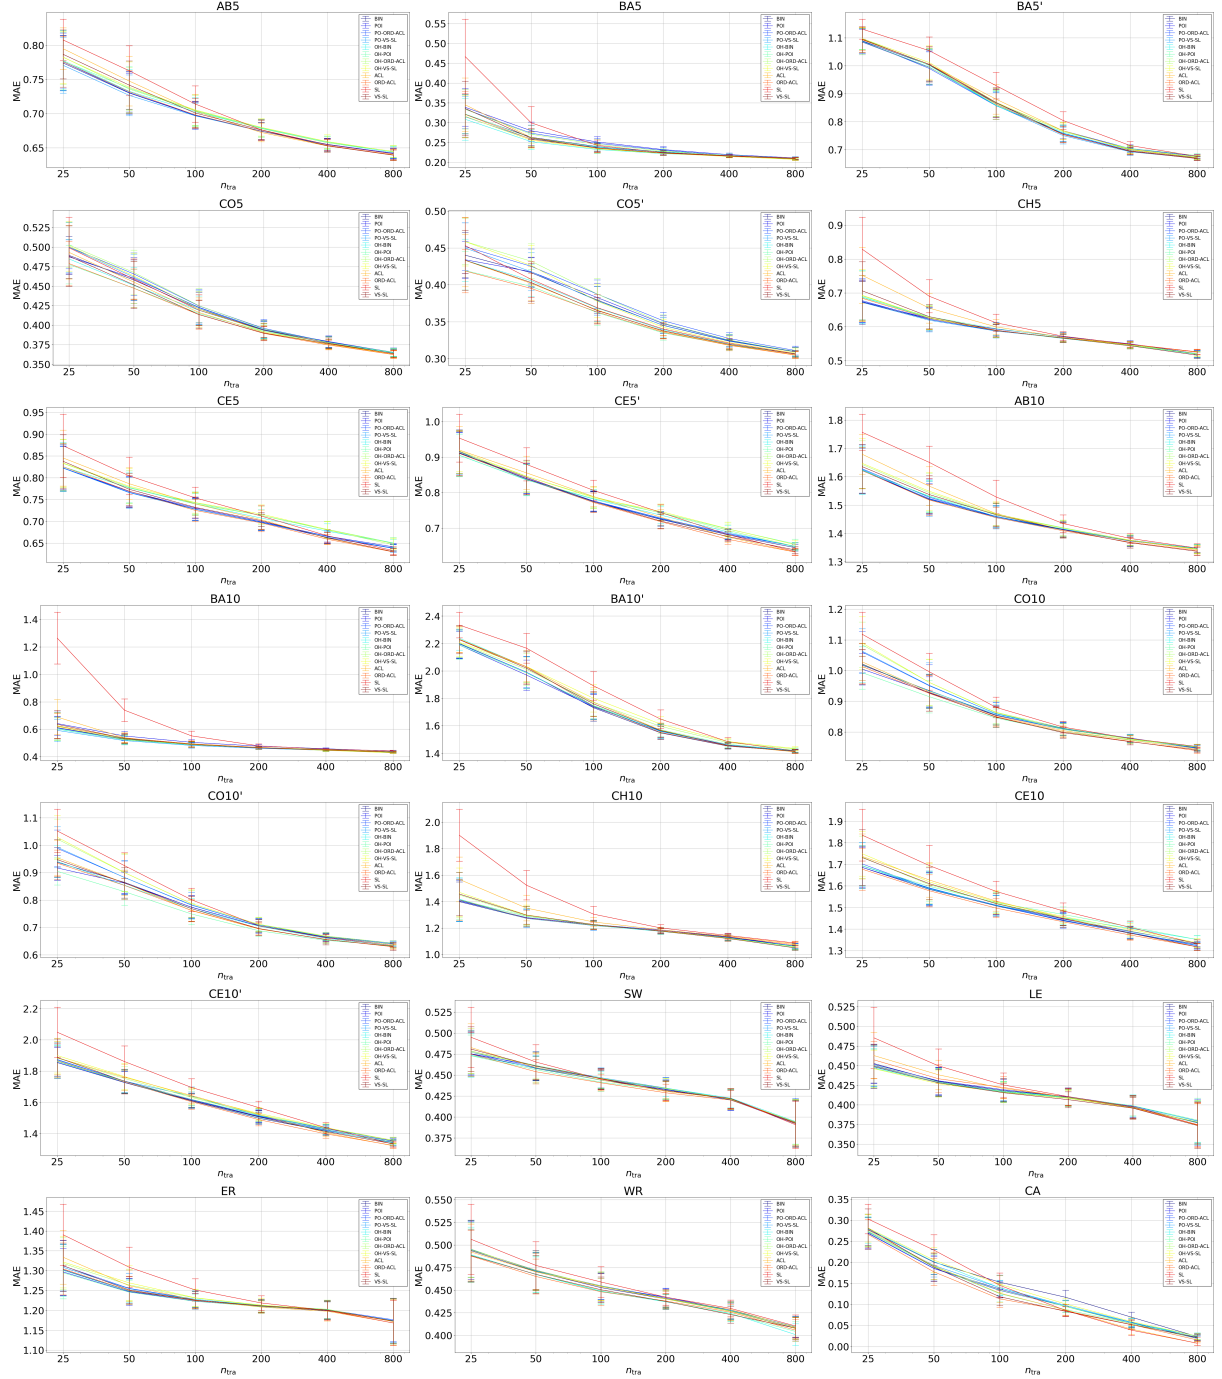


Figure 8: It shows errorbar-plots of mean and STD of the test MAEs of all methods for each dataset.

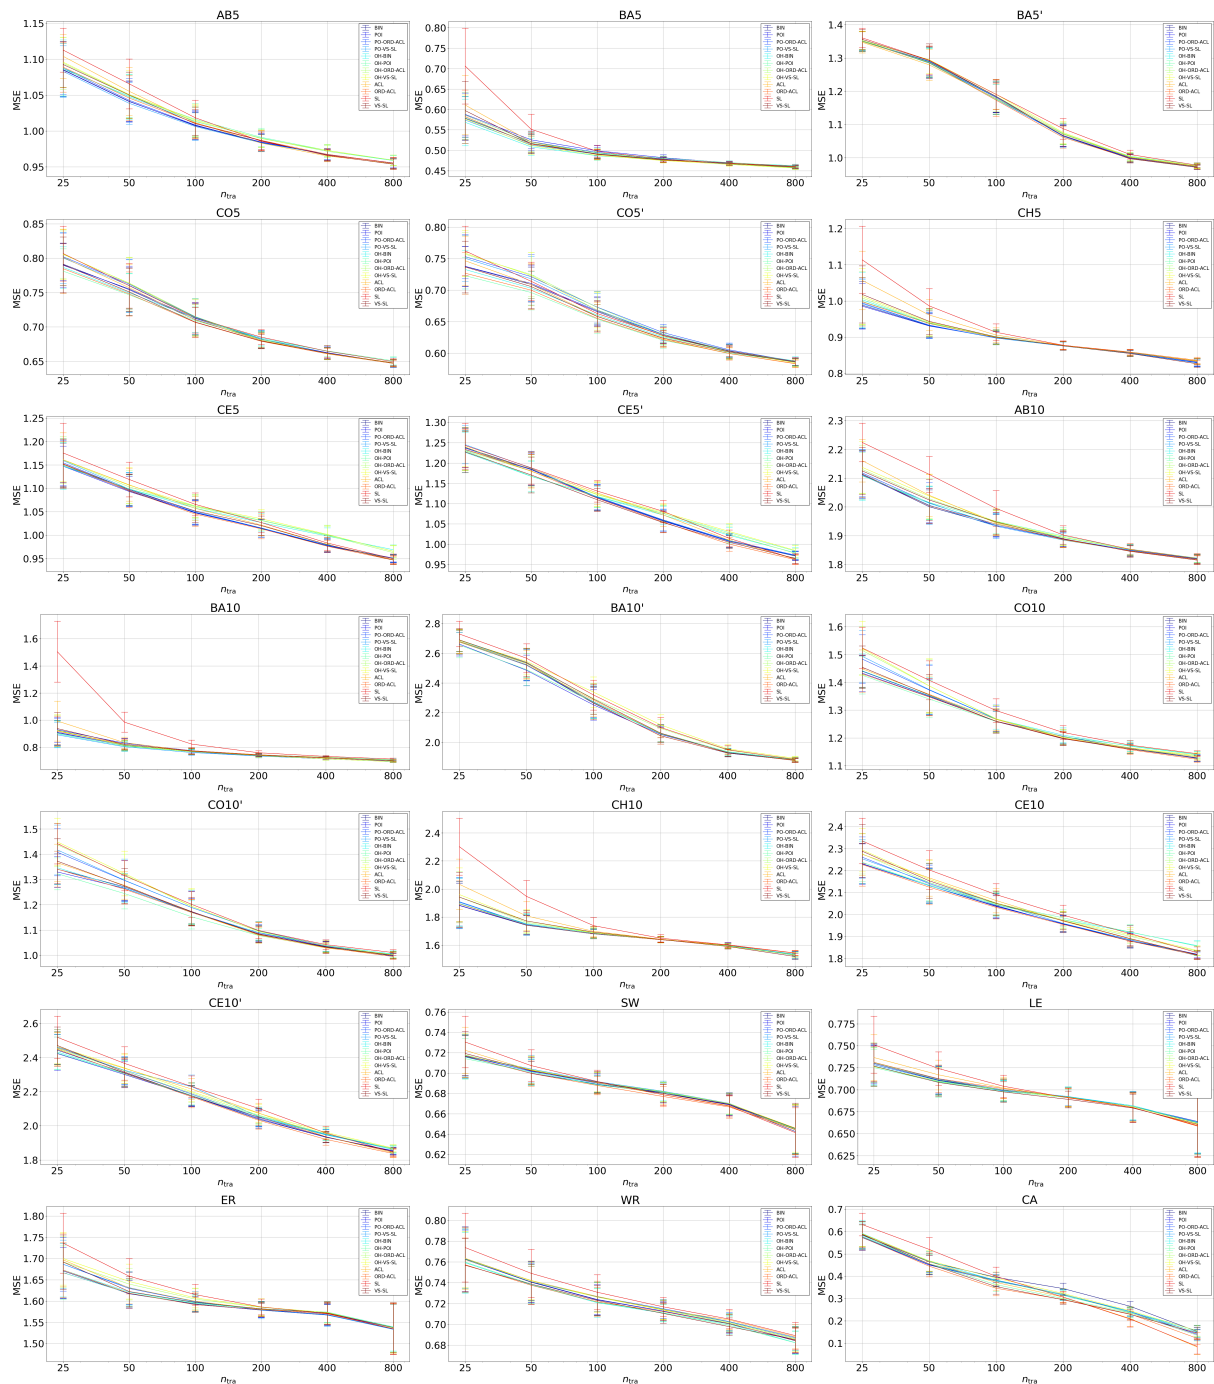


Figure 9: It shows errorbar-plots of mean and STD of the test MSEs of all methods for each dataset.