

Better Assumptions, Stronger Conclusions: The Case for Ordinal Regression in HCI

Brandon Victor Syiem
School of Computer Science
The University of Sydney
Sydney, New South Wales, Australia
brandon.syiem@sydney.edu.au

Eduardo Velloso
School of Computer Science
The University of Sydney
Sydney, New South Wales, Australia
eduardo.velloso@sydney.edu.au

Abstract

Despite the widespread use of ordinal measures in HCI, such as Likert-items, there is little consensus among HCI researchers on the statistical methods used for analysing such data. Both parametric and non-parametric methods have been extensively used within the discipline, with limited reflection on their assumptions and appropriateness for such analyses. In this paper, we examine recent HCI works that report statistical analyses of ordinal measures. We highlight prevalent methods used, discuss their limitations and spotlight key assumptions and oversights that diminish the insights drawn from these methods. Finally, we champion and detail the use of cumulative link (mixed) models (CLM/CLMM) for analysing ordinal data. Further, we provide practical worked examples of applying CLM/CLMMs using R to published open-sourced datasets. This work contributes towards a better understanding of the statistical methods used to analyse ordinal data in HCI and helps to consolidate practices for future work.

CCS Concepts

• **Human-centered computing** → **HCI design and evaluation methods**.

Keywords

Statistics, Ordinal Regression, Ordinal Data, Cumulative Link Models, Cumulative Link Mixed Models

ACM Reference Format:

Brandon Victor Syiem and Eduardo Velloso. 2026. Better Assumptions, Stronger Conclusions: The Case for Ordinal Regression in HCI. In *Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 21 pages. <https://doi.org/10.1145/3772318.3790821>

1 Introduction

Ordinal measures, such as Likert-item surveys, score-based questionnaires, and user preference rankings, are extensively used in HCI research. These measures enable researchers to gather quantitative estimates of users' subjective perceptions, psychological states, attitudes, preferences, judgements, and traits, allowing direct comparative analysis between experimental conditions, otherwise impossible through qualitative data alone [60]. Despite the ubiquity

and pertinence of such measures in HCI, there is a lack of consensus on the current statistical methods used to analyse such data. This has potential long-term consequences that can hinder HCI researchers from aggregating, contrasting, and replicating findings derived from ordinal data.

Both parametric and non-parametric statistical methods are frequently encountered for analysing ordinal data in the HCI literature. Parametric methods, such as the one-way repeated ANOVA test, offer more statistical power when compared to their commonly used non-parametric counterparts, such as the Friedman test. However, this increase in statistical power of parametric methods comes at the cost of stronger assumptions, typically on the distribution, ordering, and interval spacing of the data. Previous work has raised significant concerns in analysing ordinal data with parametric methods [12, 37]. However, these concerns largely stem from the metric (ordered and evenly spaced intervals) handling of ordinal data in widely used methods, as opposed to the method's parametric nature. Analysing ordinal data with metric assumptions has been shown to increase Type I (false positives) and Type II (false negatives) errors, and even inverse the estimated means between groups [43].

Comparatively, non-parametric methods impose fewer assumptions on data distribution and interval spacing. Importantly, *most* non-parametric methods do not treat ordinal data as metric, and are therefore seen as appropriate for analysing such data [36, 37, 55]. However, non-parametric methods are less sensitive, and could prevent researchers from detecting existing effects [58]. Moreover, ongoing debate regarding the metric/non-metric nature of data *derived* from ordinal measures, such as aggregated Likert scales [13, 49], further complicate the selection of appropriate methods. These challenges are particularly salient in HCI, where limited discourse [37, 60] and the conflation of concepts such as 'metric' and 'parametric' have produced inconsistencies in the statistical analysis of ordinal data.

Such methodological disagreement in HCI research can lead to the propagation of erroneous insights, prevent methodologically distinct but related studies from being cross-examined, and hinder the process of deriving theory from empirical results. To address these issues, we survey current HCI literature to better understand the different statistical methods used for analysing ordinal data. We discuss the underlying mechanisms and assumptions used by frequently used methods in HCI, and deliberate on their appropriateness for handling ordinal data. We then spotlight and promote the use of cumulative link (mixed) models, or CL(M)Ms, that appropriately treat ordinal data as categorical while leveraging the inherent ordering between the categories. We explain the intuition and theory behind CL(M)Ms and detail its use and interpretation



This work is licensed under a Creative Commons Attribution 4.0 International License. *CHI '26, Barcelona, Spain*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2278-3/26/04
<https://doi.org/10.1145/3772318.3790821>

through a practical demonstration on open-sourced HCI datasets using R.

We show that despite past criticisms of analysing ordinal data with metric assumptions, they continue to be frequently used in HCI research. This includes frequently used parametric tests with known metric assumptions, as well as commonly used non-parametric methods often regarded as appropriate for ordinal data. Our findings and discussion serve to better inform HCI researchers of the assumptions and pitfalls in current practices employed for statistically analysing ordinal data. Finally, our description of CL(M)Ms and demonstration of its application using R on past open-sourced HCI data provide researchers with the information needed to apply and interpret such methods — previously considered overly complex and inaccessible to HCI researchers [65].

2 Background

2.1 Ordinal Measures, Likert-scales, and HCI

Ordinal measures are **categorical data that have inherent ordering**. Take, for example, a survey item that asks participants to rate how satisfied they were with using a new image editing software on a scale of 1 to 5, with 1 being ‘Very dissatisfied’ and 5 being ‘Very satisfied’. The participant’s response — a whole number score between 1 and 5 — is an ordinal measure that clearly indicates better user satisfaction as the rating goes up. However, ordinal data **carry no metric information**. This means that the distance between ordinal levels may not be equal, and the labels are *purely categorical with known ordering*. As such, the levels 1 to 5 in our example can be replaced by any other set of ordered categories, such as A to E or non-contiguous numerical labels (e.g., 1,2,7,16, and 41), without any loss of information. However, this also suggests that typical statistical methods that involve arithmetic operations to analyse or describe data, such as the data average, cannot be meaningfully interpreted when applied to ordinal data [45].

Despite such inhospitable properties to arithmetic operations, contiguous numerical labels continue to be the most frequently used labels for representing different levels of ordinal data. This numerical representation often leads ordinal data to be interpreted as an interval or ratio scale. However, the appropriateness of treating ordinal data, and data derived from ordinal measures, as metric has stirred considerable debate across multiple scientific disciplines [13, 37, 43, 45, 60] that has spanned nearly a century [59].

This has not deterred the use of experimental methods that generate ordinal data in scientific disciplines. HCI research, for example, widely makes use of ordinal measures in Likert-scales, surveys and questionnaires, such as the NASA-TLX [33] or the system usability scale (SUS) [8]. Likert scales for instance, consist of multiple related survey items (or Likert items), each with an ordinal response. Their widespread use in multiple disciplines, including HCI, psychology, economics, and medicine, among others [43], places them at the heart of the dispute surrounding ordinal data analysis. Despite these challenges, Likert scales are invaluable for collecting quantitative estimates of subjective phenomena otherwise unobtainable through qualitative data alone [60].

Likert scales present an additional dilemma to researchers, i.e., to analyse the scale as a single entity or to analyse individual Likert items. The prescribed method for analysing responses to Likert

scales, as per Likert himself [44], was to aggregate (sum or average) responses to individual Likert items prior to analysis. This treats the data as metric, imposing assumptions so as to meaningfully apply arithmetic operations for analysis. Numerous prior works, in various disciplines, have criticised this approach [36, 43, 45], stating that averaging ordinal data cannot produce interval scaled values. Others, however, have argued that aggregated ordinal data is perfectly suitable for metric interpretations, and can be analysed using metric models given their robustness [13, 14, 55].

In contrast, researchers largely agree that individual Likert items, and Likert-type items [37] (single survey items with ordinal responses that are independent of Likert scales), should not be statistically analysed using metric models [13, 14, 37]. Liddell and Kruschke [43] present strong arguments against treating averaged Likert scales and Likert-type items as metric, and demonstrate that the intuition behind the interval interpretation of averaged ordinal data is wrong. These disputes are ongoing, and we refrain from prescribing any single interpretation. Instead, we highlight these disputes to demonstrate the extent and persistence of the disagreement surrounding ordinal data analysis, which continues to hinder analytical consistency and reproducibility within HCI.

Additionally, HCI research also frequently employs methods other than Likert scales and Likert-type items that generate ordinal data [23]. A less obvious example of data interpretable as ordinal in HCI is that of a pre-/post-experiment test on topic understanding that results in a numerical score. While such a test score within a bounded numerical range (say range represented by R) intuitively seems to be metric, there is no guarantee that a score increase from x to $x + 1$ represents the same level of increase in understanding as a score increase from y to $y + 1$ (where $\{x, y\} \in R$), as this would imply that all questions are equally difficult for all students. The extensive use of ordinal measures in HCI, combined with the inconclusive disputes surrounding its treatment, warrants a deeper reflection on the recent methods used to analyse ordinal data in HCI research. While past work has provided insights into specific HCI subtopics [60] and suggestions for analysing ordinal data in HCI [37, 65], there has been little reflection on the frequency of use of ordinal measures, the methods used to analyse them, and methodological consensus within the field.

2.2 The Parametric versus Non-parametric Debate

While concerns about ordinal data analysis primarily stem from metric assumptions, the widespread use of parametric methods that impose these assumptions has led prior work to discuss metric and parametric concepts interchangeably. Consequently, much of the debate surrounding ordinal data has regrettably shifted towards the appropriateness of parametric and non-parametric methods for analysis. For instance, commonly used parametric models in HCI, such as ANOVA and Student’s t -test, place assumptions on the distribution of the data to enable meaningful statements to be made around arithmetic estimates, such as the mean and standard deviation. This makes such models appropriate when data can be interpreted as interval or ratio scale. In addition, these methods are not invariant to monotone transformations, i.e., transformations that change the absolute magnitude of the data while preserving

relative order. This makes them unsuitable for analysing ordinal data, as ordinal data contain no metric information, and any changes to the magnitude of scale labels should not affect analysis results.

In contrast, most non-parametric models do not impose similar assumptions on data distribution, and are invariant to monotonic transformations. Instead, non-parametric models frequently encountered in HCI, such as the Friedman test, transform the data into relative ordered ranks to further process for insights. For example, the ordered data sets of {5, 2, 3} (ordinal levels: 1 - 7) and {D, B, C} (ordered levels: A - G), will be ranked as {3, 1, 2}. This makes non-parametric models agnostic to changes in magnitude as well; distinct data groups containing {7, 1, 2} and {5, 1, 2} will similarly be ranked as {3, 1, 2} regardless of how much greater 7 is to 5 in the original data. Consequently, these methods are less sensitive to detecting effects when compared to their parametric counterparts, as the former ignores differences in magnitude entirely. However, they are considered appropriate for ordinal data analysis as they impose less stringent assumptions [60].

In addition to sensitivity concerns, there is also a lack of accessible non-parametric methods for analysing multi-factorial experimental designs for HCI researchers. Prior work has addressed this challenge by highlighting existing methods and providing access to tools that enable non-parametric data analysis for multi-factorial designs [37, 65]. For instance, Kaptein et al. [37] highlights existing non-parametric methods used in the medical field for analysing Likert-type item data. Similarly, Wobbrock et al. [65] introduces the aligned ranked transform (ART) method that enables practitioners to apply familiar and powerful multi-factorial ANOVA tests on transformed non-parametric data. Despite these efforts, there is still little consensus within the HCI discipline on how to best analyse ordinal data [23].

A largely neglected concern with using non-parametric methods for ordinal data analysis within HCI relates to test-specific assumptions. Specifically, while non-parametric methods can be used to analyse ordinal data without *imposing the normality assumption* (i.e., assuming that the residuals are normally distributed), they may *still assume metric data*. Non-parametric methods, such as the Wilcoxon Rank Sum Test ¹ or the Friedman Test ², are examples of tests that do not assume normality of residuals or metric data. Therefore, these methods do not violate any of the properties of ordinal data, but do ignore the difference in distances between ordinal levels, which may reduce statistical power. In contrast, methods such as the Quade Test or the Wilcoxon Signed Rank Test for paired samples first calculate the *paired differences prior to assigning ranks*. This makes these tests unsuitable for strictly ordinal data as they perform arithmetic operations on data with no metric information [39]. Similarly, while it is suggested that the ART method can be used for ordinal data analysis [65], the process of ‘aligning’ or transforming the data in ART requires the calculation of averages and residuals from the unaltered ordinal data, and, as such, imposes metric assumptions on the data that may not be suitable for ordinal data.

Imposing such assumptions is unlikely [59] to be problematic by themselves, provided the assumptions and possible drawbacks

are made explicit to readers. However, more recent approaches have been developed that are specifically designed to leverage the information held by ordinal data, without imposing incompatible assumptions on the data, such as with the use of cumulative link (mixed) models, or CL(M)Ms. The availability of these methods, along with software and programming tools that make these methods more accessible and practical to use, presents an opportunity to reconcile ordinal data analysis practices in HCI and move closer to more robust, theoretically informed, and reproducible findings.

3 Review of Ordinal Analysis Practices in HCI

To understand current methodological practices for analysing ordinal data in HCI, we conducted a review of recent research articles published in the proceedings of the ACM CHI Conference on Human Factors in Computing Systems (CHI). We chose CHI based on its reputation as the flagship conference for HCI research, and its broader representation of HCI topics (unlike more focused venues, such as ISMAR or Ubicomp).

3.1 Sampling

As our aim was to understand *recent* statistical approaches for ordinal data analysis in HCI, we restricted our search to the year 2024 and used a broad search string: *[All Field: (“questionnaire” OR “likert”)]*. We also limited our search to full papers. We included papers that conducted user studies and reported statistical analyses of ordinal data. We excluded papers that (1) did not include ordinal data, (2) focused on scale development and not on generating insights from analysing ordinal data, and (3) did not report analyses of included ordinal data beyond summary statistics.

Our search resulted in 558 papers. We randomly sampled 100 papers from the search results for full-text analysis. From the 100 sampled papers, we excluded 6 papers, resulting in a final sample of 94 papers (details of the review steps are presented in Figure 1). For each paper, we recorded all ordinal measures collected and the associated construct. We then recorded the statistical method used to analyse ordinal measures related to each construct. Note that a single paper could report multiple studies, and each study may collect multiple ordinal measures that may be analysed using multiple statistical methods. All extracted data can be found in our supplemental material.

3.2 Findings

We separate our findings into three categories of statistical analysis methods; (1) statistical methods used to *mathematically represent the data, and can be used for parameter estimation and outcome predictions in addition to hypothesis testing* — **predictive modelling**, (2) statistical methods used *primarily for hypothesis testing involving multiple factors and/or more than two groups* — **omnibus tests**, and (3) statistical methods used *primarily for hypothesis tests between two groups, or for performing multiple two group comparisons with corrections* ³ (such as the Bonferroni correction) — **pairwise comparisons/multiple pairwise comparisons**.

¹ranks data points across both samples to be compared

²ranks data points within each repeating block across all levels of a factor

³Note that we do not extract data related to the specific corrections used during multiple pairwise comparisons

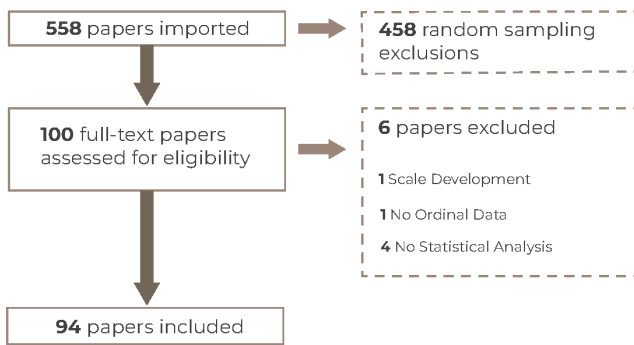


Figure 1: Overview of our review process detailing each step reported as per the PRISMA guidelines [47]. We randomly sampled 100 papers from our initial search result of 558 papers. The full text of all 100 papers was screened for eligibility. Six (6) papers were excluded, and the remaining 94 papers were included in our final sample.




Figure 2: Frequency of different predictive models reportedly used for ordinal data analysis in our sample.

Predictive modelling. Only 28 instances of predictive modelling were reported for ordinal data analysis in our sample. Figure 4 shows the frequency of the different reported models used in our sample. Eight analyses used the Cumulative Link Mixed Model (CLMM), 2 used the Cumulative Link Model (CLM), 3 uses of the Generalised Estimation Equation, 1 use of Generalised Linear Mixed Model (with a logit link function – not cumulative logit which would class it as a CLMM), 9 instances of Linear Mixed-Effects Models, 4 Linear Regression Models, and 1 analysis using the PROCESS model 7 [34].

Omnibus tests. A total of 257 omnibus tests were reported to be used for ordinal data analysis in our sample. Figure 3 shows the frequency of the different reported tests. Parametric methods comprised 71 reported uses of various Analysis of Variance (ANOVA) tests, including 17 independent-sample ANOVA, 37 repeated-measure ANOVA (ANOVA-RM), 6 mixed-effects ANOVA (ANOVA-ME), 1 Analysis of Covariance (ANCOVA), and 10 multivariate ANOVA (MANOVA). Reported non-parametric methods include 60 reported cases of the ART method to transform data prior to ANOVA analysis [65] (ART-ANOVA), 70 analyses using the Kruskal-Wallis test, and 56 analyses using the Friedman test.




Figure 3: Frequency of different omnibus tests used for ordinal data analysis in our sample.




Figure 4: Frequency of different pairwise comparisons reportedly used for ordinal data analysis in our sample.

Pairwise comparisons. 309 pairwise comparisons were reportedly conducted for ordinal data analysis in our sample. Figure 4 shows the frequency of the different reported tests. Parametric tests were dominated by paired (54) and independent-sample t-tests (30), with a minority employing the Tukey Honest Significant Difference (Tukey HSD) test (7). Non-parametric tests primarily consisted of the Wilcoxon Signed Rank test (61) for paired samples, and the Wilcoxon Rank Sum test (34) for independent samples. Other non-parametric tests observed include the Chi-squared test (5), the Dunn test (28), the Conover test (11), Games-Howell test (1), the Kolmogorov-Smirnov test (1), the Nemenyi test (2), and the ART-Contrast tests (33) following the ART-ANOVA tests. Additionally, we found three (3) instances of pairwise comparisons based on estimated marginal means (EMMs Comparison). Note that ‘EMMs Comparison’ is not a single statistical test, but refers to post-hoc pairwise comparisons of estimated marginal means derived from a fitted model [42]—observed following 1 GLMM and 2 CLM model fits in our sample⁴. Finally, we found multiple post-hoc tests in our sample that did not specify the exact test used and simply referred to pairwise comparisons as ‘post-hoc analysis/test’. This accounted for 39 instances in our sample, labelled as ‘unspecified’.

3.3 Discussion

An undecided discipline: Our review demonstrates the use of numerous statistical approaches for analysing ordinal data in HCI

⁴We classify ‘EMMs Comparison’ as parametric because estimated marginal means depend on model parameters. However, it can also be applied in non-parametric contexts, such as post-hoc comparisons following an ART-ANOVA model fit

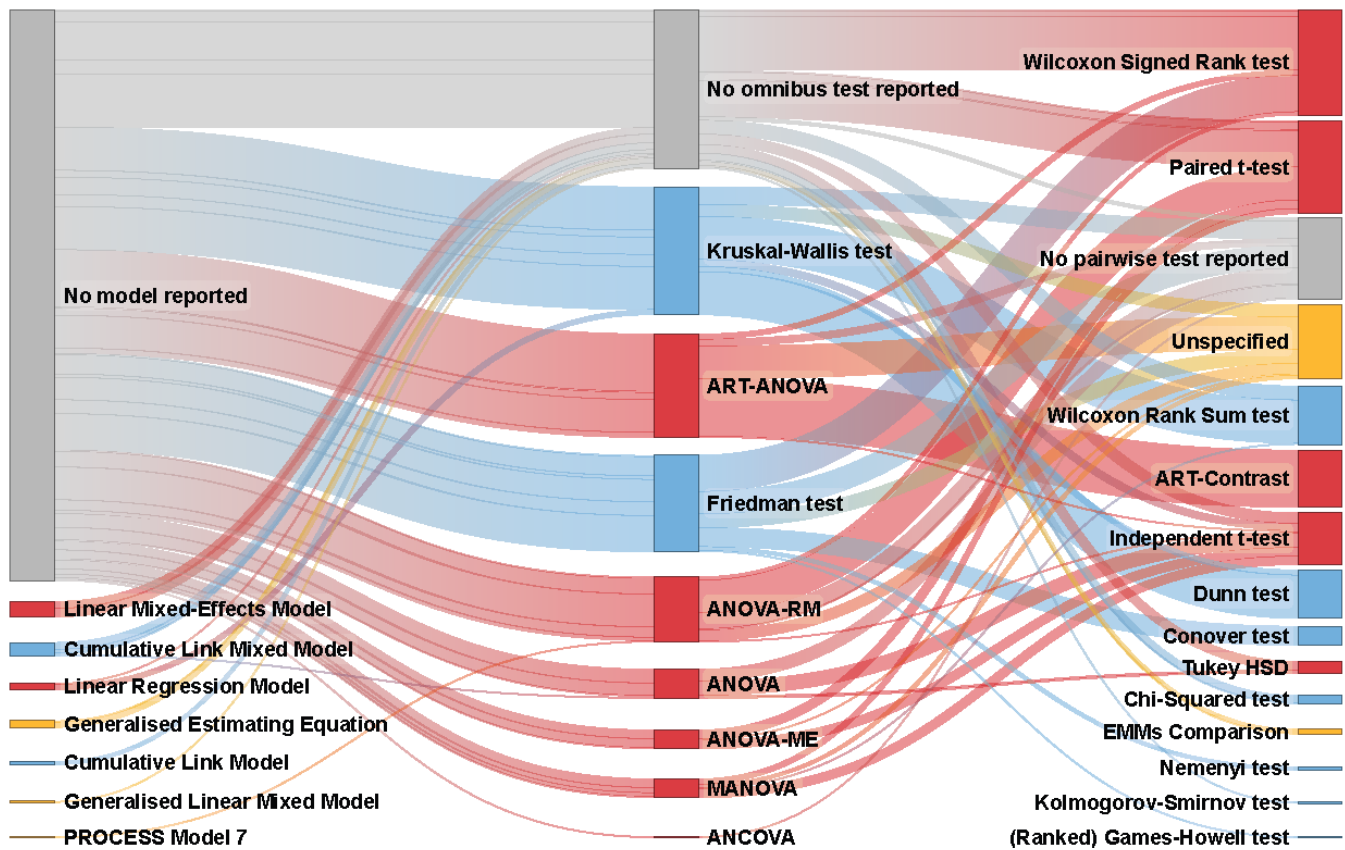


Figure 5: Sankey diagram showing the progression of statistical procedures applied to constructs measured with ordinal data. Methods are displayed in sequence from predictive models (left), to omnibus tests (centre), to pairwise tests (right). Bars labelled ‘No’ indicate constructs for which no method was reported in that category. Red nodes indicate statistical methods that impose incompatible assumptions with ordinal data, blue nodes represent ordinal data compatible methods, yellow nodes represent methods whose appropriateness for analysing ordinal data is dependent on specific factors (such as the link’ function used), and grey nodes represent absent methods in that specific category. As seen in the Figure, our sample reported limited use of predictive models (No model reported), and most models and omnibus approaches were not frequently used together (most models have links to ‘No omnibus test reported’). The Figure also clearly highlights a lack of consensus within the field; with no clear direction of which method, or sequence of methods, should be used to analyse ordinal data.

research. Both parametric and non-parametric methods continue to be frequently used, with non-parametric approaches outnumbering parametric approaches (362⁵ against 165⁶ respectively, excluding unspecified pairwise comparisons and predictive models; as models are dependent on the link function used). The increased use of non-parametric methods suggests a trend towards adopting approaches that impose fewer assumptions on ordinal data for analysis. However, we observed 154 cases of the use of the non-parametric approaches that assume data to be metric. These include 60 cases of the ART-ANOVA test, 33 cases of the post-hoc ART-Contrasts test, and 61 cases of the Wilcoxon signed rank test. This suggests that a

⁵60 ART-ANOVA + 70 Kruskal-Wallis tests + 56 Friedman tests + 33 ART-Contrast + 5 Chi-squared test + 28 Dunn test + 11 Conover test + 1 Games-Howell test + 1 Kolmogorov-Smirnov test + 2 Nemenyi test + 34 Wilcoxon Rank Sum test + 61 Wilcoxon Signed Rank test

⁶17 ANOVA + 37 ANOVA.RM + 6 ANOVA.ME + 1 ANCOVA + 10 MANOVA + 3 EMMEANS + 30 independent t-test + 54 paired t-test + 7 Tukey HSD

total of 319 out of 527 tests in our sample (excluding models and unspecified pairwise comparisons) impose metric assumptions on ordinal data. It is unclear if this is the result of widespread misinterpretation of the term ‘non-parametric’ to also *exclude metric assumptions on the data* or an informed *decision to treat ordinal data as metric*. In either case, this presents a challenge for future scholars in consistently interpreting, contrasting and building upon findings related to ordinal data that rely on different data assumptions. To better inform practitioners, we present the relevant assumptions for ordinal data analysis of the different hypothesis tests (omnibus and pairwise comparisons) found in our sample in Table 1. The table excludes predictive models as their assumptions are dependent on additional factors, such as the ‘link’ function used.

Inconsistent successive assumptions: Our findings further reveal that current HCI practices for analysing ordinal data may

Omnibus or Pairwise Method	Category	Experimental Design	Assumption		
			Metric Data	Normality of Residuals	Equal Variance
Wilcoxon Signed Rank test [39, 63]	Pairwise	Within-Subject Design	✓	✗	✗
Wilcoxon Rank Sum test [63]	Pairwise	Between-Subject Design	✗	✗	✗
ART-Contrasts [25]	Pairwise	Mixed Design	✓	✗	✗
Dunn test [22, 24]	Pairwise	Between-Subject Design	✗	✗	✗
Conover test [19, 20, 54]	Pairwise	Between/Within-Subject Design	✗	✗	✗
Chi-squared test [52]	Pairwise	Between-Subject Design	✗	✗	✗
Nemenyi test [48]	Pairwise	Within-Subject Design	✗	✗	✗
(Ranked) Games-Howell test [30, 57]	Pairwise	Between-Subject Design	✗	✗	✗
Kolmogorov-Smirnov test [4]	Pairwise	Between-Subject Design	✗	✗	✗
Kruskal-Wallis test [40]	Omnibus	Between-Subject Design	✗	✗	✗
ART-ANOVA [65]	Omnibus	Mixed Design	✓	✗	✗
Friedman test [28]	Omnibus	Within-Subject Design	✗	✗	✗
Sampled Parametric Methods	—	—	✓	✓	✓

Table 1: Table detailing common assumptions relevant to ordinal data analysis for omnibus and pairwise methods found in our sample. The table lists the category of the statistical method, the experimental design that generates data suitable for the method to analyse, and the assumptions that the method imposes. The table primarily details sampled non-parametric methods, as all parametric omnibus & pairwise tests in our sample (not all existing parametric methods) assume metric data, normality of residues, and equal variance — and are hence collectively represented by the last row. Notably, this table highlights frequently used non-parametric methods in HCI that exhibit incompatible assumptions on ordinal data. Specifically, the Wilcoxon Signed Rank test, the ART-Contrasts and the ART-ANOVA impose metric assumptions on the data (via the use of arithmetic operations on unranked data). Note that this table excludes approaches that we categorize as predictive models (e.g., CLMs or GLMs), as their assumptions depend on additional factors such as the ‘link’ function.

impose inconsistent data assumptions across the sequence of statistical procedures applied to ordinal data within the same study. Take, for example, the pairwise tests following a Friedman test (appropriate for repeated-measure non-parametric data) in our sample as depicted in Figure 5. This includes the Wilcoxon Signed Rank test, the Conover test, the Games-Howell test and the Nemenyi test, among additional unspecified tests. These tests are all appropriate for analysing non-parametric data generated through a within-subject experimental design. However, they differ in their treatment of the data, and may unintentionally impose assumptions incompatible with ordinal data that was initially avoided in the preceding omnibus test. For instance, the most frequently used test following the Friedman omnibus test in our sample was the *Wilcoxon Signed Rank test*, which assumes metric data, unlike its preceding omnibus test (see Table 1). This inconsistent treatment of ordinal data within individual sequences of analyses further complicates the interpretation of ordinal data analysis, which is already challenged by the parametric and non-parametric debate. Such inconsistencies warrant a change in the current statistical methods used for ordinal data analysis within HCI, prompting the need for methods that address the concerns stemming from over- or under-estimating the information present in ordinal data [11].

Ordinal first modelling: Specifically designed predictive modeling techniques, known as ordinal regression models (e.g., CL(M)Ms), provide greater statistical power than non-parametric approaches while avoiding the incompatible assumptions of commonly used parametric methods in HCI, such as ANOVAs. They achieve this by leveraging the ordinal structure of the data without assuming equally spaced intervals between categories. However,

such models see limited application in HCI (see Figure 2). A possible explanation for such limited use, is the focus on the use of ordinal data for hypothesis testing in recent HCI work, as corroborated by the disparity between the frequency of hypothesis tests (omnibus and pairwise tests) and predictive modelling methods in our sample (Figures 3, 4, and 2). This could further be explained by previous challenges in understanding, interpreting and applying predictive modelling methods in HCI research [65]. However, we argue that recent developments in statistical analysis software have made these methods increasingly more accessible and easy to apply. To better inform HCI practitioners in understanding and using the methods, the following sections describe the proposed CL(M)M modelling technique, and presents worked examples using CL(M)Ms to analyse ordinal data in published open source HCI data. We additionally provide an interactive tool in our supplementary material to illustrate a simple CLM model.

4 Cumulative Link (Mixed) Models

Cumulative Link (Mixed) Models, or CL(M)Ms, are a class of Generalised Linear (Mixed) Models that are designed for ordinal regression. Throughout this paper, we will often use the term ‘CL(M)M’ to refer to both cumulative link models (CLMs) and cumulative link mixed models (CLMMs). Specifically, CLMs are appropriate when only fixed effects are of interest and observations can be assumed to be independent. In contrast, CLMMs extend CLMs by incorporating random effects, allowing them to account for grouped or hierarchical data in which the independence assumption is violated (e.g., repeated-measures experimental designs). These models appropriately treat ordinal data as *categorical*, while *exploiting the*

ordered nature of the data [17]. Given such advantages, these methods are increasingly encouraged and used for modelling ordinal data in diverse fields of research, including psychology [11, 29, 43], econometrics [38], machine learning [64], and health [15]. Despite the advantages of CL(M)Ms, they have been largely overlooked in HCI, due to challenges in understanding, interpreting and applying these methods [65]. As with all statistical methods, CL(M)Ms come with their own set of assumptions that are important to understand and consider prior to their use. With the increasing availability of software packages capable of performing ordinal regression using CL(M)Ms [2, 17] in popular statistical tools, such as Stata, R, and Python, this section provides an intuitive description of the theory and assumptions of CL(M)Ms in the context of HCI research. Readers already familiar with CL(M)M theory and assumptions, or those interested primarily in practical applications of CL(M)Ms for HCI data, may skip to Section 5.

4.1 Intuition

Say that we are interested in the effect of adding a new feature to our system (e.g. faster screen frame rate) on an element of user experience (e.g. spatial presence in virtual reality). In an ideal scenario, we would have a tool that could directly measure a user's sense of spatial presence, giving us a continuous and normally distributed measure that we could analyse using a familiar method, such as an ANOVA. Figure 6 illustrates this ideal scenario.




Figure 6: Ideally, we could directly measure a user's sense of spatial presence from their thoughts (LEFT). If the measure is continuous and normally distributed (RIGHT), it would enable us to use familiar metric models of statistical analysis and comparisons.

However, constructs such as spatial presence are not directly observable, and consequently, require us to use proxy measures. Often, Likert-item styled questionnaires are used, where users provide a rating of spatial presence on an ordinal scale. Unlike metric data, ordinal data do not exhibit equal distances between ratings, i.e., the ordinal values of $3 - 2$ and $5 - 4$ may not be equal, even if they seem to arithmetically equate to 1 due to the use of contiguous numerical labels. As such, ordinal values cannot be meaningfully analysed using methods that rely on arithmetic operations. In the context of our example, this means that we cannot estimate how much more spatially present users were (on average) when using our system with, and without, the new feature by subtracting the average ordinal responses for the different system conditions. Figure 7 illustrates these differences between the continuous representation in our ideal case (left), and the ordinal measure collected in a real-world context (right). In the ideal scenario, users' spatial presence

scores with and without the added feature form two normal distributions, and the average effect of the new feature is reflected by the shift from the original distribution. In reality, however, we can only collect user spatial presence measures via ordinal ratings on a predefined scale (say from 1 to 5), resulting in distributions of ordinal data (right) that is far from normal, and whose effect is much harder to estimate.




Figure 7: Figure depicting an ideal scenario and our reality when deriving insights from ordinal data related to our system. In an ideal scenario (LEFT), we would not need ordinal measures, and we would have direct and continuous measures of users' spatial presence when using the two versions of our system. This would result in two distinct normal distributions of users' spatial presence when using our system with and without the new feature. In reality (RIGHT), spatial presence cannot be measured directly and instead a discrete score (e.g., between 1 and 5) is provided by the user to indicate their level of spatial presence under the two system variations.

CLMMs attempt to address this challenge by taking a principled approach to analysing ordinal data. These models assume that there exists a continuous *unobservable* measure (or latent variable) in the user's mind that underlies an observable ordinal measure. Since this measure is unobservable, the unit of measurement can be arbitrary. Conventionally, CL(M)Ms assume that this latent variable is distributed with a standard normal (0,1) distribution (though models can be extended into any arbitrary latent distribution). If indeed the observed scores reflect these assumed underlying latent values, then they should be statistically related. However, CL(M)Ms do not assume this relationship to be linear, only monotonic (i.e. as the latent value increases, so does the likelihood of a higher user rating). Imagine a user who initially experiences no spatial presence (i.e. their latent score should be the minimum imaginable). This user would most likely give the lowest possible rating (1 - Very low). As we intervene to increase the latent score from its absolute minimum, we would still see scores of 1 until we reach a certain level, at which point, we would start seeing scores of 2. As we keep increasing the latent score and reaching these thresholds/cutpoints, our ratings should increase accordingly.

Because the latent normal and the observed scores distributions are connected under these assumptions, we can also assume that the proportion of users with a latent value between any 2 cutpoints should be equal to the proportion of users that gave the corresponding rating. As such, we can compare the cumulative distribution of the scores to the cumulative distribution of the latent normal. In other words, if 10% of our scores were 1, and we know that in




Figure 8: Our ordinal rating increases when the latent value increases beyond a cutpoint. By looking at which cutpoints the latent value falls between, we can predict which score the user would give.

a standard normal distribution 10% of values are lower than -1.28 , then we can say that -1.28 is the cutpoint in the latent scale at which someone’s rating would go from 1 to 2. If in the same dataset we found that 15% of scores were 2, we would have to look for the $15 + 10 = 25\%$ quantile of the normal distribution to find the next cutpoint—in this case, -0.67 . We can proceed the same way for the other cutpoints. Figure 8 visualizes the selection of ordinal scores based on the cutpoints that divide the latent measure scale.

This method gives us a way to retrieve the latent cutpoints from observed scores. However, ultimately, we are interested in the causal effect of our intervention. For this example, we will assume that the intervention only shifts the mean of our latent distribution (not the variance). Under the same cutpoints, shifting the mean of the latent measure causes the area under the normal curve between the cutpoints to change—and these areas are precisely the proportion of each of the possible ratings.

Figure 9 shows this process. We infer the cutpoints from the baseline ratings, shift the mean of the normal distribution to the left (which implies an average decrease in the latent value) and see how it affects the distribution of scores. We offer a Shiny app in the supplementary material where you can set the cutpoints and observe how different effect sizes affect the distribution of scores.

So far, we described the intuition of the method generatively, that is, from the user’s mind to the scores. In essence, CL(M)Ms do the inverse of this process: they take ordinal ratings and directly estimate the effects on the latent scale. Our example uses the cumulative proportion analogies to build an intuition. In practice, however, the cutpoints are not simply derived from the cumulative proportions of observed scores, but jointly estimated from the data along with the model coefficients. Next, we provide a more formal mathematical description of how it works and a practical implementation in R.

4.2 Theory and Mathematics

As a predictive model, CL(M)Ms model the relationship between a dependent variable and predictor/independent variable(s). Let us extend the example used in section 4.1 to illustrate this. Say we ask




Figure 9: From a given distribution of baseline scores, we can infer the latent cutpoints and predict how shifting the mean of the distribution would affect the distribution of scores accordingly.

participants to use different versions of our system (a categorical independent variable represented as X) and rate their sense of spatial presence between 1 and 5, with 1 being ‘Very low’ and 5 being ‘Very high’ (an ordinal dependent variable represented as Y). We can represent this relationship as a causal graph, $X \rightarrow Y$, where the arrow pointing from X (system version) to Y (spatial presence response) indicates that X has a causal effect on Y . In this scenario, the goal of predictive modelling would be to describe how changing the system version X influences user responses about their sense of spatial presence Y , i.e., how X predicts Y . However, the ordinal response, Y , is an *imperfect measure* that the user generates about their sense of spatial presence in using a system version, based on some internal psychological measure and the available ordinal responses (elaborated in the example presented in figure 7). For example, a user’s true spatial presence experience may lie somewhere between




Figure 10: Figure illustrating the relationship between the internal continuous latent variable \tilde{Y} with the observable ordinal response Y in our spatial presence example. Here, a user determines the system to elicit ‘Very low’ spatial presence when their internal measure falls below the first cutpoint (τ_1), ‘Low’ spatial presence, when their internal measure is between the first (τ_1) and second (τ_2) cutpoints, and so on.

‘High’ (scale value 4) and ‘Very High’ (scale value 5), but the limited response options force a decision that imperfectly captures the user’s true experience. Additionally, there is no guarantee that the internal psychological distance between two ordinal responses is evenly spaced. For instance, users may find it easy to change their response from 3 (‘Moderate’) to 4 (‘High’) – small psychological distance – but may reserve the highest response of 5 (‘Very high’) to systems that they perceive as exceptional at evoking a sense of spatial presence – large psychological distance between 4 and 5 when compared to 3 and 4. This prevents meaningful analysis that relies on these distances to be consistent, such as statistical analysis methods that calculate averages or differences, and is a key cause of errors in analysing ordinal data using metric models [11, 43].

CL(M)Ms account for the internal psychological measure by assuming the presence of an underlying continuous latent (unobservable) variable (represented as \tilde{Y}). This results in an updated causal graph, $X \rightarrow \tilde{Y} \rightarrow Y$, which explicitly represents how the system version (X) affects an internal psychological measure of the user’s sense of spatial presence (\tilde{Y}), which is then categorized into an ordinal response (Y). Given this new causal structure, the relationship between the observable response, Y , and the latent variable, \tilde{Y} , can then be defined using a set of **ordered thresholds/cutpoints, which determines the value that Y categorizes to, when \tilde{Y} falls within certain cutpoints**. In our example, the 5 different responses (‘Very low’ to ‘Very high’) will need 4 cutpoints (represented as $\tau = \{\tau_1, \tau_2, \tau_3, \tau_4\}$) to categorize all 5 available ordinal response, i.e. $Y = 1$ (‘Very low’) when $\tilde{Y} \leq \tau_1$, $Y = 2$ (‘Low’) when $\tau_1 < \tilde{Y} \leq \tau_2$, and so on, with the highest value, $Y = 5$ (‘Very high’) manifesting when $\tau_4 < \tilde{Y}$. This process is illustrated in Figure 10 (and with a familiar Likert item example in figure 8). In general, if we have K different ordinal responses, we will need $K - 1$ cutpoints. Formally:

$$Y = k \text{ if } \tau_{k-1} < \tilde{Y} \leq \tau_k \quad (1)$$

This conceptualization of the latent *continuous* variable, \tilde{Y} , enables CL(M)Ms to mathematically express Y in terms of \tilde{Y} and τ .

CL(M)Ms can then model the continuous latent variable \tilde{Y} in relation to X using linear regression. Collectively, these two steps enable CL(M)Ms to describe the relationships represented in the causal graph: $X \rightarrow \tilde{Y} \rightarrow Y$.

To express Y in terms of \tilde{Y} and τ , CL(M)Ms must first ensure that the observable ordinal data (Y) is appropriately treated as *categorical*, avoiding over-representation of information (as with metric modelling approaches) [11]. However, treating ordinal data as categorical exposes only frequencies/probabilities as quantitative information, which are devoid of any ordering (under-representing available information as with most non-parametric approaches [11]). For example, the probability set $\{0.2, 0.1, 0.7\}$ corresponding to an ordered dataset with levels $\{A, B, C\}$ provides no information about the ordering of the data. To remedy this, CL(M)Ms make use of *cumulative probabilities* in place of individual probabilities. This ensures that higher ordered categories are always associated with higher cumulative probabilities and vice-versa. Extending our example, the ordered dataset with levels $\{A, B, C\}$ and probabilities $\{0.2, 0.1, 0.7\}$, will exhibit cumulative probabilities of $\{0.2, 0.3, 1.0\}$ ⁷, preserving order. This also enables us to extract individual probabilities from cumulative probabilities using:

$$P(Y = k) = P(Y \leq k) - P(Y \leq k - 1) \quad (2)$$

Now, from (1), we can infer that $Y \leq k$ when $\tilde{Y} \leq \tau_k$. Therefore, we can relate the cumulative probability $P(Y \leq k)$ with the cumulative probability $P(\tilde{Y} \leq \tau_k)$ which can be calculated using a cumulative distribution function over \tilde{Y} , represented as $F(\tau_k)$. Here, the function F depends on the assumed distribution of \tilde{Y} . Two commonly used distributions include the standard logistic distribution and the standard normal distributions, which have been referred to by the names ordered logit models and ordered probit models respectively [32]. The choice of assumed distribution influences the interpretation of the results but often results in similar parameter estimates and model fits [46]. For the purposes of this paper, we assume that \tilde{Y} follows a standard normal distribution. Details related to the different distributions and their corresponding cumulative distribution functions can be found in prior work [1, 11]. Given the function F , we can now express the relationship between Y and \tilde{Y} as:

$$P(Y \leq k) = P(\tilde{Y} \leq \tau_k) = F(\tau_k) \quad (3)$$

From equation (2) and (3), we can determine individual probabilities of ordinal responses based on the cumulative distribution function (F) and the cutpoints (τ):

$$P(Y = k) = F(\tau_k) - F(\tau_{k-1}) \quad (4)$$

However, equation (4) does not explain the influence of system version (X) on spatial presence response (Y). To account for this, CL(M)Ms use linear regression to model the continuous latent variable \tilde{Y} with respect to X ⁸:

$$\tilde{Y} = X^T \beta + \epsilon \quad (5)$$

Where the coefficient β in the linear predictor ($X^T \beta$) describes the change in \tilde{Y} with respect to unit change in X , and ϵ is the

⁷Each value in this array is obtained by adding up all preceding values in the original array.

⁸While we illustrate using a single independent variable, CL(M)Ms can handle multiple independent variables, where X would represent a matrix of column vectors.




Figure 11: Figure showing example distributions of \tilde{Y} for $X = \text{Version}_1$ ($\tilde{Y} = X_1^T \beta_1 + \epsilon$) and for $X = \text{Version}_2$ ($\tilde{Y} = X_2^T \beta_2 + \epsilon$). The figure presents an example where users are more inclined to respond with a higher spatial presence rating to Version_2 compared to Version_1 . This is depicted by the larger proportion of the curve (area under the curve) for Version_2 that falls between cutpoints (τ) that define higher ordinal responses, when compared to the curve for Version_1 . For example, area under the curve for $\tau_4 < \tilde{Y}$, defining the probability of a ‘Very high’ user response, is larger for Version_2 than for Version_1

random error describing the assumed distribution of \tilde{Y} (standard normal in this example). To elaborate, let us assume that our spatial presence study example tests two system versions ($X_1 = \text{Version 1}$ and $X_2 = \text{Version 2}$), with the first version having slower screen frame rate than the second version. This can be represented in our linear model of \tilde{Y} as two binary (0 or 1) variables X_1 and X_2 , indicating the presence or absence of that system version. Therefore, for this example, we can specify equation (5) as:

$$\tilde{Y} = X_1^T \beta_1 + X_2^T \beta_2 + \epsilon \quad (6)$$

Since a user can only experience one version at a time in our experiment that compares the two versions, either $X_1^T \beta_1$ or $X_2^T \beta_2$ will be equal to zero. Therefore, our linear predictor for \tilde{Y} can be described as:

$$\tilde{Y} = \begin{cases} X_1^T \beta_1 + \epsilon, & \text{if } X = \text{Version}_1 \\ X_2^T \beta_2 + \epsilon, & \text{if } X = \text{Version}_2 \end{cases}$$

If Version 1 and 2 yield different effects on the latent spatial presence variable \tilde{Y} (i.e., $\beta_1 \neq \beta_2$), then \tilde{Y} will exhibit different *distribution positions* for different system versions. As such, we can estimate the causal effect of switching from Version 1 to Version 2 by subtracting their respective means. Figure 11 illustrates the distribution of the latent spatial presence variable \tilde{Y} when users respond with a higher spatial presence rating for Version 2.

Collectively, we now have equation (3) that describes the relationship between the user’s response (Y) and their internal spatial presence measure (\tilde{Y}), and equation (5) that models a linear relationship between the internal measure (\tilde{Y}) and the system version (X). Therefore, we can model the probability of a user response (Y) being less than a specified value k given the linear predictor ($X^T \beta$)

of \tilde{Y} as:

$$\begin{aligned} P(Y \leq k | X^T \beta) &= P(\tilde{Y} \leq \tau_k | X^T \beta) = P(X^T \beta + \epsilon \leq \tau_k) \\ &= P(\epsilon \leq \tau_k - X^T \beta) = F(\tau_k - X^T \beta) \end{aligned} \quad (7)$$

In practice, equation (7) is expressed in terms of the inverse cumulative distribution function (F^{-1}) which is the ‘link’ function in generalised linear models to connect the linear predictor with the response value (see Agresti and Kateri [1] for a detailed description of these functions). Given (7), we can update equation (4) to consider the linear predictor ($X^T \beta$) of \tilde{Y} to express the probability of Y at individual responses as:

$$P(Y = k | X^T \beta) = F(\tau_k - X^T \beta) - F(\tau_{k-1} - X^T \beta) \quad (8)$$

This gives us the final equation that describes the relationship between the system version (X) and the observable spatial presence user rating (Y).

4.3 Assumptions

Like any statistical analysis method, CL(M)Ms impose a set of assumptions that should be carefully considered prior to their application. While we briefly mention these assumptions in sections 4.1 & 4.2, we elaborate on them here for further consideration. Additionally, tools to test for, and relax, some of the discussed assumptions are available with most software implementations for ordinal data analysis (see examples in R [17] and Stan [11]). If readers are not inclined towards making certain assumptions, and if these assumptions cannot be relaxed, they can refer to our table 1 for support in selecting alternative statistical tests to analyse their ordinal data.

Latent Variable Assumption. A fundamental assumption of CL(M)Ms relevant to HCI is that there exists an internal psychological measure (or latent variable) with an assumed distribution (standard normal in our ‘probit link’ examples) that dictates, based on cutpoints, the ordinal response of a user to a specific survey question or Likert-item. The assumption of underlying latent variables is prominent and relevant in many fields, and particularly in psychology and the social sciences [5] where studies often focus on phenomena with directly unobservable causes. However, it is still worth considering whether such an assumption aligns with the researchers’/analysts’ beliefs prior to the use of CL(M)Ms (or similar) for their analysis.

Latent Equal-Variance Assumption. Heretofore, we have described CL(M)Ms using simplified examples that only illustrate changes in the location (distributional position) of the underlying latent variables \tilde{Y} (see figure 7, 9, & 11). While standard CL(M)Ms assume constant variance of \tilde{Y} , they can be extended to account for unequal variances by modelling the scale (standard deviation, σ) of \tilde{Y} within the model [17]. Practically, σ can be incorporated through an additional regression factor in the model, often expressed as its inverse [11] ($\alpha = 1/\sigma$). Accordingly, equation (7) takes the form: $P(Y \leq k | X^T \beta, \alpha) = F(\alpha \times (\tau_k - X^T \beta))$. Where, $\alpha = \exp(X^T \beta_\alpha)$ to ensure $\alpha > 0$, and $X^T \beta_\alpha$ is the linear predictor for α . Modelling the scale as an additional regression factor enables the variance of \tilde{Y} to also be dependent on the variables X . More in-depth details can be found in Bürkner and Vuorre [11].

Proportional Odds Assumption. Another important assumption of CL(M)Ms is dubbed the *proportional odds assumption*. This

assumption refers to the unchanging coefficient estimates (β) across all cutpoint estimates ($\tau = \{\tau_1, \tau_2, \dots, \tau_{k-1}\}$), i.e., the coefficients (β) remain the same $\forall k = \{1, 2, \dots, K - 1\}$ in $P(Y = k|X^T\beta)$ in equation (8). *Equivalently, this assumes that the cutpoints do not individually vary based on the independent variables (X).* Effectively, this suggests that an intervention shifts user latent scores by the same amount across all cutpoints, and therefore has an equal effect on the likelihood of moving user responses from 1 to 2, 2 to 3, and so on. However, cases exist where an intervention's effect cannot be completely captured by a constant shift in the location of the original distribution, but also by how the distribution is 'stretched' or 'compressed' at different points. One way to test if this assumption holds is by checking for the equality of the scale σ of the latent variable \tilde{Y} under different conditions, provided the model accounts for unequal variances, i.e., if the conditions only affect the location of the distribution, then the proportional odds assumption should hold. Statistical tests, such as the Brant test [7] can also be used to verify if this assumption holds. This assumption can also be partially relaxed [17, 27] using the partial proportional odds model extension, which allows a subset of the independent variables to not assume proportional odds [53]. Here, an additional regression element ($P^T\gamma_k$) can then be incorporated into the model, where P is a subset of the independent variables X , and γ_k are the corresponding regression coefficients associated with P for **the specific ordered category k** . The complete mathematical description around (partial) proportional odds is beyond the scope of this paper, and readers are referred to Peterson and Harrell Jr [53] for details.

While detailed mathematical knowledge of CL(M)Ms is not required for their use, a fundamental understanding of the latent model and its relationship to the ordinal data is essential for interpreting results and assessing assumptions. The following section moves from the theoretical aspects of CL(M)Ms to practical examples of analysing ordinal data using these models.

5 Analysing Ordinal Data with CL(M)Ms in R

While historically less accessible [11, 65], several software packages now exist that enable ordinal regression modelling in popular statistical tools used in HCI. These include the `statsmodel` package for Python [56], the `ologit` and `oprobit` functions for Stata ⁹, and the `repolr` [51] and `brms` package in R [10], among others. These packages provide different interfaces for fitting CL(M)Ms and may format model outputs differently. However, the underlying concepts (see Section 4) behind CL(M)Ms remain unchanged.

For the purposes of this paper, we demonstrate the practical use and interpretation of CL(M)Ms in HCI using R's ordinal package [18]. Specifically, we present two examples, using CLMs and CLMMs respectively, to reanalyse open-sourced data provided with previous CHI papers by Fitton et al. [26] and Chen et al. [16]. These papers were sampled in our initial search of the literature (see Section 3), but were excluded in our review during random sampling (Figure 1). Both examples make use of CL(M)Ms with a 'probit' link to enable simpler result interpretation on the standard normal latent scale, as opposed to a log-odds interpretation when using the 'logit' link. As mentioned in section 4.2 and in prior work [46], both link functions return similar results. Furthermore, for simplicity,

our examples will assume equal variances and proportional odds. However, these assumptions can be relaxed when using the 'ordinal' package by specifying the 'scale' and 'nominal' function variables respectively. Specific details for available functions in the ordinal package can be found in Christensen [17]. Finally, we present only a subset of the analysis plots in this section for illustrative purposes. Additional plots can be found in our Appendix A.1.

Pre-processing and analysis replication: For each paper, we first load the available data into an R data frame and organise the data into a long-table format. This format eases use with R *formulae* structure used in the 'ordinal::clm()/clmm()' function. We then apply the analysis methods reported in the original paper to ensure that our pre-processing preserves data integrity and reproduces the findings reported. If test-statistics and p-values are provided in the paper, we ensure that these values are replicated in addition to alignment with significant/non-significant reports. For brevity, we do not reiterate the findings of the original analysis reported in the paper. Readers can refer to the original article or the analysis scripts provided in our supplementary material for additional details.

5.1 CLM Example: Fitton et al. [26]

Study Design & Aim: The paper by Fitton et al. [26] explores how observing and performing a psychomotor task in virtual reality (VR) affects learning and retention. Four *learning conditions* were evaluated: active learning (ACTIVE), observational learning with a self-avatar (SELF), observational learning with a minimally similar avatar (MINIMAL), and observational learning with a dissimilar avatar (DISSIMILAR). These conditions were evaluated **between-subjects** using a puzzle assembly task. Immediately following the learning conditions in VR, participants were assessed in completing the puzzle in VR (immediate VR retention task) and in two real-world transfer assessments with 3D printed puzzle pieces (immediate near transfer task, immediate far transfer task). The near transfer task mirrored the learning condition in VR and presented the printed puzzle pieces as they appeared in VR (same colour coded pieces and initial relative piece pose), while the far transfer task differed from the learning condition by presenting non-coded puzzle pieces in different initial relative poses and included a second task of counting audio tones presented at random intervals. These tasks were performed by all participants, i.e., a *within-subject* factor. The participants then returned after 10-14 days to complete the delayed VR retention task, delayed near transfer task, and the delayed far transfer task (*within-subject*).

Outcome measures: The paper presents several recorded outcome measures, including task completion times, number of pieces correctly assembled ¹⁰, and questionnaire responses (Vividness of Imagery, Intrinsic Motivation Inventory, and Self-efficacy scores, among others). We focus on a subset of the ordinal responses analysed in the paper for illustrating the use of CLMs in HCI. Specifically, we re-analyse the:

- (1) Usefulness of training (Perceived Usefulness), 1 Likert-type item.

⁹<https://www.stata.com/features/overview/logistic-regression/>

¹⁰This was analysed using an ordinal regression model in the paper, but other ordinal data included in the paper were not.

- (2) Simulated Task Load Index (SIM-TLX) survey, 9-item Likert-scale.

This subset of measures were chosen as the raw, unprocessed data is provided, i.e., not averaged or weighted. Both measures were recorded **between-subjects** which enables analysis using CLMs.

Analysis used in the paper: Fitton et al. [26] report a series of one-way ANOVAs to analyse the effect of learning condition on the ‘Perceived Usefulness’ measure and on individual sub-scales of the SIM-TLX survey. The paper reports no significant differences between learning conditions on the ‘Perceived Usefulness’ measure and across all SIM-TLX sub-scales. We were able to reproduce these results in our analysis by following the methods reported in the paper.

Modelling ‘Perceived Usefulness’ using CLMs: We first organise data related to the ‘Perceived Usefulness’ measure (USEFULNESS) and the corresponding learning condition (CONDITION) into an R data frame. The first 5 rows of the data frame is presented using the ‘head()’ function:

```
head(usefulness_dataframe, 5)
  Condition Usefulness
  <fct>      <ord>
1 Dissimilar 5
2 Self       4
3 Dissimilar 2
4 Active     5
5 Self       5
```

We can then fit a CLM with a probit link to the data, and display an overview of the model results by using the ‘clm()’ function followed by the ‘summary()’ function:

```
usefulness.clm <- clm(
  Usefulness ~ 1 + Condition,
  data = usefulness_dataframe,
  link = "probit")

summary(usefulness.clm)

formula: Usefulness ~ 1 + Condition
data:    usefulness_dataframe

link threshold nobs logLik AIC niter max.grad cond.H
probit flexible 102 -113.99 241.97 4(0) 1.47e-08 5.0e+01

Coefficients:
              Estimate Std. Error z value Pr(>|z|)
ConditionDissimilar -0.32161    0.32288  -0.996  0.319
ConditionSelf       0.04704    0.33048   0.142  0.887
ConditionMinimal   -0.49092    0.31466  -1.560  0.119

Threshold coefficients:
Estimate Std. Error z value
1|2 -2.5897    0.4409 -5.874
2|3 -1.8801    0.2965 -6.340
3|4 -1.1437    0.2519 -4.540
4|5 -0.2450    0.2358 -1.039
```

Table 2: Frequency table of user responses (R) to the ‘Perceived Usefulness’ survey in each learning condition (C) of Fitton et al.’s [26] paper.

C \ R	R				
	Y = 1	Y = 2	Y = 3	Y = 4	Y = 5
Active	0	1	3	6	16
Self	0	0	3	7	15
Minimal	1	0	4	12	9
Dissimilar	0	3	3	6	13

The ‘Coefficients’ table and the ‘Threshold Coefficients’ table present the primary results of model fitting.

Interpreting Results: The ‘Coefficients’ table presents the estimated change (β parameters) in the user’s internal psychological measure of usefulness (latent variable, \tilde{Y}) when the learning condition is changed from a reference condition. In this case, we set our reference condition as ‘CONDITION = ACTIVE’, which, as described in Section 4, has a standard normal distribution, $N(\mu = 0, \sigma^2 = 1)$. Therefore, the estimates can be interpreted as follows. The location (mean = μ) of the user’s internal measure of usefulness (\tilde{Y}) distribution is shifted by -0.322 ($p = 0.319$) standard units (i.e. standard deviations) when changing the learning condition from ACTIVE to DISSIMILAR, $+0.047$ ($p = 0.887$) units from ACTIVE to SELF, and -0.491 ($p = 0.119$) units from ACTIVE to MINIMAL. When assuming equal variances of the latent variables, this has a similar interpretation to Cohen’s d coefficient. Note that the summary presents p-values ≥ 0.5 for each of the estimates (β), indicating that the estimates related to each of the non-reference learning conditions (DISSIMILAR, SELF, MINIMAL) do not significantly differ from the reference condition (ACTIVE), i.e., these p-values are significance tests for the parameter estimates being zero (reference).

The ‘Threshold Coefficients’ table presents the cut-point/threshold values (τ) relative to \tilde{Y} at the reference condition (CONDITION = ACTIVE). These cutpoints determine the observable ordinal response measure of perceived usefulness (Y) as described by equation (1). Specifically, user’s will respond with a ‘Perceived Usefulness’ (Y) measure of $Y = 1$ when their internal measure of usefulness (\tilde{Y}) is -2.590 units below the standard normal mean (i.e., zero), $Y = 2$ when $-2.590 < \tilde{Y} \leq -1.880$, $Y = 3$ when $-1.880 < \tilde{Y} \leq -1.144$, $Y = 4$ when $-1.144 < \tilde{Y} \leq -0.245$, and $Y = 5$ when $-0.245 < \tilde{Y}$. These estimates suggest that users are more likely to respond with $Y = 5$ for CONDITION = ACTIVE given that the *average* value of the internal psychological measure of usefulness (0 by assumption of standard normal distribution) is greater than the last cutpoint ($\tau_4 = -0.245$). The frequency table (Table 2) for the responses based on the different learning conditions provides further support for this result – indicating that the most common response for ‘Perceived Usefulness’ in the ACTIVE condition is $Y = 5$ (with 16/26 responses, or 61.5% of all responses).

To inspect the other learning conditions, we shift the location of the \tilde{Y} distribution from the standard normal (location at zero) by the coefficient estimate corresponding to the learning condition of

interest, as presented in the ‘Coefficients table’, i.e., $\tilde{Y}_{Condition} \sim N(0 + Estimate_{Condition}, 1)$ ¹¹. For instance, if we were interested in $CONDITION = MINIMAL$, we would shift the location of \tilde{Y} from the standard normal location ($\mu = 0$) by -0.491 units ($\mu = 0 - 0.491 = -0.491$). This suggests that for $CONDITION = MINIMAL$, user’s will respond with $Y = 1$ when $\tilde{Y} - 0.491 \leq -2.590$, $Y = 2$ when $-2.590 < \tilde{Y} - 0.491 \leq -1.880$, $Y = 3$ when $-1.880 < \tilde{Y} - 0.491 \leq -1.144$, $Y = 4$ when $-1.144 < \tilde{Y} - 0.491 \leq -0.245$, and $Y = 5$ when $-0.245 < \tilde{Y} - 0.491$. Note that this shifts the mean of \tilde{Y} for $CONDITION = MINIMAL$ between cutpoints 3|4 (τ_3) and 4|5 (τ_4), suggesting that the most likely response, given a normal distribution, is 4 (as $\tau_3 < mean(\tilde{Y}_{Minimal}) \leq \tau_4 \Rightarrow -1.144 < -0.491 \leq -0.245$). This is also supported by Table 2, which indicates that the most frequent response for $CONDITION = MINIMAL$ is $Y = 4$ with 12/26 of the total responses (46.2%).

Modelling SIM-TLX sub-scales using CLMs: As we have familiarised ourselves with CLM fitting and result interpretation in the previous sections, this example highlights differences in insights that may arise between statistical approaches traditionally used in HCI and CLMs. Specifically, we fit a CLM to each individual sub-scale of the SIM-TLX data provided by Fitton et al. [26] to mirror their analysis using one-way ANOVAs. Contrary to their conclusion of no significant effects of learning condition on any of the SIM-TLX sub-scales, our analysis using CLMs suggests that parameter estimates for the sub-scales related to ‘Physical demands’ and ‘Stress’ shows significant divergence from zero for $CONDITION: SELF$, i.e., significant differences between learning condition $SELF$ and the reference condition $ACTIVE$. We present the ‘Coefficients’ table returned by the ‘summary()’ functions for the CLMs fit to the ‘Physical demands’ and ‘Stress’ sub-scale data. The ‘Threshold Coefficients’ table for the SIM-TLX measures are not presented here as details related to threshold/cutpoints interpretation is extensively discussed in the previous example. However, these details can be reproduced using our R scripts provided in the supplementary materials.

```
# SIM-TLX: Physical Demands
Coefficients:
      Estimate Std. Error z value Pr(>|z|)
ConditionDissimilar -0.5196  0.2845  -1.826  0.0678 .
ConditionSelf       -0.6751  0.2884  -2.341  0.0192 *
ConditionMinimal   -0.1907  0.2799  -0.681  0.4958

# SIM-TLX: Stress
Coefficients:
      Estimate Std. Error z value Pr(>|z|)
ConditionDissimilar  0.09148  0.28315  0.323  0.7466
ConditionSelf       0.59792  0.28613  2.090  0.0366 *
ConditionMinimal    0.43877  0.28198  1.556  0.1197
```

These findings suggest that the distribution for the user’s internal measure of ‘Physical Demand’ and ‘Stress’ for the condition $SELF$ is significantly different from $ACTIVE$. These findings differ

¹¹Note that the same results can be achieved by shifting the thresholds/cutpoints (τ) by the relevant estimate, instead of the location of \tilde{Y} , as per equation (7)

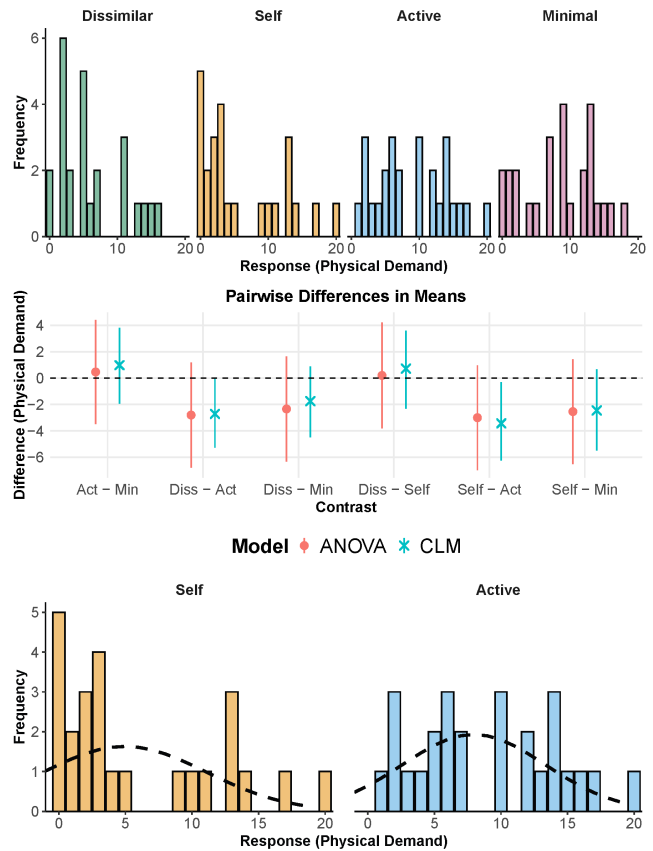


Figure 12: Visualizations relevant to Fitton et al.’s [26] paper. (Top) Frequency of user responses on the ‘SIM-TLX: Physical Demand’ subscale across learning conditions. (Middle) Confidence intervals for pairwise differences in mean responses between conditions under ANOVA and CLM models. (Bottom) Frequency plots for ‘Physical Demand’ in the $SELF$ and $ACTIVE$ conditions, which differed significantly under CLMs. Dashed curves show fitted linear (metric) model distributions based on condition means and standard deviations, highlighting how metric models can miss skewed responses. For example, $SELF$ responses are mostly minimal (0–3), but the metric model mean is inflated by a spike at 13.

from the original analysis reported by Fitton et al. [26]. To better illustrate these differences, figure 12 presents relevant plots related to the analysis of ‘Physical Demand’. Specifically, figure 12 presents the response frequency histograms, pairwise mean difference confidence intervals (CIs) under the ANOVA and CLM analysis, and a side-by-side comparison of the distributions (discrete and linear model fit) for ‘Physical Demand’ responses under conditions where ANOVA and CLM analysis disagreed ($SELF$ and $ACTIVE$). Note that CIs for both models were obtained using the emmeans package in R [42]. Additionally, as CL(M)Ms report estimates on the latent scale, a bootstrapping approach [21] was employed to obtain CIs on the response scale for the CLM model to allow direct comparison with the ANOVA-based CIs in this example. Differences in CIs between

the models suggest only a minor, but observable, disagreement. The side-by-side comparison, shown in the last subfigure, illustrates how metric model fits (based on the data's mean and standard deviations) can appear similar even when ordinal responses differ. For example, SELF responses cluster at 0–3, but the metric model mean is skewed upward by a spike at 13, producing a distribution similar to condition ACTIVE¹². Similar examples with more prominent model disagreements, caused by ordinal responses at the extremes, can be found in Liddell and Kruschke's [43] paper.

Finally, if there is a need to predict estimates in relation to a reference condition other than the default (ACTIVE) condition, the 'relevel()' function in the R 'stats' package can be used to set the desired reference condition. While beyond the scope of this paper, post-hoc analysis of CL(M)Ms can be achieved by using appropriate methods, such as estimated marginal means [42] (emmeans R package), which would enable tests like multiple pairwise comparisons between the different conditions.

5.2 CLMM Example: Chen et al. [16]

Study Design & Aim: Chen et al.'s [16] paper explores the effects of different mechanism for adjusting display table height during follow-along home fitness instructional video viewing. Specifically, they investigate three conditions: FIXED — where the display table height cannot be adjusted, MANUAL — where the user can manually adjust the display table height using buttons attached to the table, and AUTOMATIC — where the height is automatically adjusted using a custom motion tracking system that aims to optimize viewing angles during fitness routines. The study employed a *within-subject* design with randomized order for the conditions.

Outcome measures: During the study, participant's head angle flexion was recorded, and an extended NASA-TLX survey [33] (scale 0 - 10) was administered after each condition. The extended TLX survey included two additional sub-scales related to 'Degree of Support', and 'Degree of Participation'. The paper reports the use of a weighting procedure on the NASA-TLX measures to capture individual differences in prioritizing workload dimensions. Weighting NASA-TLX measures is a common [33], yet disputed practice [3, 50, 61], and transforms the ordinal workload measures into interval scale which is incompatible for analysis using a CL(M)M. However, as Chen et al. [16] also open-source the raw NASA-TLX measures, and because the *purpose of this example is to illustrate the use of CLMMs for analysing data generated in HCI*, we focus solely on the raw ordinal workload measures.

Analysis used in the paper: The paper reports the use of a Friedman test with post-hoc Wilcoxon Signed Rank tests for pairwise comparisons between conditions of the *weighted NASA-TLX measures* Chen et al. [15]¹³. While we focus on the *raw NASA-TLX measures*, we first run the analysis reported in the paper to ensure that we have correctly imported and organized the data in R. Our analysis using the reported Friedman test reproduced the exact

¹²This distributions did not appear normal. We confirmed this using a Shapiro-Wilk's test, and ran a Kruskal-Wallis non-parametric test between SELF and ACTIVE, which also returned a non-significant difference.

¹³Additional analysis of the NASA-TLX measures with respect to participant related demographic factors were also conducted. We do not replicate these analyses as demographic data was not open-sourced.

p-values that were reported in the paper. However, we were unable to reproduce the pairwise significant differences reported in the paper using post-hoc Wilcoxon Signed Rank tests. This may be because of a mismatch between our post-hoc test parameters and those used by Chen et al. [16], as the paper does not specify these details (such as the type of post-hoc correction used, if any). However, we conclude that we correctly imported the data into our R environment, as we reproduced the exact p-values reported for the Friedman test in the paper. We provide our analysis scripts in the supplementary material for transparency.

Modelling Raw NASA-TLX sub-scales using CLMMs: The first 5 rows of our data frame containing the data from Chen et al.'s [16] paper (named *tlx_df* in R) is presented using the 'head()' function. It contains the raw NASA-TLX user score/response (*score*), information regarding the experimental condition (*condition*), information regarding the NASA-TLX sub-scale (*tlx_construct*), and the participant identifier (*participant_id*).

```
head(tlx_df, 5)
  condition tlx_construct participant_id score
<fct>      <chr>         <fct>      <dbl>
1 Fixed    Mental Demand    p1           5
2 Fixed    Mental Demand    p2           3
3 Fixed    Mental Demand    p3           2
4 Fixed    Mental Demand    p4           8
5 Fixed    Mental Demand    p5           2
```

Fitting a CLMM to the data can be achieved by using the 'ordinal' package's 'clmm()'. Similar to the clm() function described in section 5.1, we can use the 'summary()' function to get an overview of the results. To illustrate the results returned by a CLMM, we present the output of fitting a CLMM to one example sub-scale (*Physical Demand*) of the raw NASA-TLX measures provided in Chen et al.'s [16] paper. Results for the remaining sub-scales can be found in our supplementary materials.

```
# dataframe with only "Physical Demand" measures
pd_df <- tlx_df[
  tlx_df$tlx_construct == "Physical Demand",
]

physical_demand.clmm <- clmm(
  score ~ 1 + condition + (1|participant_id),
  data = pd_df,
  link = "probit")

summary(physical_demand.clmm)

formula: score ~ 1 + condition + (1 | participant_id)
data:    pd_df

link threshold nobs logLik AIC niter max.grad cond.H
probit flexible 90 -161.75 343.50 706(2737) 3.18e-04 1.0e+02

Random effects:
Groups Name Variance Std.Dev.
participant_id (Intercept) 1.871 1.368
Number of groups: participant_id 30

Coefficients:
Estimate Std. Error z value Pr(>|z|)
```

```

conditionFixed  0.4326    0.2777    1.558    0.11925
conditionManual 0.7551    0.2844    2.655    0.00793 **
---
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Threshold coefficients:
  Estimate Std. Error z value
2|3 -2.2670    0.4800 -4.723
3|4 -1.1828    0.3703 -3.194
4|5 -0.6183    0.3458 -1.788
5|6 -0.1894    0.3371 -0.562
6|7  0.3812    0.3390  1.124
7|8  1.3784    0.3651  3.775
8|9  2.5397    0.4396  5.777

```

Similar to the results returned by the ‘clm()’ function, the ‘clmm()’ summary returns the primary results of model fitting within the ‘Coefficients’ and ‘Threshold Coefficients’ tables. These can be interpreted in the same way as for the results return by the ‘clm()’ summary, and is detailed in section 5.1. In this example, the results suggest that the parameter estimate for the *Manual* condition is significantly different from zero¹⁴, i.e., the user’s internal psychological measure of physical demand (\bar{Y}) for the *Manual* condition is different from the reference condition (*Automatic*) by an estimate that is significantly different from zero¹⁵. The ‘Threshold Coefficients’ table also suggest that the most likely response for the reference (*Automatic*) condition is $Y = 6$ as the assumed location of the distribution ($\mu = 0$) falls between cutpoints 5|6 (τ_5) and 6|7 (τ_6). Note that insufficient data related to the lower responses (0 and 1) leads the ‘clmm()’ summary to return a ‘Threshold Coefficients’ table with only 7, out of 10, cutpoints, i.e., given 11 response categories (0 - 10) for the ‘Physical Demand’ sub-scale used in the paper, there should be $(11 - 1) = 10$ cutpoints to partition the psychological measure (\bar{Y}) in 11 categories (see section 4).

Additionally, the CLMM model summary also returns a ‘Random effects’ table which details group specific random intercept and slope variations (via the presented ‘Variance’ and ‘Std.Dev’). In this case, our formula specifies that we want to model for the random intercept for each participant – signified by the $(1|participant_id)$ in our formula. The results suggest that participant’s internal measure of physical demand varied by a standard deviation of 1.368 units. This indicates a large variation between participants’ psychological measures of ‘Physical demand’, given the comparatively small values for cutpoint (τ) and parameter (β) estimates presented in the ‘Coefficients’ and ‘Threshold Coefficients’ tables.

5.3 Reporting CL(M)M Analysis Results

CLM/CLMM models report estimates (coefficients, std. error, etc.) relative to the latent variable scale, and is dependent on the ‘link’ function used. Consequently, these values should be interpreted and reported considering these factors. As mentioned previously, both the ‘logit’ and ‘probit’ link functions return similar model fits but the returned estimates should be interpreted differently. We briefly describe ‘logit’ model result interpretation and reporting, before detailing ‘probit’ results and reporting.

¹⁴Greater than zero as the estimate is positive.

¹⁵We refrain from making any comparisons with the analysis results presented in the original paper [16] as our analysis focuses on the raw NASA-TLX measures as opposed to the weighted NASA-TLX measures used in the paper.

Ordinal ‘Logit’ models return estimates as log-odds ratios. The coefficient estimate for a specific condition relative to the reference can be interpreted as the odds of that condition resulting in a higher response category than the reference condition [9]. Let us remodel the ‘SIM-TLX:Physical Demand’ measure presented in section 5.1 using a CLM with a ‘Logit’ model. We forego presenting all estimates for brevity, but only present the significant coefficient for the SELF condition relative to the reference (ACTIVE) for this illustration. The ‘logit’ model returns a log-odds value of -1.1951 , which when exponentiated (inverse log function) gives us the odds of observing higher ratings in the SELF condition compared to the ACTIVE condition. This can be reported as: Users in the SELF condition had an odds ratio of $\exp(-1.1951) = 0.303$ compared to the ACTIVE condition, indicating they were significantly – approximately 3.3 times ($1/0.303$) or 70% ($1 - 0.303$ or $100 - 30.3$) – **less likely** to report higher physical demand than in the ACTIVE condition.

Ordinal ‘Probit’ models return estimates on a standard normal scale for the assumed latent variable. This allows for direct interpretation of the effects on the latent scale, but prevents easy interpretation on the response or outcome scale. Let us take the same example ‘SIM-TLX:Physical Demand’ used for illustration of the ‘logit’ model. Here, the ‘probit’ model returns a coefficient estimate of -0.6751 for the SELF condition against the reference ACTIVE condition. This is interpreted as a *significant reduction on the underlying latent score that determines the response of the user*, suggesting that users are less likely to report higher ratings for the SELF condition compared to the ACTIVE condition. This can be reported as: Users in the SELF condition were significantly less likely to report higher physical demand than those in the ACTIVE condition, with a decrease of -0.6751 on the latent score scale. The reliance on the latent scale makes ‘probit’ model results easier to report for audiences familiar with latent models, foregoing the need to perform additional functions, such as exponentiation in case of ‘logit’ models. However, this also implies that unfamiliar audiences would require appropriate prior introductions. As such, we recommend ‘probit’ model use be accompanied by a brief acknowledgement of the underlying latent variable, along with appropriate references to more detailed literature on latent models (an example can be seen in the paper by Brailsford et al. [6]).

In addition to the above, reporting coefficient and threshold coefficient estimate means and standard errors for both ordinal ‘probit’ and ‘logit’ model results greatly help familiar readers in eliciting more detailed insights. We also recommend reporting assumptions, validity tests for assumptions, and if methods were employed to relax certain assumptions. Finally, we recommend accompanying these numbers and text with appropriate visualisations, such as for the frequency/proportions of the data, and confidence intervals of model estimates as shown in Figure 12 (top & middle).

6 Discussion

In this paper, we have shown the extent of divergent practices in analysing ordinal data within current HCI literature, discussed the associated concerns with the prevalent methods currently employed, and detailed the theory and application of CL(M)Ms as a more suitable method for ordinal data analysis with HCI related

examples. Despite the stated advantages of CL(M)Ms over traditional statistical approaches for ordinal data analysis in HCI, certain limitations and assumptions should be considered.

Importantly, CL(M)Ms impose several assumptions that influence how results are interpreted, and these should be evaluated for validity and suitability. Take for example our re-analysis of the ‘Perceived Usefulness’ measure provided by Fitton et al. [26] in section 5.1. Here, we used a CLM with the ‘probit’ link function, and assumed proportional odds and equal variance. Use of the ‘probit’ function forces interpretations on the latent scale. Specifically, effects of different conditions on ‘Perceived Usefulness’ are estimated as shifts in a standard normal latent distribution, and cutpoints ‘divide’ these latent distributions into proportions that correspond to the ‘Perceived Usefulness’ ordinal categories. This is contrary to the use of the ‘logit’ link which enables interpretation as odds ratios (see section 5.3). Additionally, the proportional odds assumption enables us to inspect how estimates for conditions influence the latent distribution positions, without considering how different conditions affect the dividing cutpoints. This is evident in our elaboration, which highlights the shift in the latent distribution (\bar{Y}) and inspecting how the ‘static’ cutpoints (\mathcal{T}) predict proportions of the ordinal response Y . However, such assumptions are often violated and should be tested for, and relaxed when necessary and possible (detailed in section 4.3). This is also true for the equal variance of latent distribution assumption. But this can be relaxed using model extensions discussed in section 4.3.

Both a strength and a limitation of CL(M)Ms lies in their *design for analysing strictly ordinal data*. CL(M)Ms treat data as categorical with inherent ordering, enabling ordinal data analysis without requiring contiguous numerical labels for the ordered categories. As such, CL(M)Ms can be perfectly fit to data with ordinal responses that use non-numerical or non-contiguous labels (such as {A,B,C} or {1,7, 91}) as long as the order is defined. However, treating data as categorical also means that CL(M)Ms cannot handle ordinal data that has been transformed to interval or ratio scale, as such scales imply infinite categories, and therefore, infinite cutpoint estimates. This is evident by our inability to apply CLMMs to analyse the *weighted* NASA-TLX measures in section 5.2, prompting us to instead focus on the *raw* NASA-TLX measures provided by Chen et al. [16]. Though possible, CL(M)Ms can also struggle to converge when applied to ordinal data with a large number of ordered categories (say a questionnaire with a slider response ranging from 0 to 100). In such cases where ordinal data needs to be transformed to metric data, or when the number of ordered categories is substantially large, it may be worth considering non-parametric alternatives that do not impose incompatible assumptions on ordinal data (see table 1), or more specialized methods, such as beta regression models [31]. While some prior work argue for the applicability of parametric methods with metric assumptions for analysing ordinal data under metric interpretations [13, 14, 59, 60], the evidence for potential errors provided by Liddell and Kruschke [43] serves as a strong deterrent, and should be carefully considered when opting for such approaches.

Finally, while we demonstrated Frequentist implementations of CL(M)Ms in this paper, these methods also have Bayesian counterparts. The *brms* package in R [10] uses a similar model specification

syntax as the ordinal package. The biggest difference is in the requirement for the specification of the priors for the main effects (as in any Bayesian analysis) and for the cutpoints. An example of a default uninformative prior is offered by Kurz [41], who uses normal distributions centred at the $1/(N-1)$ quantiles of the standard normal (where N is the number of possible scores), which implies a uniform prior for the distribution of scores. Examples of Bayesian analyses using ordinal regression models can be found in prior HCI work, including the use of CL(M)Ms with probit [6, 62] and logit [35] links.

7 Conclusion

As with most practices, it is pertinent to occasionally reflect on the methods that we use in HCI and update them where necessary. The case of ordinal data analysis within HCI is a prime example where such reflection can spotlight divisive opinions and approaches. Our review of the current HCI literature illustrates this well, highlighting not only disagreements in the use of parametric and non-parametric methods for ordinal data analysis between research articles, but also inconsistencies in the imposed data assumptions for ordinal data analysis between successive statistical procedures within the same study. Such inconsistencies hinder scientific progress by increasing the chances of propagating erroneous insights, introducing challenges to reproducible findings, and limiting comparisons between analytically distinct but related experiments. To address these issues, we propose the use of cumulative link (mixed) models, a form of ordinal regression modelling widely recommended for ordinal data analysis in disciplines such as psychology [11, 17, 43]. We explain the theory and mathematical foundations of CL(M)Ms, highlight its advantages over traditional methods used in HCI, and present practical examples of applying CL(M)Ms to published open-sourced ordinal data in HCI using the ‘ordinal’ package in R [18]. Through these examples and the ensuing discussions, we provide HCI researchers with the necessary knowledge of when and how to apply CL(M)Ms to their data. We hope that our paper presents a convincing argument for the HCI community to adopt statistical approaches that are suitable for the intended data, and more broadly, to continue to reflect on, and update, our practices.

References

- [1] Alan Agresti and Maria Kateri. 2011. Categorical data analysis. In *International encyclopedia of statistical science*. Springer, 206–208.
- [2] M Agüena, C Avestruz, C Combet, S Fu, R Herbonnet, A I Malz, M Penna-Lima, M Ricci, S D P Vitenti, L Baumont, H Fan, M Fong, M Ho, M Kirby, C Payerne, D Boutigny, B Lee, B Liu, T McClintock, H Miyatake, C Sifón, A von der Linden, H Wu, M Yoon, and The LSST Dark Energy Science Collaboration. 2021. CLMM: a LSST-DESC cluster weak lensing mass modeling library for cosmology. *Monthly Notices of the Royal Astronomical Society* 508, 4 (10 2021), 6092–6110. doi:10.1093/mnras/stab2764 arXiv:https://academic.oup.com/mnras/article-pdf/508/4/6092/41074412/stab2764.pdf
- [3] Ebrahim Babaei, Tilman Dingler, Benjamin Tag, and Eduardo Velloso. 2025. Should we use the NASA-TLX in HCI? A review of theoretical and methodological issues around Mental Workload Measurement. *International Journal of Human-Computer Studies* (2025), 103515.
- [4] Vance W Berger and YanYan Zhou. 2014. Kolmogorov–smirnov test: Overview. *Wiley statsref: Statistics reference online* (2014).
- [5] Kenneth A Bollen. 2002. Latent variables in psychology and the social sciences. *Annual review of psychology* 53, 1 (2002), 605–634.
- [6] Joe Brailsford, Frank Vetere, and Eduardo Velloso. 2024. Exploring the association between moral foundations and judgements of AI behaviour. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–15.

- [7] Rollin Brant. 1990. Assessing proportionality in the proportional odds model for ordinal logistic regression. *Biometrics* (1990), 1171–1178.
- [8] John Brooke et al. 1996. SUS-A quick and dirty usability scale. *Usability evaluation in industry* 189, 194 (1996), 4–7.
- [9] J. Bruin. 2011. *newtest: command to compute new test @ONLINE*. <https://stats.oarc.ucla.edu/other/mult-pkg/faq/ologit/>
- [10] Paul-Christian Bürkner. 2017. brms: An R package for Bayesian multilevel models using Stan. *Journal of statistical software* 80 (2017), 1–28.
- [11] Paul-Christian Bürkner and Matti Vuorre. 2019. Ordinal regression models in psychology: A tutorial. *Advances in Methods and Practices in Psychological Science* 2, 1 (2019), 77–101.
- [12] Michael Calver and Douglas Fletcher. 2020. When ANOVA isn't ideal: Analyzing ordinal data from practical work in biology. *The American Biology Teacher* 82, 5 (2020), 289–294.
- [13] James Carifio and Rocco Perla. 2008. Resolving the 50-year debate around using and misusing Likert scales. 1150–1152 pages.
- [14] James Carifio and Rocco J Perla. 2007. Ten common misunderstandings, misconceptions, persistent myths and urban legends about Likert scales and Likert response formats and their antidotes. *Journal of social sciences* 3, 3 (2007), 106–116.
- [15] Mingyang Chen, Peng Chen, Xu Gao, and Chao Yang. 2020. Examining injury severity in truck-involved collisions using a cumulative link mixed model. *Journal of Transport & Health* 19 (2020), 100942.
- [16] Xinyu Chen, Yuqi Li, Jintao Chen, Jiabao Li, Chong Wang, and Pinyan Tang. 2024. Enhancing Home Exercise Experiences with Video Motion-Tracking for Automatic Display Height Adjustment. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 310, 13 pages. doi:10.1145/3613904.3642936
- [17] Rune Haubo B Christensen. 2018. Cumulative link models for ordinal regression with the R package ordinal. *Submitted in J. Stat. Software* 35 (2018), 1–46.
- [18] Rune H. B. Christensen. 2023. *ordinal—Regression Models for Ordinal Data*. <https://CRAN.R-project.org/package=ordinal> R package version 2023.12-4.1.
- [19] WJ Conover and RL Iman. 1979. *Multiple-comparisons procedures. Informal report*. Technical Report. Los Alamos National Lab.(LANL), Los Alamos, NM (United States).
- [20] William Jay Conover. 1999. *Practical nonparametric statistics*. John Wiley & Sons.
- [21] Anthony Christopher Davison and David Victor Hinkley. 1997. *Bootstrap methods and their application*. Number 1. Cambridge university press.
- [22] Alexis Dinno. 2015. Nonparametric pairwise multiple comparisons in independent groups using Dunn's test. *The Stata Journal* 15, 1 (2015), 292–300.
- [23] Pierre Dragicevic, Yvonne Jansen, Abhira Sarma, Matthew Kay, and Fanny Chevalier. 2019. Increasing the transparency of research papers with explorable multiverse analyses. In *proceedings of the 2019 chi conference on human factors in computing systems*. 1–15.
- [24] Olive Jean Dunn. 1964. Multiple comparisons using rank sums. *Technometrics* 6, 3 (1964), 241–252.
- [25] Lisa A Elkin, Matthew Kay, James J Higgins, and Jacob O Wobbrock. 2021. An aligned rank transform procedure for multifactor contrast tests. In *The 34th annual ACM symposium on user interface software and technology*. 754–768.
- [26] Isabel Sophie Fitton, Elizabeth Dark, Manoela Milena Oliveira da Silva, Jeremy Dalton, Michael J Proulx, Christopher Clarke, and Christof Lutteroth. 2024. Watch This! Observational Learning in VR Promotes Better Far Transfer than Active Learning for a Fine Psychomotor Task. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 721, 19 pages. doi:10.1145/3613904.3642550
- [27] Benjamin French and Matthew S Shotwell. 2022. Regression models for ordinal outcomes. *Jama* 328, 8 (2022), 772–773.
- [28] Milton Friedman. 1937. The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the American Statistical Association* 32, 200 (1937), 675–701.
- [29] Filippo Gambarota and Gianmarco Altoè. 2024. Ordinal regression models made easy: A tutorial on parameter interpretation, data simulation and power analysis. *International Journal of Psychology* 59, 6 (2024), 1263–1292.
- [30] Paul A Games and John F Howell. 1976. Pairwise multiple comparison procedures with unequal n's and/or variances: a Monte Carlo study. *Journal of Educational Statistics* 1, 2 (1976), 113–125.
- [31] Jason Geller, Matti Vuorre, Chelsea Parlett, and Robert Kubinec. 2025. A Beta Way: A Tutorial For Using Beta Regression in Psychological Research. doi:10.31234/osf.io/d6v5t_v2
- [32] William H Greene and David A Hensher. 2010. *Modeling ordered choices: A primer*. Cambridge University Press.
- [33] Sandra G Hart and Lowell E Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of empirical and theoretical research. In *Advances in psychology*. Vol. 52. Elsevier, 139–183.
- [34] Andrew F Hayes. 2012. PROCESS: A versatile computational tool for observed variable mediation, moderation, and conditional process modeling.
- [35] Silas Hsu, Vinay Koshy, Kristen Vaccaro, Christian Sandvig, and Karrie Karahalios. 2025. Placebo Effect of Control Settings in Feeds Are Not Always Strong. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–16.
- [36] Susan Jamieson. 2004. Likert scales: How to (ab) use them? *Medical education* 38, 12 (2004), 1217–1218.
- [37] Maurits Clemens Kaptein, Clifford Nass, and Panos Markopoulos. 2010. Powerful and consistent analysis of likert-type rating scales. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 2391–2394.
- [38] Ani Katchova. 2020. *Econometrics Academy*. <https://sites.google.com/site/econometricsacademy/econometrics-models/ordered-probit-and-logit-models?authuser=0>
- [39] Diana Eugenie Kornbrot. 1990. The rank difference test: A new and meaningful alternative to the Wilcoxon signed ranks test for ordinal data. *Brit. J. Math. Statist. Psych.* 43, 2 (1990), 241–264.
- [40] William H Kruskal and W Allen Wallis. 1952. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association* 47, 260 (1952), 583–621.
- [41] Solomon Kurz. [n. d.]. Causal Inference with Ordinal Regression. <https://solomonkurz.netlify.app/blog/2023-05-21-causal-inference-with-ordinal-regression/>. Accessed: 2025-08-28.
- [42] Russell V. Lenth. 2025. *emmeans: Estimated Marginal Means, aka Least-Squares Means*. <https://CRAN.R-project.org/package=emmeans> R package version 1.10.7.
- [43] Torrin M Liddell and John K Kruschke. 2018. Analyzing ordinal data with metric models: What could possibly go wrong? *Journal of Experimental Social Psychology* 79 (2018), 328–348.
- [44] Rensis Likert. 1932. A technique for the measurement of attitudes. *Archives of psychology* (1932).
- [45] Helen M Marcus-Roberts and Fred S Roberts. 1987. Meaningless statistics. *Journal of Educational Statistics* 12, 4 (1987), 383–394.
- [46] Peter McCullagh. 1980. Regression models for ordinal data. *Journal of the Royal Statistical Society: Series B (Methodological)* 42, 2 (1980), 109–127.
- [47] David Moher, Alessandro Liberati, Jennifer Tetzlaff, Douglas G Altman, and PRISMA Group*. 2009. Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement. *Annals of internal medicine* 151, 4 (2009), 264–269.
- [48] Peter Bjorn Nemenyi. 1963. *Distribution-free multiple comparisons*. Princeton University.
- [49] Geoff Norman. 2010. Likert scales, levels of measurement and the “laws” of statistics. *Advances in health sciences education* 15 (2010), 625–632.
- [50] Thomas E Nygren. 1991. Psychometric properties of subjective workload measurement techniques: Implications for their use in the assessment of perceived mental workload. *Human factors* 33, 1 (1991), 17–33.
- [51] Nick Parsons. 2016. *repolr: Repeated Measures Proportional Odds Logistic Regression*. <https://CRAN.R-project.org/package=repolr> R package version 3.4.
- [52] Karl Pearson. 1900. X. On the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 50, 302 (1900), 157–175.
- [53] Bercedis Peterson and Frank E Harrell Jr. 1990. Partial proportional odds models for ordinal response variables. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 39, 2 (1990), 205–217.
- [54] Thorsten Pohlert. 2024. *PMCMRplus: Calculate Pairwise Multiple Comparisons of Mean Rank Sums Extended*. <https://CRAN.R-project.org/package=PMCMRplus> R package version 1.9.12.
- [55] Mariah Schrum, Muyleng Ghuy, Erin Hedlund-Botti, Manisha Natarajan, Michael Johnson, and Matthew Gombolay. 2023. Concerning trends in likert scale usage in human-robot interaction: Towards improving best practices. *ACM Transactions on Human-Robot Interaction* 12, 3 (2023), 1–32.
- [56] Skipper Seabold and Josef Perktold. 2010. statsmodels: Econometric and statistical modeling with python. In *9th Python in Science Conference*.
- [57] Mital C Shingala and Arti Rajyaguru. 2015. Comparison of post hoc tests for unequal variance. *International Journal of New Technologies in Science and Engineering* 2, 5 (2015), 22–33.
- [58] Sidney Siegel. 1957. Nonparametric statistics. *The American Statistician* 11, 3 (1957), 13–19.
- [59] Miguel Alejandro Silan. 2025. When can we treat Likert type data as interval? (2025).
- [60] Laura South, David Saffo, Olga Vitek, Cody Dunne, and Michelle A Borkin. 2022. Effective use of Likert scales in visualization evaluations: A systematic review. In *Computer Graphics Forum*, Vol. 41. Wiley Online Library, 43–55.
- [61] Kai Virtanen, Heikki Mansikka, Helmiina Kontio, and Don Harris. 2022. Weight watchers: NASA-TLX weights revisited. *Theoretical issues in ergonomics science* 23, 6 (2022), 725–748.

- [62] Samangi Wadinambiarachchi, Ryan M Kelly, Saumya Pareek, Qiushi Zhou, and Eduardo Velloso. 2024. The effects of generative ai on design fixation and divergent thinking. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*. 1–18.
- [63] Frank Wilcoxon. 1945. Individual comparisons by ranking methods. *Biometrics bulletin* 1, 6 (1945), 80–83.
- [64] Andre Williams. 2025. Cumulative Link Ordinal Outcome Neural Networks: An Evaluation of Current Methodology. *Journal of Data Analysis and Information Processing* 13, 2 (2025), 182–198.
- [65] Jacob O. Wobbrock, Leah Findlater, Darren Gergle, and James J. Higgins. 2011. The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (Vancouver, BC, Canada) (CHI '11)*. Association for Computing Machinery, New York, NY, USA, 143–146. doi:10.1145/1978942.1978963

A Appendix

A.1 Additional Plots




Figure 13: Visualizations relevant to Fitton et al.'s [26] paper. (Top) Frequency of user responses on the 'SIM-TLX: Stress' sub-scale across learning conditions. (Middle) Confidence intervals for pairwise differences in mean responses between conditions under ANOVA and CLM models. (Bottom) Frequency plots for 'Stress' in the SELF and ACTIVE conditions, which differed significantly under CLMs. Dashed curves show fitted linear (metric) model distributions based on condition means and standard deviations.

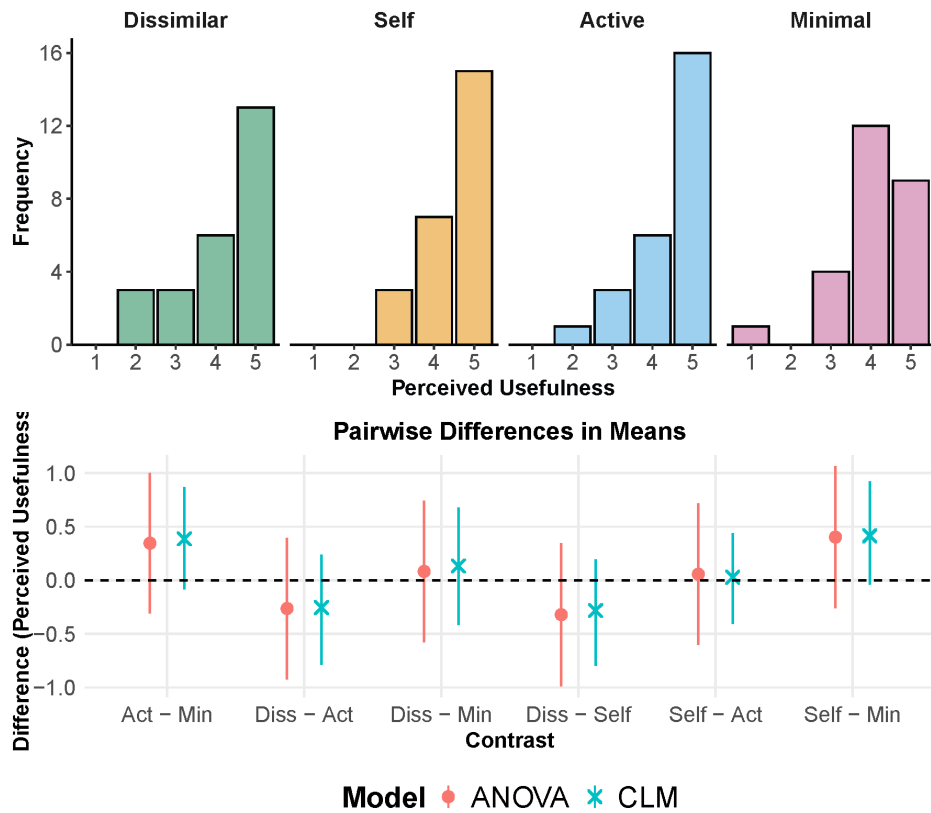


Figure 14: Visualizations relevant to Fitton et al.'s [26] paper. (Top) Frequency of user responses to the 'Perceived Usefulness' measure across learning conditions. (Middle) Confidence intervals for pairwise differences in mean responses between conditions under ANOVA and CLM models. No disagreements were found between the ANOVA and CLM model analyses.

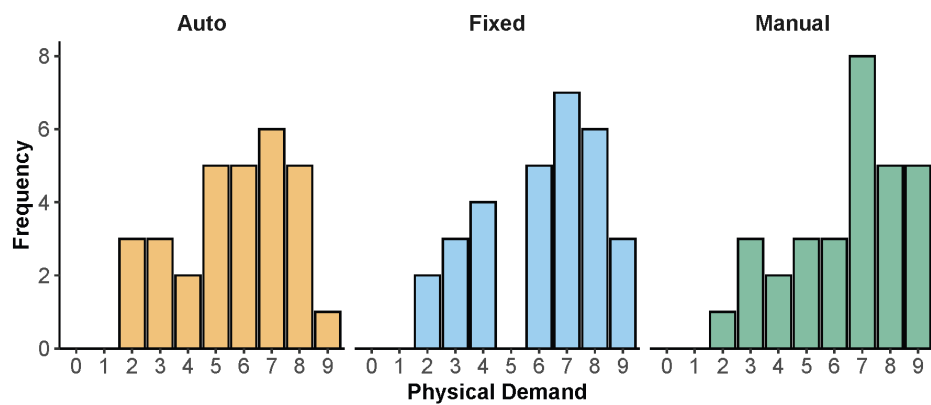


Figure 15: Visualizations relevant to Chen et al.'s [16] paper. (Top) Frequency of user responses on the 'NASA-TLX: Physical Demand' sub-scale across conditions. Comparative plots such as mean difference confidence intervals between modelling approaches are not presented as Chen et al.'s [16] paper uses weighted NASA-TLX scores which cannot be analysed using an ordinal regression model like CLMM. This prevented us from plotting any comparative plots between statistical analysis methods used in the paper and our re-analysis using a CLMM.




Figure 16: Sankey diagram showing the progression of statistical procedures applied to constructs measured with ordinal data. Methods are displayed in sequence from predictive models (left), to omnibus tests (centre), to pairwise tests (right). Bars labelled 'No' indicate constructs for which no method was reported in that category. Each node represents a statistical method and is coloured uniquely. This plot is the same as figure 5 but with unique colours for each statistical method observed in our sample for easy identification.