
Fair Knowledge Tracing in Second Language Acquisition

Weitao Tang
Monash University
Weitao.Tang1@monash.edu

Dr. Guanliang Chen*
Monash University
guanliang.chen@monash.edu

Shuaishuai Zu
Renmin University of China
Zushuaishuai@gmail.com

Jiangyi Luo
Australian National University
Jiangyi.Luo@anu.edu.au

Abstract

In the domain of second-language acquisition, predictive modeling serves as a pivotal tool for facilitating educators in implementing diversified teaching strategies, thereby garnering extensive research attention. Despite the prevalent focus on model accuracy in most existing studies, the exploration into model fairness remains substantially underexplored. Model fairness pertains to the equitable treatment of different groups by machine learning algorithms. It ensures that the model's predictions do not exhibit unintentional biases against certain groups based on attributes such as gender, ethnicity, age, or other potentially sensitive characteristics. In essence, a fair model should produce outcomes that are impartial and do not perpetuate existing prejudices, ensuring that no group is systematically disadvantaged. In this research, we evaluate the fairness of two predictive models based on second-language learning, utilizing three tracks from the Duolingo dataset: en_es (English learners who speak Spanish), es_en (Spanish learners who speak English), and fr_en (French learners who speak English). We measure

- (i) algorithmic fairness among different clients such as iOS, Android and Web and
- (ii) algorithmic fairness between developed countries and developing countries.

Our findings indicate:

- 1) Deep learning exhibits a marked advantage over machine learning when applied to knowledge tracing based on second language acquisition, owing to its heightened accuracy and fairness.
- 2) Both machine learning and deep learning algorithms exhibit a noticeable bias favoring mobile users over their non-mobile counterparts.
- 3) Compared to deep learning algorithms, machine learning algorithms showcase a more pronounced bias against developing countries.
- 4) To strike a balanced approach in terms of fairness and accuracy, deep learning is identified as being more apt for making predictions in the en_es and es_en tracks, whereas machine learning emerges as a more suitable option for the fr_en.

This study underscores the necessity to delve deeper into the realms of fairness in predictive models, ensuring equitable educational strategies across diverse clients and countries.

*Corresponding author: guanliang.chen@monash.edu

1 Introduction

The historical reliance on traditional metrics like student performance, feedback, and participation to shape educational methodologies was fraught with potential inaccuracies [1]. Such inaccuracies mainly arise because humans cannot process the vast amount of data generated during the educational process like machines can. A deeper delve suggests that these inaccuracies might be influenced by a range of psychological factors [2]. For instance, educators might be influenced by

1. The limits of short-term memory.
2. Feelings of fatigue or boredom.
3. A positivity bias, where there is a predisposition towards favoring positive feedback from students. (This means that an educator might harbor an undue expectation for positive comments, thereby skewing their interpretation of student feedback)

As a result, a single poor examination outcome might prompt an educator to simplify the test content, risking a dilution of educational standards. To mitigate these issues, in 1994, Corbett and Anderson introduced the concept of "knowledge tracing" [3]. The crux of this approach was to harness computational technology to model various student indicators such as grades, feedback, participation, and other features. By doing so, predictions about students' future performance could be made, enabling the crafting of more appropriate educational strategies.

Employing computational techniques effectively reduces the chances of human-induced errors, enhancing the precision of outcomes. In recent times, educational technology increasingly incorporates artificial intelligence to employ data and predictive models [4, 5, 6, 7, 8]. This provides tailored support and insights for students, educators, and administrators alike [9, 10]. This process closely resembles an adaptive system. As students accumulate knowledge, the system makes predictions regarding their understanding and offers varied learning materials tailored to their needs [11].

As outlined by [12], the Intelligent Tutoring System (ITS Figure 1) functions as an adaptive feedback mechanism. The ITS poses questions to students and, based on their responses, discerns their knowledge state. This knowledge state offers insight into the student's proficiency in various skills. For instance, if a student consistently demonstrates over 90% accuracy in addition tasks, one could infer mastery over the skill of addition. However, the same student might struggle with problems combining addition and subtraction, indicating a lack of proficiency in the latter. Certain skills might be interdependent; for example, mastering calculus presupposes proficiency in basic arithmetic operations like addition and subtraction. As learning progresses, students might forget certain concepts, with more complex skills generally having a higher forgetting rate.

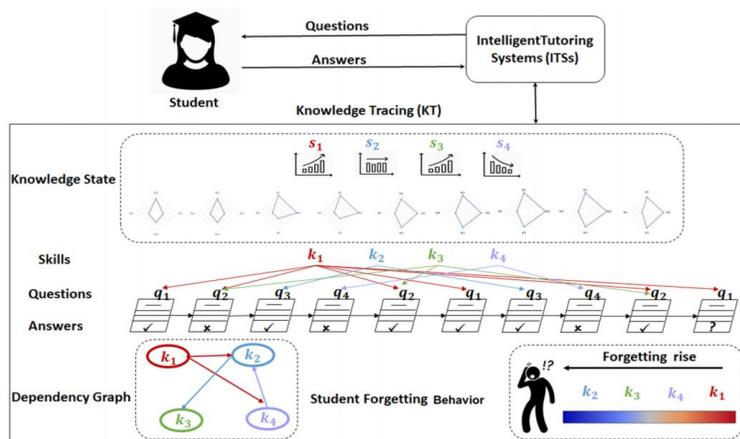


Figure 1: Intelligent tutoring system [12].

While contemporary knowledge tracing algorithms have achieved commendable accuracy, there remains a pressing need to explore their fairness. For instance, gender discrepancies could lead to algorithmic biases [13, 14, 15, 16]. Similar biases, such as racial prejudices, have been observed in

criminal justice record systems [17]. As concerns over fairness in predictive educational models grow [18], ensuring algorithmic equity becomes pivotal for creating an inclusive educational environment. For diverse student populations, investigating the influence of factors like native language background on algorithmic fairness is crucial.

Second language acquisition is a vital area of study, shaping our understanding of cognitive processes, cultural exchange, and global communication. The ability to acquire a second language not only bridges cultural and geographical divides but also offers individuals a competitive edge in today’s globalized world. It fosters cognitive benefits, including problem-solving skills, creativity, and memory enhancements. With the increasing reliance on technology to aid language learning, ensuring the fairness and effectiveness of these tools is paramount.

The motivation behind this research arises from the pressing need to ensure equity in technological tools employed for second language acquisition. While machine learning and deep learning models promise unprecedented benefits in tailoring learning experiences, there’s a pressing need to scrutinize potential biases they may harbor. The objective of this study is to delve into fairness concerns in machine learning and deep learning models using the Duolingo dataset. Specifically, we intend to utilize the Gradient Boosted Decision Trees (GBDT) algorithm and the Multi-task Learning algorithm—both of which have demonstrated remarkable results in the Duolingo shared task². Our primary goal is to ascertain the extent to which these algorithms exhibit algorithmic biases. We investigate biases concerning students’ demographic attributes, particularly "country" and "client." Here, "country" is categorized into developed and developing nations, highlighting socio-economic disparities. Simultaneously, "client" delineates the platforms students use for learning, encompassing iOS, Android, and Web interfaces. By understanding these biases, this research endeavors to offer insights that can direct developers in forging more equitable language learning tools and platforms.

2 Related Work

In the scope of this project, a compilation of 16 research articles was undertaken, all of which explore the application of knowledge tracing on the Duolingo dataset. These research contributions can be broadly bifurcated into three primary categories: those employing Traditional Machine Learning-based Algorithms (7 articles), and those utilizing Deep Learning-based Algorithms (9 articles), with the objective of predicting student performance metrics. A notable 13 of these papers presented predictions across all three tracks, while a pair exclusively targeted the English track, and a singular article was dedicated to investigating both the Spanish and French tracks. Subsequent sections will delve into a meticulous analysis of the models adopted in these articles, along with an exploration of their respective performance metrics in a detailed manner.

2.1 Methodological Approaches

To better compare the models in these 16 papers, two metrics will be introduced. One is the F1 score, which represents the harmonic mean of Precision and Recall. The other is AUC (Area Under the Curve) [19]. AUC is the area under the ROC (Receiver Operating Characteristic) curve, which primarily describes the relationship between the False Positive Rate and the True Positive Rate. Both values are closer to 1, indicating a better model.

2.2 Comparison between all Algorithms

Table 1 illustrates that models integrating deep learning or combining it with machine learning tend to outperform, whereas standalone machine learning models might fall short in certain tasks. This points towards a necessity for neural networks in Second Language Acquisition (SLA) tasks to delve deeper into feature extraction, considering the Duolingo dataset may present complexities beyond the capacity of traditional machine learning models. Notably, [20] delivered superior outcomes by employing an encoder-decoder architecture and multi-task learning, which facilitated the use of shared representation layers and the extraction of features even in data-scarce scenarios. Conversely, the LM-KT model presented in [21] exhibited subpar performance, potentially due to intrinsic model constraints and a lack of comprehensive model optimization by the authors. Hence, the selection of an appropriate model emerges as pivotal in the tracing of knowledge in SLA. Both high-achieving and

²Data for the 2018 Duolingo Shared Task on Second Language Acquisition Modeling (SLAM)

underachieving models have vitally informed the evolution of SLA knowledge tracing. In summary, within the realm of Machine Learning (ML) algorithms, the most notable results were achieved by [22], while [20] claimed the pinnacle of performance among Deep Learning (DL) algorithms, thanks to its utilization of multi-task learning.

Table 1: Comparison of all model’s performance

(a) English Track			(b) Spanish Track		
Literature	AUC↓	F1↓	Literature	AUC↓	F1↓
[20]●	.864	.564	[20]●	.839	.530
[23]●	.863	.564	[23]●	.838	.531
[24]●	.861	.561	[24]●	.838	.530
[25]●	.861	.559	[25]●	.835	.524
[22]△	.859	.468	[22]△	.835	.420
[26]△	.854	-	[28]●	.824	.439
[27]●	.853	.479	[29]△	.818	.390
[28]●	.848	.476	[30]●	.807	.435
[29]△	.846	.414	[31]△	.803	.375
[30]●	.841	.479	[32]△	.801	.344
[31]△	.829	.424	[34]△	.791	.452
[32]△	.821	.389	[33]△	.790	.338
[33]△	.821	.376	[35]●	.788	.306
[34]△	.817	.462	[21]●	.750	-
[35]●	.815	.329	[27]●	-	-
[21]●	-	-	[26]△	-	-

(c) French Track

Literature	AUC↓	F1↓
[20]●	.860	.579
[23]●	.859	.575
[24]●	.857	.573
[25]●	.854	.569
[22]△	.854	.493
[29]△	.843	.487
[28]●	.839	.502
[30]●	.835	.508
[31]△	.823	.442
[32]△	.815	.415
[34]△	.813	.502
[33]△	.811	.431
[35]●	.809	.406
[21]●	.730	-
[27]●	-	-
[26]△	-	-

△ means using Machine Learning Algorithm, ● means using Deep Learning Algorithm

3 Experiment

3.1 Experimental Environment and Set Up

All experiments were conducted on a high-performance Alienware X17R2 laptop. The specifications of the device are as follows:

Device: Alienware X17R2 Laptop

CPU: Intel i9-12900HK

Memory: 64GB DDR5

The experiments utilized TensorFlow for both training the models on the CPU.

To fully explore the fairness of algorithms on the Duolingo dataset, we selected the best-performing GBDT algorithm in machine learning [22] and the best-performing Multi-task learning in deep learning [20] from Table 1. Based on this, we further explore the accuracy and fairness of these two algorithms.

Accuracy metrics Consistent with previous knowledge tracing modeling based on second language acquisition, we employed two metrics to gauge the predictive accuracy of knowledge tracing: Area Under the Curve (AUC) and F1 score.

Fairness metrics The study [18] is noted to be the inaugural exploration into identifying apt metrics for evaluating the fairness of predictive models in educational research. Specifically, a metric named Absolute Between-group ROC Area (ABROCA) was introduced, aiming to measure the predictive unfairness across various demographic groups by calculating the definite integral between the ROC curves of two observed groups. Significantly, ABROCA possesses two advantages:

1. It accounts for performance disparities across the entire threshold spectrum, superseding other fixed-threshold approaches.
2. It can be easily calculated from predictive results without the need for additional data or metrics.

Thus, this metric was utilized in our research. It’s crucial to note that a lower ABROCA value signifies reduced algorithmic unfairness in the predictive model. We executed training using GBDT and multi-task learning across three tracks: `en_es`, `es_en`, and `fr_en`. The data post-training were categorized in two dimensions:

1. By countries, segmented into:
 - Developed Countries
 - Developing Countries
2. By client, segmented into:
 - Android
 - iOS
 - Web

Testing was subsequently conducted on the three tracks, employing AUC and F1 score as metrics to evaluate precision for developed and developing countries, and Android, iOS, and Web clients respectively. Concurrently, ABROCA was utilized to assess fairness.

3.2 Accuracy

3.2.1 Accuracy on three tracks (`en_es`, `es_en`, `fr_en`)

Table 2: Accuracy on Different Tracks

Track	Model	F1 Score	AUC
EN_ES	Multi-task learning	0.554	0.858
	GBDT	0.469	0.856
ES_EN	Multi-task learning	0.519	0.830
	GBDT	0.421	0.833
FR_EN	Multi-task learning	0.571	0.855
	GBDT	0.493	0.851

1. Dataset Size Consideration:

- The *fr_en* track corresponds to a smaller dataset compared to *en_es* and *es_en*. The complex nature of the Multi-task learning model may not be as beneficial for smaller datasets, resulting in suboptimal performance.

2. Training Time:

- The Multi-task learning model, trained using TensorFlow-CPU, took approximately 11 days, making it a time-intensive approach. In contrast, the GBDT model took just an hour. This stark difference in training time suggests that while Multi-task learning may offer performance benefits in some scenarios, it comes at the cost of significantly increased computational time.

3. Performance Across Tracks:

- *EN_ES Track*: The Multi-task learning model outperformed GBDT in F1 score for this track. However, their AUCs are very similar, suggesting both models are comparable in terms of ranking the predictions.
- *ES_EN Track*: While the Multi-task learning model still leads in F1 score, the gap is narrower. Both models' AUCs remain close.
- *FR_EN Track*: Multi-task learning exceeds GBDT in F1 score. However, considering the significantly shorter training time of GBDT and its competitive performance, GBDT might be a more feasible choice, especially for the smaller *FR_EN* dataset.

3.2.2 Accuracy on Client (ios, android, web)

Table 3: Accuracy on Different Clients

Track	Model	F1 Score	AUC
CLIENT_IOS_EN_ES	Multi-task learning	0.535	0.86
	GBDT	0.434	0.86
CLIENT_ANDROID_EN_ES	Multi-task learning	0.551	0.859
	GBDT	0.468	0.858
CLIENT_WEB_EN_ES	Multi-task learning	0.584	0.849
	GBDT	0.507	0.842
CLIENT_IOS_ES_EN	Multi-task learning	0.528	0.831
	GBDT	0.431	0.833
CLIENT_ANDROID_ES_EN	Multi-task learning	0.524	0.84
	GBDT	0.427	0.841
CLIENT_WEB_ES_EN	Multi-task learning	0.475	0.8
	GBDT	0.371	0.811
CLIENT_IOS_FR_EN	Multi-task learning	0.421	0.727
	GBDT	0.498	0.858
CLIENT_ANDROID_FR_EN	Multi-task learning	0.4	0.719
	GBDT	0.474	0.857
CLIENT_WEB_FR_EN	Multi-task learning	0.44	0.713
	GBDT	0.504	0.824

• Performance Across Clients:

- *iOS*: Multi-task learning consistently shows a higher F1 score across all tracks for iOS users. This might indicate that iOS users, who are often associated with higher spending power, benefit more from the complexity of the Multi-task learning model.
- *Android*: The performance between the models is closer on this platform. For some tracks, Multi-task learning has a slight edge in F1 score.
- *Web*: The F1 score of the Multi-task learning model is generally higher than that of GBDT, but GBDT sometimes leads in AUC, particularly for the *FR_EN* track.

3.2.3 Accuracy on Country (developed countries and developing countries)

Table 4: Accuracy on Different Countries

Track	Model	F1 Score	AUC
COUNTRY_DEVELOPED_EN_ES	Multi-task learning	0.553	0.863
	GBDT	0.476	0.865
COUNTRY_DEVELOPING_EN_ES	Multi-task learning	0.555	0.856
	GBDT	0.467	0.853
COUNTRY_DEVELOPED_ES_EN	Multi-task learning	0.525	0.832
	GBDT	0.428	0.835
COUNTRY_DEVELOPING_ES_EN	Multi-task learning	0.480	0.818
	GBDT	0.368	0.820
COUNTRY_DEVELOPED_FR_EN	Multi-task learning	0.427	0.726
	GBDT	0.497	0.851
COUNTRY_DEVELOPING_FR_EN	Multi-task learning	0.410	0.729
	GBDT	0.491	0.854

• **Performance Across Countries:**

- *Developed Countries:* The performance of both models is closely matched in developed countries. Interestingly, for the *FR_EN* track, GBDT has a distinct advantage in AUC.
- *Developing Countries:* In developing countries, Multi-task learning generally leads in F1 score. However, their AUCs are again closely matched. This suggests that while the more complex model may predict certain classes better, both models rank their predictions similarly.

3.3 Fairness

Understanding the model’s performance in terms of fairness is especially crucial for platforms like Duolingo, where feedback directly influences learners’ progress and motivation. Any disparity can lead to differing user experiences, potentially leading to decreased trust in the platform or even reduced learning outcomes for a particular user group.

3.3.1 Fairness on Client

Figure 2 shows the ABROCA between iOS and Android in *en_es* track for GBDT. Figure 3 shows the ABROCA between iOS and Android in *en_es* track for Multi-task learning.

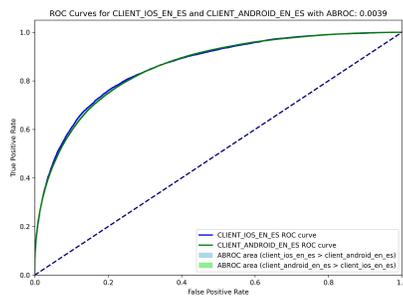


Figure 2: Fairness between iOS and Android in *en_es* track for GBDT

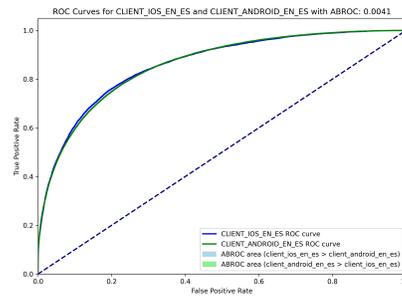


Figure 3: Fairness between iOS and Android in *en_es* track for Multi-task learning

To better elucidate the differences in ABROCA with respect to client-side performance between GBDT and Multi-task learning, we present the following Table 5.

Overall, the ABROCA on the Android and iOS platforms is noticeably smaller than in the other categories across the three linguistic tracks, highlighting the inherent similarity between the Android

Table 5: Fairness on Client

Client 1	Client 2	Track	GBDT(ABROCA)	Multi-task learning(ABROCA)
Android	Web	en_es	0.0226	0.0202
ios	Android	en_es	0.0039	0.0041
ios	Web	en_es	0.0236	0.0198
Android	Web	es_en	0.0307	0.0406
ios	Android	es_en	0.0082	0.0078
ios	Web	es_en	0.0227	0.0318
Android	Web	fr_en	0.0333	0.0073
ios	Android	fr_en	0.002	0.0091
ios	Web	fr_en	0.0341	0.0138

and iOS mobile platforms in the context of algorithmic fairness. Pertinently, both the GBDT and Multi-task learning algorithms exhibit discernible biases against Web users. Despite the notable that the ABROCA of the GBDT end is larger than that of the Multi-task learning, indicating a somewhat larger discriminatory impact against non-mobile users, it’s imperative to note that neither algorithm provides an entirely impartial solution when considering platform-based disparities. Thus, while differences in the extent of bias exist between the two algorithms, our analysis leads us to the indispensable conclusion that both GBDT and Multi-task learning algorithms present prejudicial outcomes against non-mobile (Web) users, thereby necessitating a further exploration into mechanisms that can mitigate such biases to ensure fair and equitable knowledge tracing across all platforms.

3.3.2 Fairness on Country

Figures 4 shows the ABROCA between developed country and developing country in en_es track for GBDT. Figure 5 shows the ABROCA between developed country and developing country in en_es track for Multi-task learning. To better elucidate the differences in ABROCA with respect to

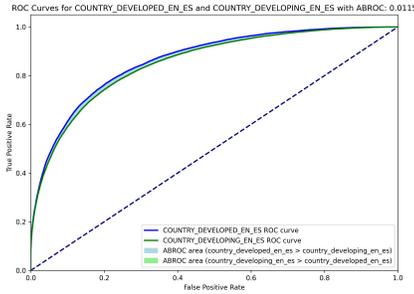


Figure 4: Fairness between developed country and developing country in en_es track for GBDT

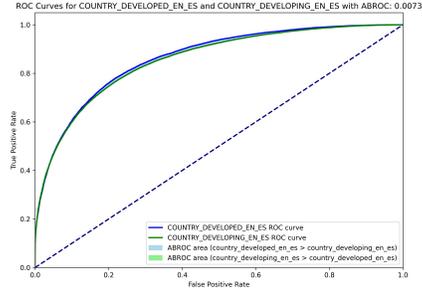


Figure 5: Fairness between developed country and developing country in en_es track for Multi-task learning

country-side performance between GBDT and Multi-task learning, we present the following table 6.

Table 6: Fairness on Country(developed country vs developing country)

Track	GBDT(ABROCA)	Multi-task learning(ABROCA)
en_es	0.0115	0.0073
es_en	0.0154	0.0139
fr_en	0.0026	0.0062

We can observe that Multi-task learning demonstrates greater fairness than GBDT in the en_es and es_en tracks, while it exhibits less fairness in the fr_en track. This indicates that for the fr_en track, GBDT possesses superior accuracy and fairness. Conversely, for the en_es and es_en tracks,

Multi-task learning holds a greater advantage. Viewing holistically, GBDT exhibits more bias against developing countries compared to Multi-task learning.

4 Discussion and Conclusion

This paper examines the fairness and accuracy of two knowledge tracing algorithms, GBDT and Multi-task learning, in the domain of second language acquisition across three linguistic tracks: `en_es`, `es_en`, and `fr_en`. Our comprehensive analysis reveals key insights into the interplay between algorithm choice and its impact on fairness and accuracy.

Firstly, our findings underscore that deep learning, exemplified here by Multi-task learning, tends to be more apt for knowledge tracing based on second language acquisition than machine learning algorithms like GBDT, owing predominantly to its higher accuracy and fairness, especially in the `en_es` and `es_en` tracks. Secondly, a discernible discriminatory tendency against non-mobile end-users was observed in both machine learning and deep learning algorithms. Furthermore, our third finding points towards a more pronounced bias inherent in machine learning algorithms, particularly towards developing countries, as opposed to deep learning algorithms.

To assess the potential algorithmic bias towards developing countries and users of web and android platforms (excluding iOS), we employed ABROCA as our fairness measurement tool. The fairness and accuracy analysis starkly illuminated the superiority of Multi-task learning in terms of algorithmic fairness compared to the GBDT algorithm.

Despite Multi-task learning overwhelmingly emerging as the preferable option due to its enhanced accuracy and algorithmic fairness, our fourth finding elucidates a nuanced approach to algorithm selection across different linguistic tracks. While Multi-task learning would be the optimal choice for `en_es` and `es_en` due to its fairness and predictive proficiency, for `fr_en` track, the GBDT algorithm is recommended due to its superior accuracy and fairness, despite Multi-task learning's broader applicability and success.

In summary, while Multi-task learning broadly demonstrates enviable aptitude in knowledge tracing, especially when balancing fairness and accuracy, the GBDT algorithm's selective applicability, such as in `fr_en` track scenarios, indicates that a tailored, context-specific approach to algorithm selection offers an optimal strategy in ensuring equitable and accurate knowledge tracing across varied linguistic and platform-specific contexts.

5 Limitations

We acknowledge several limitations in our study. Firstly, our analysis was confined to a single predictive task: knowledge tracing in second language acquisition, utilizing only the Duolingo dataset. This limitation restricts the generalizability of our findings across different datasets and contexts. To enhance the universality and robustness of our results, future studies should explore a broader range of machine learning and neural network algorithms and apply them to other predictive tasks in varied datasets. Additionally, our examination of demographic differences was limited to two specific student groups: Client and Country. Future research should extend this analysis to include other demographic factors, such as age and gender, to provide a more comprehensive evaluation of algorithmic fairness across diverse student populations.

References

- [1] Margaret Price, Karen Handley, Jill Millar, and Berry O'donovan. Feedback: all that effort, but what is the effect? *Assessment & Evaluation in Higher Education*, 35(3):277–289, 2010.
- [2] Phil F Culverhouse, Robert Williams, Beatriz Reguera, Vincent Herry, and Sonsoles González-Gil. Do experts make mistakes? a comparison of human and machine identification of dinoflagellates. *Marine ecology progress series*, 247:17–25, 2003.
- [3] Albert T Corbett and John R Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction*, 4:253–278, 1994.
- [4] Kenneth Holstein, Bruce M McLaren, and Vincent Alevén. Intelligent tutors as teachers' aides: exploring teacher needs for real-time analytics in blended classrooms. In *Proceedings of the seventh international learning analytics & knowledge conference*, pages 257–266, 2017.
- [5] Francesc Pedro, Miguel Subosa, Axel Rivas, and Paula Valverde. Artificial intelligence in education: Challenges and opportunities for sustainable development. 2019.
- [6] Ido Roll and Ruth Wylie. Evolution and revolution in artificial intelligence in education. *International Journal of Artificial Intelligence in Education*, 26:582–599, 2016.
- [7] Rose Luckin and Wayne Holmes. Intelligence unleashed: An argument for ai in education. 2016.
- [8] Marie Bienkowski, Mingyu Feng, and Barbara Means. Enhancing teaching and learning through educational data mining and learning analytics: An issue brief. *Office of Educational Technology, US Department of Education*, 2012.
- [9] RS Baker and PS Inventado. Educational data mining and learning analytics. en ja larusson & b. white (eds.), *learning analytics: From research to practice* (pp. 61-75), 2014.
- [10] Rosemary Luckin and Mutlu Cukurova. Designing educational technologies in the age of ai: A learning sciences-driven approach. *British Journal of Educational Technology*, 50(6):2824–2838, 2019.
- [11] John F Pane, Daniel F McCaffrey, Mary Ellen Slaughter, Jennifer L Steele, and Gina S Ikemoto. An experiment to evaluate the efficacy of cognitive tutor geometry. *Journal of Research on Educational Effectiveness*, 3(3):254–281, 2010.
- [12] Ghodai Abdelrahman, Qing Wang, and Bernardo Nunes. Knowledge tracing: A survey. *ACM Comput. Surv.*, 55(11), feb 2023.
- [13] Amit Datta, Michael Carl Tschantz, and Anupam Datta. Automated experiments on ad privacy settings: A tale of opacity, choice, and discrimination. *arXiv preprint arXiv:1408.6491*, 2014.
- [14] Tom Simonite. Probing the dark side of google's ad-targeting system. *MIT Technology Review*, 2015.
- [15] Rodrigo Duran, Lassi Haaranen, and Arto Hellas. Gender differences in introductory programming: Comparing moocs and local courses. In *Proceedings of the 51st ACM Technical Symposium on Computer Science Education*, pages 692–698, 2020.
- [16] Saif Rayyan, Daniel T Seaton, John Belcher, David E Pritchard, and Isaac Chuang. Participation and performance in 8.02 x electricity and magnetism: The first physics mooc from mitx. *arXiv preprint arXiv:1310.3173*, 2013.
- [17] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine bias. In *Ethics of data and analytics*, pages 254–264. Auerbach Publications, 2022.
- [18] Josh Gardner, Christopher Brooks, and Ryan Baker. Evaluating the fairness of predictive student models through slicing analysis. In *Proceedings of the 9th international conference on learning analytics & knowledge*, pages 225–234, 2019.
- [19] Andrew P Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern recognition*, 30(7):1145–1159, 1997.
- [20] Yong Hu, Heyan Huang, Tian Lan, Xiaochi Wei, Yuxiang Nie, Jiarui Qi, Liner Yang, and Xian-Ling Mao. Multi-task learning for low-resource second language acquisition modeling. In *Web and Big Data: 4th International Joint Conference, APWeb-WAIM 2020, Tianjin, China, September 18-20, 2020, Proceedings, Part I 4*, pages 603–611. Springer, 2020.

- [21] Megha Srivastava and Noah Goodman. Question generation for adaptive education. *arXiv preprint arXiv:2106.04262*, 2021.
- [22] Alexander Rich, Pamela Osborn Popp, David Halpern, Anselm Rothe, and Todd Gureckis. Modeling second-language learning from a psychological perspective. In *Proceedings of the thirteenth workshop on innovative use of NLP for building educational applications*, pages 223–230, 2018.
- [23] Sherry Ruan, Wei Wei, and James Landay. Variational deep knowledge tracing for language learning. In *LAK21: 11th International Learning Analytics and Knowledge Conference*, pages 323–332, 2021.
- [24] Anton Osika, Susanna Nilsson, Andrii Sydoruk, Faruk Sahin, and Anders Huss. Second language acquisition modeling: An ensemble approach. *arXiv preprint arXiv:1806.04525*, 2018.
- [25] Shuyao Xu, Jin Chen, and Long Qin. Cluf: A neural model for second language acquisition modeling. In *Proceedings of the thirteenth workshop on innovative use of NLP for building educational applications*, pages 374–380, 2018.
- [26] Florian Sense, Ryan Wood, Michael G Collins, Joshua Fiechter, Aihua Wood, Michael Krusmark, Tiffany Jastrzembski, and Christopher W Myers. Cognition-enhanced machine learning for better predictions with limited data. *Topics in cognitive science*, 14(4):739–755, 2022.
- [27] Álvaro J Jiménez Palenzuela, Flavius Frasinca, and Maria Mihaela Truşcă. Modeling second language acquisition with pre-trained neural language models. *Expert Systems with Applications*, 207:117871, 2022.
- [28] Masahiro Kaneko, Tomoyuki Kajiwara, and Mamoru Komachi. Tmu system for slam-2018. In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 365–369, 2018.
- [29] Yves Bestgen. Predicting second language learner successes and mistakes by means of conjunctive features. In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 349–355, 2018.
- [30] Zheng Yuan. Neural sequence modelling for learner error prediction. In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 381–388, 2018.
- [31] Brendan Tomoschuk and Jarrett Lovelett. A memory-sensitive classification model of errors in early second language learning. In *Proceedings of the thirteenth workshop on innovative use of NLP for building educational applications*, pages 231–239, 2018.
- [32] Guanliang Chen, Claudia Hauff, and Geert-Jan Houben. Feature engineering for second language acquisition modeling. In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 356–364, 2018.
- [33] Nihal V Nayak and Arjun R Rao. Context based approach for second language acquisition. In *Proceedings of the thirteenth workshop on innovative use of NLP for building educational applications*, pages 212–216, 2018.
- [34] Sigrid Klerke, Héctor Martínez Alonso, and Barbara Plank. Grotoco@ slam: Second language acquisition modeling with simple features, learners and task-wise models. In *Proceedings of the thirteenth workshop on innovative use of NLP for building educational applications*, pages 206–211, 2018.
- [35] Jill-Jënn Vie. Deep factorization machines for knowledge tracing. *arXiv preprint arXiv:1805.00356*, 2018.

A Appendix / supplemental material

All the graphs related to fairness on Client are listed below

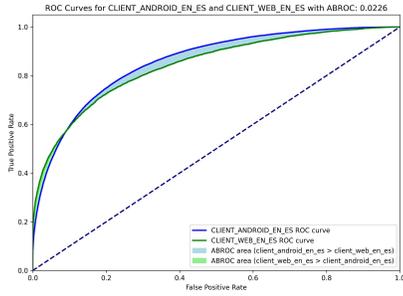


Figure 6: Fairness between android and web in en_es track for GBDT

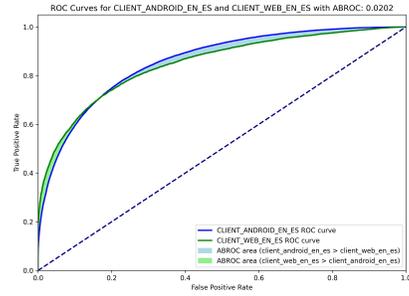


Figure 7: Fairness between android and web in en_es track for Multi-task learning

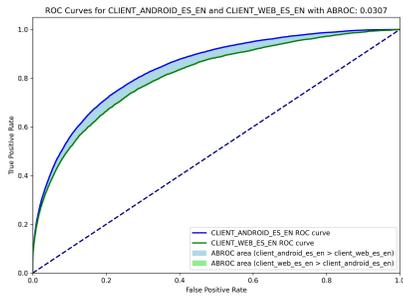


Figure 8: Fairness between android and web in es_en track for GBDT

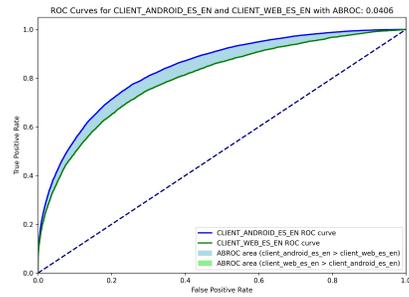


Figure 9: Fairness between android and web in es_en track for Multi-task learning

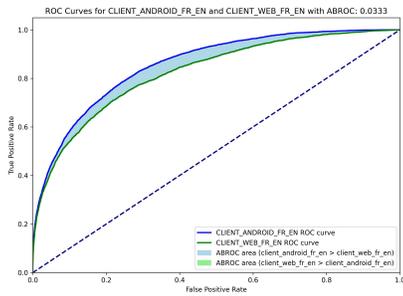


Figure 10: Fairness between android and web in fr_en track for GBDT

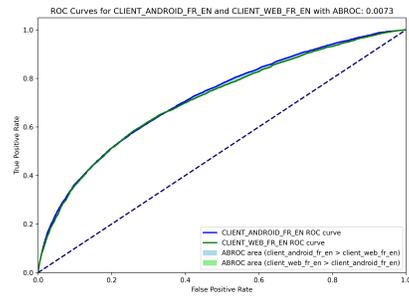


Figure 11: Fairness between android and web in fr_en track for Multi-task learning

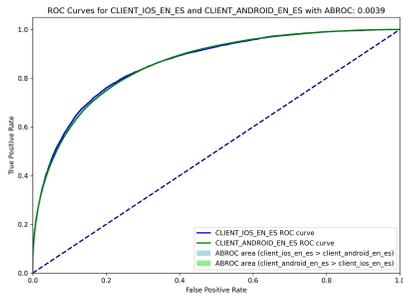


Figure 12: Fairness between ios and android in en_es track for GBDT

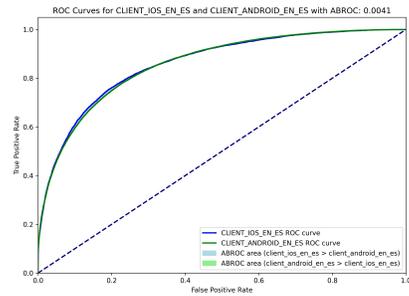


Figure 13: Fairness between ios and android in en_es track for Multi-task learning

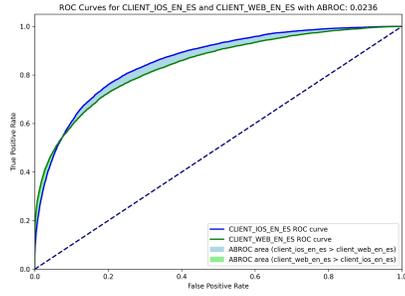


Figure 14: Fairness between ios and web in en_es track for GBDT

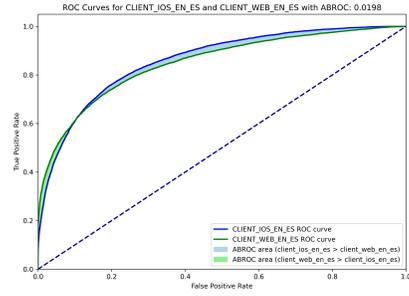


Figure 15: Fairness between ios and web in en_es track for Multi-task learning

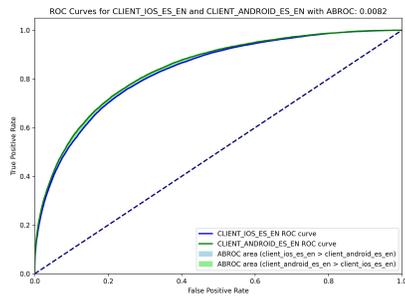


Figure 16: Fairness between ios and android in es_en track for GBDT

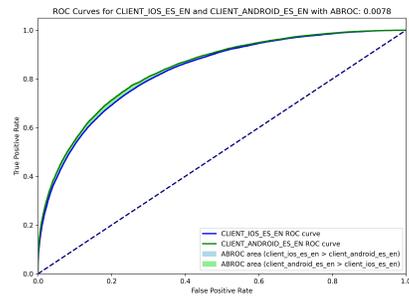


Figure 17: Fairness between ios and android in es_en track for Multi-task learning

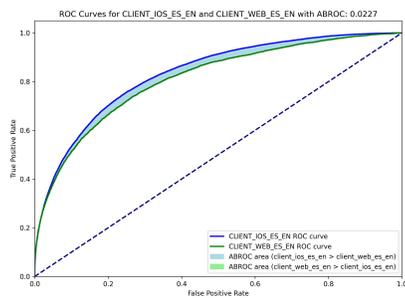


Figure 18: Fairness between ios and web in es_en track for GBDT

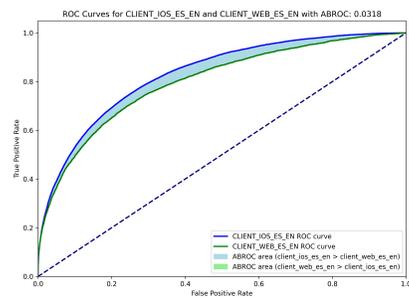


Figure 19: Fairness between ios and web in es_en track for Multi-task learning

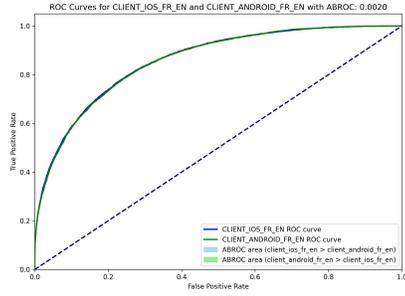


Figure 20: Fairness between ios and android in fr_en track for GBDT

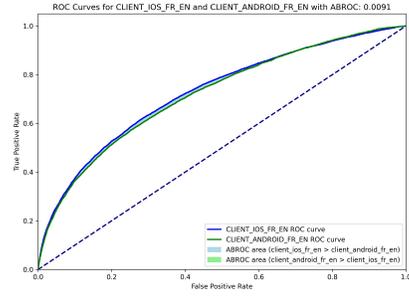


Figure 21: Fairness between ios and android in fr_en track for Multi-task learning

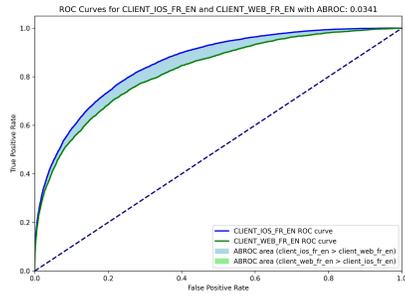


Figure 22: Fairness between ios and web in fr_en track for GBDT

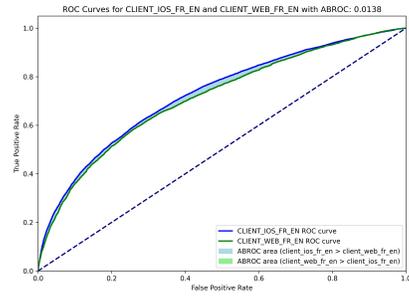


Figure 23: Fairness between ios and web in fr_en track for Multi-task learning

All the graphs related to fairness on Country are listed below

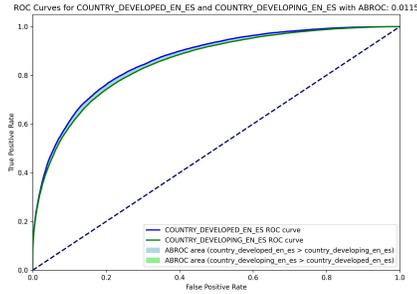


Figure 24: Fairness between developed and developing in en_es track for GBDT

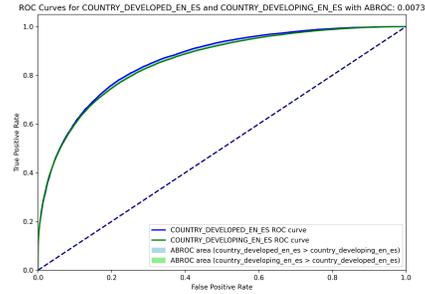


Figure 25: Fairness between developed and developing in en_es track for Multi-task learning

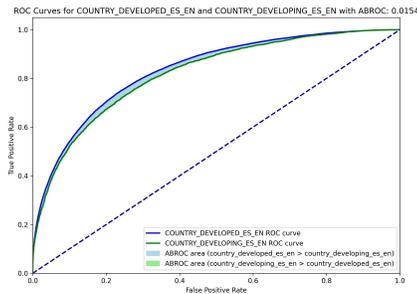


Figure 26: Fairness between developed and developing in es_en track for GBDT

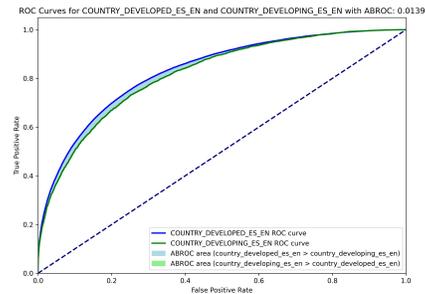


Figure 27: Fairness between developed and developing in es_en track for Multi-task learning

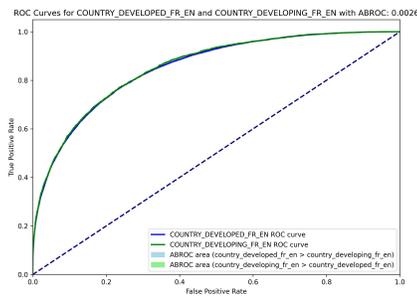


Figure 28: Fairness between developed and developing in fr_en track for GBDT

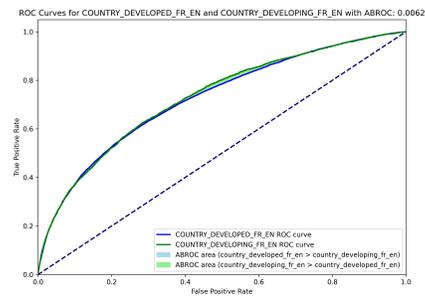


Figure 29: Fairness between developed and developing in fr_en track for Multi-task learning