

Equitable Ability Estimation in Neurodivergent Student Populations with Zero-Inflated Learner Models

Niall Twomey^{*}
Kidsloop Ltd, UK
niall.twomey@kidsloop.live

Sarah McMullan
auticon, UK; Kidsloop Ltd, UK
sarah.mcmullan@auticon.co.uk

Anat Elhalal
Kidsloop Ltd, UK
anat.elhalal@kidsloop.live

Rafael Poyiadzi
Kidsloop Ltd, UK
rafael.poyiadzi@kidsloop.live

Luis Vaquero
Kidsloop Ltd, UK
luis.vaquero@kidsloop.live

ABSTRACT

At present, the educational data mining community lacks many tools needed for ensuring equitable ability estimation for Neurodivergent (ND) learners. On one hand, most learner models are susceptible to under-estimating ND ability since confounding contexts cannot be held accountable (*e.g.* consider dyslexia and text-heavy assessments), and on the other, few (if any) existing datasets are suited for appraising model and data bias in ND contexts. In this paper we attempt to model the relationships between context (delivery and response types) and performance of ND students with zero-inflated learner models. This approach facilitates simulation of several expected ND behavioural traits, provides equitable ability estimates across all student groups from generated datasets, increases interpretability confidence, and can significantly increase the quality of learning opportunities for ND students. Our approach consistently out-performs baselines in our experiments and can also be applied to many other learner modelling frameworks.

Keywords

Neurodiversity, Zero-Inflated Models, Learner Models, Item Response Theory, Data Simulation

1. INTRODUCTION

In the UK, it is estimated that 15% of the population are ND, having neurological functions that differ from what is considered typical [22]. Neurodiversity covers the range of differences in individual brain function and behavioural traits, regarded as part of normal variation in the human population [37]. Each Neurodivergent Condition (NDC) uniquely affects how information is absorbed, processed, and communicated [30, 4]. Our objective is to adapt Learner Models (LMs) for the individual requirements of a number of NDCs

^{*}Corresponding author.

N. Twomey, S. McMullan, A. Elhalal, R. Poyiadzi, and L. Vaquero. Equitable ability estimation in neurodivergent student populations with zero-inflated learner models. In A. Mitrovic and N. Bosch, editors, *Proceedings of the 15th International Conference on Educational Data Mining*, pages 571–577, Durham, United Kingdom, July 2022. International Educational Data Mining Society.

© 2022 Copyright is held by the author(s). This work is distributed under the Creative Commons Attribution NonCommercial NoDerivatives 4.0 International (CC BY-NC-ND 4.0) license.

<https://doi.org/10.5281/zenodo.6853000>

in learning environments, focusing specifically on dyslexia, dyscalculia and Sensory Processing Disorder (SPD) (with prevalences of 10%, 6% and 5-15% respectively [7, 36, 9]).

Achievement gaps due to NDCs occur early in life and persist through adolescence into adulthood [8]. In many cases, impeded learning opportunities for ND students result from unsuitable learning contexts or lack of adequate student support rather than intrinsically low student ability [29]. However, as learning begins to move further into the digital space [14, 34], LMs, which are statistical models of student attainment, will use historic performance to estimate student ability. Owing to a legacy of potentially poor learning contexts, the ability of ND students tends to be underestimated by LMs since they are not equipped to distinguish between context- and ability-based explanations of performance. Without deliberate effort, therefore, it is very likely that LMs will become biased and offer inequitable recommendations for ND students. On the other hand, opportunities to quell these achievement gaps before they grow are at hand in smart learning environments if LMs are empowered to reason about alternative explanations of performance.

LM research is highly active in the Educational Data Mining (EDM) community. State-of-the-art approaches include deep neural networks [33, 11, 28], and nonparametric Bayesian methods [15]. We find that the literature is sparse for inclusive LMs applied to ND populations, and we were unable to find many bespoke models or datasets (real or synthetic) even in recent literature reviews [1, 21]. Kohli *et. al.* [16] introduced an approach for identifying dyslexic students based on historic patterns of behaviour and artificial neural networks. Mejia *et. al.* [26] approached the task by estimating learner's cognitive deficit specifically for students with dyslexia or reading difficulties. Ensuring the equity of LM is an important area of research, and learning interfaces can be improved by offering multiple assessment Delivery and Response Type (DRT) [29]. Other works have elaborated further on scores and metrics for ethical and equitable recommendation systems with broad stakeholders, including dyslexic students [25]. Equity is also explored along explainability and interpretability axes. Some classical LMs are readily interpretable and offer intuitive explanations of datasets [31, 24], though caution must be exercised to avoid over-interpreting models [13].

ND students face at least two additional hurdles than Neurotypical (NT) students in learning environments: 1) their ability is inaccurately modelled due to LMs shortcomings; and 2) choosing the most suitable learning context for them to express their true ability is rarely considered. Furthermore, the EDM community currently lacks datasets and simulation tools for developing LMs and assessing equity for NDC contexts. We address these three limitations in this work, by motivating and defining equitable LMs for ND students (Sec 2), defining a simulation environment (Sec 2.2), and demonstrating strong performance in our results and conclusions (Secs 3 and 4).

2. METHODS

Due to a lack of available datasets that include ND students, we explore equitable estimation in simulations. Our model combines the use of Zero-Inflated Models (ZIMs) [17] and Item Response Theory (IRT) [2, 20]. Our assumption is that DRT choices will affect the quality of learning opportunities for ND students, with unsuitable DRT resulting in lower Learning Quality Factor (LQF). Without considering the suitability of DRTs for students, LMs risk recommending low-quality learning opportunities and mis-interpreting poor performance on these as an indication of low student ability. The model and simulation procedure proposed is designed to be used to identify the best DRTs for each student, and prevent underestimation of abilities.

2.1 IRT-based Zero-Inflated Learner Model

Our proposed approach, Zero-inflated Learner Models (ZILMs), shown in Eqn (1), builds on the assumption that there are two explicit explanations of zeros: 1) low ability relative to difficulty (low p); and 2) low LQF (high π). With this formulation, a zero from a student with high ability with in an unsuitable DRT can be explained by the poor LQF since π has high responsibility for the outcome [3].

$$Pr(Y = y) = \begin{cases} \pi + (1 - \pi) \cdot (1 - p) & \text{if } y = 0 \\ (1 - \pi) \cdot p & \text{if } y = 1 \end{cases} \quad (1)$$

In our setting, p is based on IRT, and π (which reflects LQFs) is parameterised by item, NDC and DRT features (*c.f.* Sec 2.2), resulting in IRT-based ZILM (IRT-ZILM).

IRT was chosen as the base LM in IRT-ZILM over alternative options as: 1) IRT is well-understood and simple to interpret; 2) Bayesian Knowledge Tracing (BKT) is known to have over- and under-estimation problems [6, 18] that may muddle our understanding of equity for ND students; 3) several technical hurdles need to be overcome to incorporate our approach into BKT; and 4) although Deep Knowledge Tracing (DKT) [33] models can probably learn latent representations that correlate to DRT preferences, this is at the expense of control and interpretation of the effects.

2.2 Simulations

In the simulated dataset, we assume that the ability of ND and NT students are drawn from the same distribution, meaning that ability and NDCs are independent. The NDCs considered in this initial work are dyslexia, dyscalculia, and SPD. These chosen conditions reflect a wide range of effects from different delivery and response types, but this work could be applied to others.

Table 1: Description of parameter distributions used to generate synthetic dataset. Each parameter was randomly assigned from distributions. Users are given an intrinsic ability and the possibility of one or more ND conditions. Items are assigned a difficulty, discrimination, guessing, subject, content type, information density, delivery type and response type. Information density describes how much information is provided—0.1 represents only a few words, 1 is a large block of text—designed to reflect how clearly an item is presented.

Parameter	Value (Range)	Probability
Ability	$(-\infty, \infty)$	$\mathcal{N}(0, 1)$
ND condition	Dyslexia, Dyscalculia, SPD	0.1, 0.06, 0.11
Difficulty	$(-2, 2)$	uniform
Discrimination	$(0.5, 4)$	uniform
Guessing	$(0, 0.15)$	uniform
Subject	Maths, English	0.5, 0.5
Content type	Letter, Digit, Both	M: 0.1, 0.5, 0.6, E: 1, 0, 0
No. attempts	20	fixed
Info. density	$(0.1, 1)$	$\mathcal{N}(0.35, 0.15)$
Delivery type	Read, Listen, Both	0.3, 0.3, 0.4
Response type	Written, Speak, Click Picture,	0.4, 0.2, 0.2, 0.2
	Click Read	

Datasets are created based on the parameters outlined in Table 1. These features contribute to the estimation of LQFs and the probability a user will respond to an item. For example, a dyslexic user’s learning quality is impacted by delivery types involving reading letters, and response types involving reading letters to click the correct answer(s) or writing an answer that includes letters. A dyscalculic user is affected by delivery and response types involving digits. And someone with SPD is impacted when the delivery involves both reading and listening with either letters and/or digits, as this can cause sensory overload [29].

Collectively, these features are used to describe the suitability of DRT to a variety of NDCs, which we now relate back to Eqn (1). If a poorly chosen DRT is selected for a ND student, this will result in poor learning opportunities due

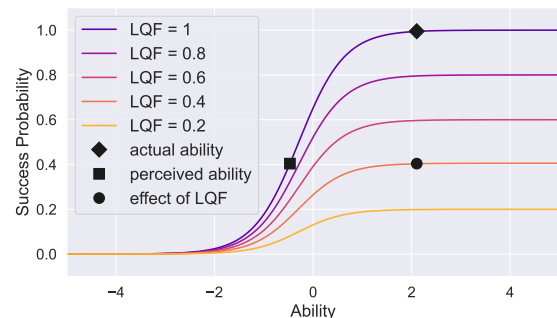


Figure 1: This image shows that LQFs can make the perceived ability of an affected student much lower than their true unobserved ability. \diamond shows a student’s true ability, \circ shows the impact of low LQFs, and \square shows the perceived ability if LQF is not considered.

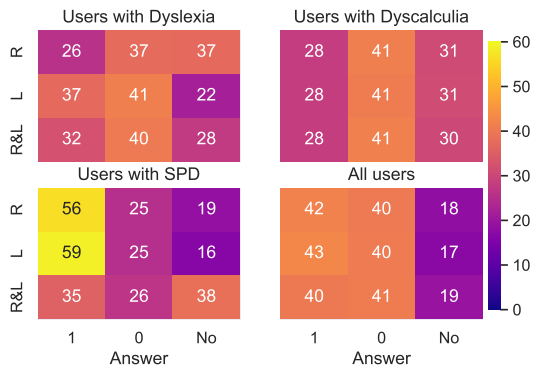


Figure 2: Comparison of attempt outcomes for each NDC (dyslexia: upper left; dyscalculia: upper right; SPD: lower left; all: lower right) when a single delivery type is used for all items (R and L correspond to ‘read’ and ‘listen’ respectively).

to a low LQF (*i.e.* large π). However, if a suitable DRT is selected for a student, the suitability is reflected in higher LQFs. Synthesising datasets that adapt to DRTs and NDCs requires specification of the weight vectors to adapt π to context (*e.g.* ‘reading’ should increase π / reduce LQF for dyslexic but not for dyscalculic students). Although specification of weight vectors is a subjective process, it allows us to express our intuition and instincts about the influential pathways. These are fully described in our implementation¹.

The effect of LQFs on an item’s characteristic curve can be seen in Fig 1. As the LQF decreases, the upper asymptote is reduced, indicating that their opportunity to learn from the interaction is compromised. With this, we interpret LQF as a measure of the contextual inequity.

3. RESULTS AND DISCUSSION

There are four main questions we want to explore in this work: 1) how much are ND users learning opportunities impacted by poor DRTs; 2) is it possible to identify users with potential NDC based on their performance on items with a range of DRTs; 3) is it possible to estimate user true abilities, accounting for any poor performance due to other factors; and 4) can student learning quality and success be improved through active selection of DRTs?

3.1 How are ND users impacted?

Fig 2 shows how ND student performance is affected if a learning environment only delivers information in a single format. Across the full neurodiverse population, the mean performance is approximately the same for all learning material formats. There are also no observable differences in performance for users with dyscalculia. However, for users with dyslexia or SPD there are noticeable differences. For users with dyslexia, they answer 6–11% more attempts correctly and are able to attempt 9–15% more items when the item has a listening component. For users with SPD, they answer an item correctly, and are able to attempt, 19–24% more attempts when the item is only delivered in one format compared to multiple formats. The probability of a

¹github.com/niall-twomey/zero-inflated-learner-models

user succeeding at an item is can be drastically effected by a poor learning quality.

3.2 Can NDCs be identified from interactions?

To investigate if users with a potential NDCs can be identified from the interactions, we have compared individuals mean performance in different subjects and on items with different delivery types (Fig 3). When Maths and English are compared (Fig 3 left), dyscalculic users have attempted more English items than Maths (large spike on ‘Not answered’). Additionally, when Maths is attempted, there is a lower success rate than in English (dip in ‘Correct’). Their performance in terms of ‘Incorrect’ counts in English and Maths are equivalent. However, this tally is achieved with 30% fewer attempts, indicating poor performance in Maths, further illustrating the effect of their NDC (*i.e.* 10/20 *vs.* 5/15). The most noticeable effects between read *vs.* listen DRT (Fig 3 middle) are seen by a clear increase in number of not answered items and decrease in the number of correct answers for ‘dyslexia’ and ‘dyslexia & SPD’ students. SPD students are unaffected by these DRTs. Comparing the ‘read & listen’ and ‘read’ delivery types (Fig 3 right), there are features seen with the dyslexia users, as above, but the SPD users now show a significant difference in performance, with large increases on ‘not answered’ and decreases on ‘correct’. So, by comparing individual students’ performance in different subjects and DRTs, it’s possible to identify the ND students and their condition. In practices, these comparisons could be used to identify what contexts a student may be struggling with, and additional support they may need.

3.3 Can a user’s true ability be estimated?

One aspect of ensuring each user gets suitable learning material is understanding their true ability. Fig 4 compares the performance of classical IRT and our IRT-ZILM model for parameter recovery. With IRT, most of the ability values are under-estimated, particularly for students with 1 or 2 NDCs (Fig 4a). Under-estimated ability makes sense given our expected inflated zero counts. However, the bias of under-estimated ability for ND students is concerning given that ND and NT abilities were drawn from the same distribution. On the other hand, IRT-ZILM is a much better estimator of true abilities (Fig 4b). Additionally, there is no obvious gap in ability estimates for students with NDCs compared to NT students. Table 2 summarises the predictive accuracy of the considered models. Although the performance of all models is approximately equivalent (only small gains for our approach) the lack of distorted recovered parameters may indicate stronger reliability of IRT-ZILM.

Table 3 summarise parameter estimation using Pearson and Spearman correlation coefficients, and have included linear KTM [39] (using contextual features) as another baseline. KTM, like IRT also under-estimates ND ability, and IRT-ZILM is a significantly better estimator of the true parameters.

3.4 Can learning quality be improved?

We explore the effect of actively selecting DRTs to improve LQFs and the number of successful learning attempts for ND students in Table 4. The table shows the potential that selecting the most suitable DRT can have on learning quality, with large lifts on students with 1 or 2 NDCs.

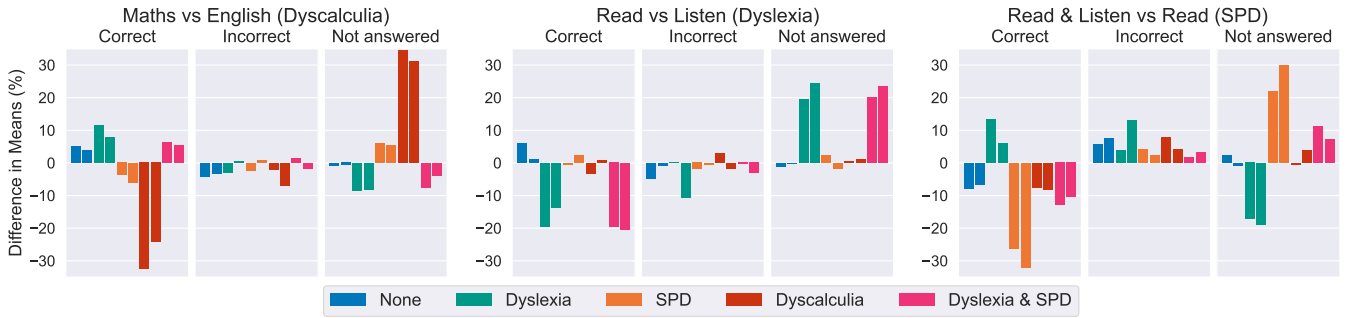


Figure 3: Performance differences for selected students across subject- and NDC-oriented contexts. Bar chart colour indicates NDCs. Large positive and negative values in the bar charts indicates that a group has been affected by the context. While every NDC is affected (indicated in parentheses in subfigure captions), no significant effects present for NT students.

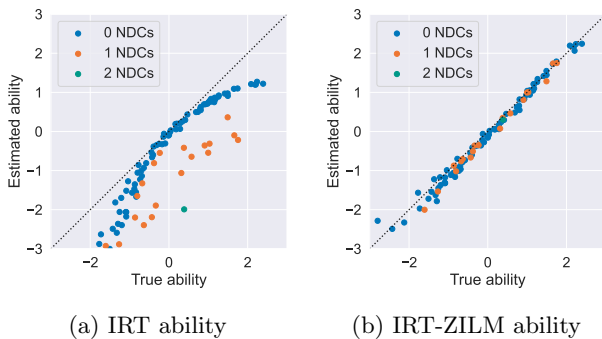


Figure 4: Scatter plots of true vs. estimated ability parameters from IRT and IRT-ZILM. Perfect estimation will place all points on diagonal. IRT is biased against ND students, while IRT-ZILM parameter estimation is very reliable.

Table 2: Predictive test metrics. Similar performance obtained with all models, though IRT-ZILM is slightly more performant than baselines.

Metric	IRT	KTM	IRT-ZILM
Accuracy	0.734	0.742	0.753
F1	0.559	0.567	0.583
NLL	0.513	0.499	0.494
Brier Score	0.170	0.166	0.163

3.5 How can this model be applied?

As already discussed, comparing user interactions in different contexts can identify students who may need additional support in specific areas. Often, high achieving ND students needs can be overlooked since their performance doesn't tend to require interventions. With IRT-ZILM, support/adaptions can be put in place early to enable them to reach their full potential since this model is less susceptible to the biases of traditional LMs. IRT-ZILM can be used to better estimate a students true ability, by adapting it to contexts and underestimating their DRTs preferences. This can help identify and explain causes for underperforming students. By understanding which DRTs a student struggles to engage with, alternative items can be provided to

Table 3: Pearson and Spearman correlation coefficients between true and recovered parameters. Values of 1 indicate perfect matches. IRT-ZILM parameter estimation is the most accurate across both metrics for all parameters.

	Pearson			Spearman		
	IRT	KTM	IRT-ZILM	IRT	KTM	IRT-ZILM
Ab	0.839	0.955	0.993	0.929	0.966	0.996
Diff	0.394	0.686	0.953	0.413	0.707	0.954
Disc	0.270	0.544	0.932	0.234	0.610	0.942

Table 4: Increase (and decrease) of learning opportunities obtained with active (and adversarial) DRT selection.

	1 NDC	2 NDCs
Baseline	0.391	0.123
Lift	1.432 ↑	1.898 ↑
Drop	0.248 ↓	0.014 ↓

help them reach their full potential. These insights can also be used by teachers to explore if the DRTs of their content can be expanded to create an accessible learning environment for all. Education traditionally has taken a one size fits all approach. By harnessing models that incorporate contextual understanding, learning can be tailored to each student, reaching many of those who may previously have felt dejected in learning, as their needs weren't being met.

4. CONCLUSIONS

Our application of zero-inflated models in learning contexts offers a rich simulation environment of neurodivergent conditions in question answering settings, unbiased evaluations of neurodivergent learners, encourages increased learning quality, and more reliably recovers unbiased ability parameters. On the basis of our successful results we believe that further study and exploration of zero-inflated learner models can yield an inclusive framework for equitable, explainable, and reliable learner models in diverse educational data mining contexts. Future work will expand on the experimentation to new contexts, and the model to new domains.

References

- [1] A. Abyaa, M. Khalidi Idrissi, and S. Bennani. Learner modelling: systematic review of the literature from the last 5 years. *Educational Technology Research and Development*, 67(5):1105–1143, 2019.
- [2] M. A. Barton and F. M. Lord. An upper asymptote for the three-parameter logistic item-response model*. *ETS Research Report Series*, 1981(1):i–8, 1981.
- [3] C. M. Bishop and N. M. Nasrabadi. *Pattern recognition and machine learning*, volume 4. Springer, 2006.
- [4] L. E. Boyd, K. Day, N. Stewart, K. Abdo, K. Lamkin, and E. Linstead. Leveling the playing field: Supporting neurodiversity via virtual realities. *Technology & Innovation*, 20(1-2):105–116, 2018.
- [5] Y. Chen, X. Li, J. Liu, and Z. Ying. Item response theory—a statistical framework for educational and psychological measurement. *arXiv preprint arXiv:2108.08604*, 2021.
- [6] A. T. Corbett and J. R. Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction*, 4(4):253–278, 1994.
- [7] J. Crisfield. *The Dyslexia Handbook, 1995*. British Dyslexia Association, 1995.
- [8] E. Ferrer, B. A. Shaywitz, J. M. Holahan, K. E. Marchione, R. Michaels, and S. E. Shaywitz. Achievement gap in reading is present as early as first grade and persists through adolescence. *The Journal of pediatrics*, 167(5):1121–1125, 2015.
- [9] A. Galiana-Simal, M. Vela-Romero, V. M. Romero-Vela, N. Oliver-Tercero, V. García-Olmo, P. J. Benito-Castellanos, V. Muñoz-Martinez, and L. Beato-Fernandez. Sensory processing disorder: Key points of a frequent alteration in neurodevelopmental disorders. *Cogent Medicine*, 7(1):1736829, 2020.
- [10] A. Gelman and J. Hill. *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press, 2006.
- [11] T. Gervet, K. Koedinger, J. Schneider, T. Mitchell, et al. When is deep learning the best approach to knowledge tracing? *Journal of Educational Data Mining*, 12(3):31–54, 2020.
- [12] J. A. Greene, L.-J. Costa, and K. Dellinger. Analysis of self-regulated learning processing using statistical models for count data. *Metacognition and Learning*, 6(3):275–301, 2011.
- [13] K. Holstein and S. Doroudi. Equity and artificial intelligence in education: Will” aided” amplify or alleviate inequities in education? *arXiv preprint arXiv:2104.12920*, 2021.
- [14] B. D. Homer and J. L. Plass. Using multiple data streams in executive function training games to optimize outcomes for neurodiverse populations. In *International Conference on Human-Computer Interaction*, pages 281–292. Springer, 2021.
- [15] M. Khajah, R. V. Lindsey, and M. C. Mozer. How deep is knowledge tracing? *arXiv preprint arXiv:1604.02416*, 2016.
- [16] M. Kohli and T. Prasad. Identifying dyslexic students by using artificial neural networks. In *Proceedings of the world congress on engineering*, volume 1, pages 1–4, 2010.
- [17] D. Lambert. Zero-inflated poisson regression, with an application to defects in manufacturing. *Technometrics*, 34(1):1–14, 1992.
- [18] J. I. Lee and E. Brunskill. The impact on individualizing student models on necessary practice opportunities. *International Educational Data Mining Society*, 2012.
- [19] C.-S. Li. Identifiability of zero-inflated poisson models. *Brazilian Journal of probability and Statistics*, 26(3):306–312, 2012.
- [20] W.-W. Liao, R.-G. Ho, Y.-C. Yen, and H.-C. Cheng. The Four-Parameter Logistic Item Response Theory Model As a Robust Method of Estimating Ability Despite Aberrant Responses. *Social Behavior and Personality: an international journal*, 40(10):1679–1694, Nov. 2012.
- [21] Q. Liu, S. Shen, Z. Huang, E. Chen, and Y. Zheng. A survey of knowledge tracing. *arXiv preprint arXiv:2105.15106*, 2021.
- [22] A. Lollini. Brain equality: Legal implications of neurodiversity in a comparative perspective. *NYUJ Int’l L. & Pol.*, 51:69, 2018.
- [23] B. E. Magnus and Y. Liu. A zero-inflated box-cox normal unipolar item response model for measuring constructs of psychopathology. *Applied psychological measurement*, 42(7):571–589, 2018.
- [24] V. Mandalapu, J. Gong, and L. Chen. Do we need to go deep? knowledge tracing with big data. *arXiv preprint arXiv:2101.08349*, 2021.
- [25] M. Marras, L. Boratto, G. Ramos, and G. Fenu. Equality of learning opportunity via individual fairness in personalized recommendations. *International Journal of Artificial Intelligence in Education*, pages 1–49, 2021.
- [26] C. Mejia, S. Gomez, L. Mancera, and S. Taveneau. Inclusive learner model for adaptive recommendations in virtual education. In *2017 IEEE 17th International Conference on advanced learning technologies (ICALT)*, pages 79–80. IEEE, 2017.
- [27] A. Menon, B. Van Rooyen, C. S. Ong, and B. Williamson. Learning from corrupted binary labels via class-probability estimation. In *International conference on machine learning*, pages 125–134. PMLR, 2015.
- [28] S. Pandey and G. Karypis. A self-attentive model for knowledge tracing. *arXiv preprint arXiv:1907.06837*, 2019.

- [29] T. Papathoma, R. Ferguson, F. Iniesto, I. Rets, D. Vogiatzis, and V. Murphy. Guidance on how learning at scale can be made more accessible. In *Proceedings of the seventh ACM conference on learning@ Scale*, pages 289–292, 2020.
- [30] A. Patrick. *The Memory and Processing Guide for Neurodiverse Learners: Strategies for Success*. Jessica Kingsley Publishers, 2020.
- [31] R. Pelánek. Bayesian knowledge tracing, logistic models, and beyond: an overview of learner modeling techniques. *User Modeling and User-Adapted Interaction*, 27(3):313–350, 2017.
- [32] M. Perello-Nieto, R. Santos-Rodriguez, D. Garcia-Garcia, and J. Cid-Sueiro. Recycling weak labels for multiclass classification. *Neurocomputing*, 400:206–215, 2020.
- [33] C. Piech, J. Bassen, J. Huang, S. Ganguli, M. Sahami, L. J. Guibas, and J. Sohl-Dickstein. Deep knowledge tracing. *Advances in neural information processing systems*, 28, 2015.
- [34] J. L. Plass and S. Pawar. Toward a taxonomy of adaptivity for learning. *Journal of Research on Technology in Education*, 52(3):275–300, 2020.
- [35] E. S. Roemmele. A flexible zero-inflated poisson regression model. 2019.
- [36] R. S. Shaley, J. Auerbach, O. Manor, and V. Gross-Tsur. Developmental dyscalculia: prevalence and prognosis. *European child & adolescent psychiatry*, 9(2):S58–S64, 2000.
- [37] J. Singer. Why can’t you be normal for once in your life? from a problem with no name to the emergence of a new category of difference. *Disability discourse*, pages 59–70, 1999.
- [38] N. Smits, O. Öğreden, M. Garnier-Villarreal, C. B. Terwee, and R. P. Chalmers. A study of alternative approaches to non-normal latent trait distributions in item response theory models used for health outcome measurement. *Statistical Methods in Medical Research*, 29(4):1030–1048, 2020.
- [39] J.-J. Vie and H. Kashima. Knowledge tracing machines: Factorization machines for knowledge tracing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 750–757, 2019.
- [40] L. Wang. Irt–zip modeling for multivariate zero-inflated count data. *Journal of Educational and Behavioral Statistics*, 35(6):671–692, 2010.
- [41] Z.-H. Zhou. A brief introduction to weakly supervised learning. *National science review*, 5(1):44–53, 2018.

Appendix

This section gives supplementary details of our proposed model, see Sec 2.1.

Delivery and Response Weakening

We adapt learner models for NDCs by taking inspiration from techniques used in Weakly Supervised Machine Learning (WSML) [41]. Our approach is to model the interplay between item DRTs and NDCs. Let a binary random variable be drawn from a Bernoulli distribution, $y \sim \text{Ber}(p)$, and let us assume that a label flipping process acts upon y and this results in observations of the corrupted labels, \tilde{y} . The mixing matrix, M , is defined as follows:

$$\begin{pmatrix} 1 - q_0 & q_1 \\ q_0 & 1 - q_1 \end{pmatrix} = \begin{pmatrix} Pr(\tilde{Y} = 1|Y = 1) & Pr(\tilde{Y} = 1|Y = 0) \\ Pr(\tilde{Y} = 0|Y = 1) & Pr(\tilde{Y} = 0|Y = 0) \end{pmatrix}$$

The $q_{\tilde{y}}$ variables can be selected using prior knowledge and assumptions on the data distributions [27, 32]. In our setting, we are interested particularly in the contexts when learning of ND students is being sabotaged by the environment, *i.e.* q_0 . We therefore model q_0 (previously introduced as a global parameter) and parameterise it with ND, LQF and interaction features.

IRT-based Zero-Inflated Learner Model

Our IRT-ZILM merges LMs and ZILM as follows:

$$Pr(Y = y | \mathbf{x}) = \begin{cases} \pi(\mathbf{x}_\pi) + (1 - \pi(\mathbf{x}_\pi))(1 - p(\mathbf{x}_p)) & \text{if } y = 0 \\ (1 - \pi(\mathbf{x}_\pi))p(\mathbf{x}_p) & \text{if } y = 1 \end{cases}$$

where π and p from Eqn (1) are now functions leveraging ND/LQF/content features (\mathbf{x}_π) and LM/collaborative features (\mathbf{x}_p).

By separating the functional contribution of confounders (π) and ability (p) in IRT-ZILM, we hope to unambiguously decouple these aspects from each other and improve interpretability and explainability. The model is learnt by gradient descent of negative log likelihood of the training data to optimise all parameters. In WSML it is common to learn in a two-step process, for example, by iteratively fixing and optimising IRT and weak label weights [32].

ZIMs have been used to account for excess zeros in many counting tasks using Poisson and negative binomial models [17, 40, 38, 23], and in learning analytics as statistical counting models in self-regulated learning [12]. An important property of statistical models is identifiability as it allows for the precise estimation of the values of its parameters [10, Sec 4.5]. Parallel theoretical analysis has considered identifiability of the counting model parameters [35] and the mixture components [19]. It is worth noting that IRT also suffers from identifiability problems (*c.f.* [5, p.6] and [10, Sec 14.]) but using priors or regularisation can alleviate these.

As far as we are aware, this is the first work to incorporate ZIM in this manner. Choosing IRT as the base LM in IRT-ZILM over alternative options is motivated for several reasons. Firstly, IRT is well-understood and simple to interpret, and using this model as a platform to demonstrate new properties of equity in this early work carries the same benefits. Secondly, BKT is known to have over- and under-estimation problems [6, 18] which may muddle our under-

standing of equity for ND students. Additionally, several technical hurdles need to be overcome, notably adaptation for contextualised individualisation in mixed graphs. Finally, although DKT [33] models can probably learn latent representations that correlate to DRT preferences, this is at the expense of control and interpretation of the effects.

Extra Results

Fig 5 shows this effect for four user/item pairs. For example, the first student should be 60% (orange) successful on this item, however, their LQF is 0.25 (blue), so their success rate drops to 15% (green). Therefore, LQF can be interpreted as a measure of the contextual inequity in these settings.

Although the purpose of this research is to provide equitable estimates of student ability and to provide enabling technology that selects the most appropriate DRT for students, we note that we may also identify students that need additional support in specific areas by recognising potentially unidentified NDCs. We can approach this by creating two models: let \mathcal{M}_0 be the model for a student’s reported NDC state (the ‘null’ model), and let \mathcal{M}_1 be a model trained on data assuming an alternative NDC state (the ‘alternative’ model). Since we have already shown that metrics and likelihood is improved with IRT-ZILM, a statistical hypothesis test can be performed on both likelihoods to determine whether the null or alternative NDC offers a better explanation of data. We leave further elaboration of this approach as future work since it is outside the scope of our direct objectives.




Figure 5: IRT predictions (f) combining with complementary LQFs (π) producing IRT-ZILM predictions (p).

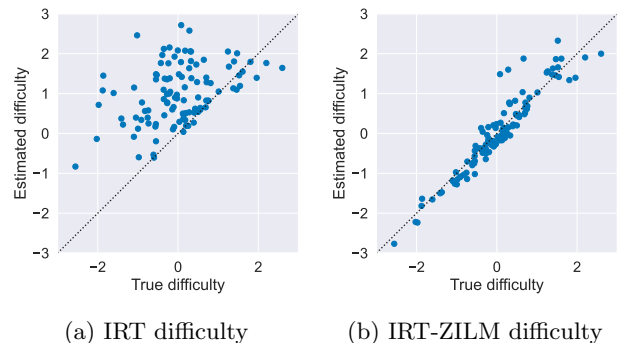


Figure 6: Scatter plots of true *vs.* estimated difficulty parameters from IRT and IRT-ZILM. Perfect estimation will place all points on diagonal. Estimation from IRT-ZILM is significantly more accurate than IRT.