Bias in MOOC Completion Prediction: A Fairness Analysis Across African Countries and Socio-Economic Contexts

Jade Mai Cock
LIP6, CNRS, Sorbonne University
Paris, France
Jade-Mai.Cock@lip6.fr

Sébastien Lallé
LIP6, CNRS, Sorbonne University
Paris, France
sebastien.lalle@lip6.fr

Mélina Verger
INSA Lyon, LIRIS
Lyon, France
melina.verger@insa-lyon.fr

François Bouchet
LIP6, CNRS, Sorbonne University
Paris, France
francois.bouchet@lip6.fr

ABSTRACT

Low completion rates are a well-documented challenge in MOOCs, and Machine Learning (ML)-based predictors have been proposed as a way to detect at-risk learners and support them through personalised interventions. Yet, evidence shows that ML models trained on such data can be biased against certain groups of learners, often the historically discriminated ones. In particular, usage patterns can substantially vary across world regions due to contextual factors, which can affect predictions. In this work, we analyse ML models trained on a MOOC dataset including African and OECD learners to examine whether, and for whom, biases arise. Our findings show that learners from Middle and Western Africa and from countries with low HDI and low literacy rates are affected by significant biases in ML predictions. Our findings highlight the need to systematically evaluate ML unfairness in African educational contexts.

Keywords

Fairness, Biases, Machine learning, MOOCs, Africa

1. INTRODUCTION

The widespread adoption of ML in education, e.g., for predicting course success or dropout risk, has raised concerns about their fairness towards historically discriminated groups [7]. ML unfairness arises when model behaviours and accuracy substantially differ across learner groups [8, 21]. When such models guide important decisions like program admission, resource allocation, and interventions for at-risk learners, they risk perpetuating existing biases and hindering the academic outcomes for specific groups of learners [4].

Recent educational data mining research has explored fairness in ML models on educational datasets [3, 10, 11, 26, 29]. Yet, surveys and opinion papers highlight a key limitation: an overwhelming amount of studies focuses on North American and European learners, with few considering African contexts [8, 17, 21]. Yet, African educational contexts have distinctive features, including the high prevalence of mobile or offline learning [16], stark rural-urban contrasts [14], unequal access to connectivity, electricity and computers [16], widespread reliance on open digital spaces and digital universities to mitigate infrastructure limitations [24], decreasing gender parity with higher levels of educational level increases [1], and overcrowded classrooms [12]. These may affect the target population and platform interactions, potentially reflecting contextual constraints rather than learners' engagement, abilities or interests.

Understanding this gap is critical for fairness-aware analysis, particularly in MOOCs which serve international and highly diverse learners, and are widely adopted in Africa [12, 16, 28]. While ML studies to predict drop-out on MOOC data is extensive [2, 18], fairness analyses are scarce [10, 13, 22]. As a result, the promise of democratizing access to knowledge across continents through MOOCs cannot be fully realized if known contextual differences are not accounted for from the course design stage onward. In this respect, indiscriminately combining data from African and Western learners may exacerbate existing biases rather than mitigate them.

Within prior work, the only openly available MOOC dataset used in a fairness study is the GDP-MOOC dataset published by Lallé et al. [13]. This dataset involves students from 23 African countries, mostly located in Western, Northern and Middle Africa. This work [13] showed that ML models can accurately predict MOOC completion, but were biased against African learners. However, they treated all African learners as a single homogeneous group. Building on this, we propose a fine-grained fairness analysis of biases across African countries and socio-economic characteristics in this dataset. Beyond MOOCs, ML biases have been analysed for automatic scoring of Nigerian essays (which accounted for only $< 0.03\%$ of the data) [6], and in the PISA 2015 data, but without specific results for Africa (Al-

Jade Mai Cock, Sébastien Lallé, Mélina Verger, and François Bouchet. Bias in MOOC Completion Prediction: A Fairness Analysis Across African Countries and Socio-Economic Contexts. In Anthony Botelho, Maria Mercedes T. Rodrigo, Adish Singla, Hiroaki Ogata, Hyojeong So, and Young Hoan Cho (eds.) Proceedings of the 19th International Conference on Educational Data Mining, Seoul, Republic of Korea, June, 2026, pp. 605–610. International Educational Data Mining Society (2026).

https://doi.org/10.5281/zenodo.21040101

geria and Tunisia) [15]. This led to this research question: *Which specific countries or socio-economic characteristics help explain the biases toward Africa of ML models predicting course success in the GDP-MOOC dataset?*

2. METHOD

2.1 Data and Machine Learning Setup

We used the GDP-MOOC open dataset in this paper (`https://github.com/Mocahteam/GdPMOOC_Fairness`) [13]. It originates from a French MOOC on Project Management lasting five weeks (2022 offering) delivered on the OpenEdX platform. The dataset includes weekly aggregated statistical summaries (e.g. counts) derived from OpenEdX log data, the students' country, and their final grade. Course completion (pass/fail) is defined as achieving a final grade of at least 50%. The GDP-MOOC dataset corresponds to the Spring 2022 offering. It includes complete data for 912 learners (460 are in Africa, 452 in OECD countries in Western Europe and North America), among whom 380 (39%) successfully passed the course. Learners provided informed consent for their data to be used for research purposes, and all data are fully anonymized per the EU's General Data Protection Regulation (GDPR).

We replicate the ML pipeline proposed in [13] and shared in their GitHub to train three supervised ML models and predict in a binary way whether learners would *succeed* or *fail* to obtain the certificate at the end of the MOOC. The features used to train the models are the concatenation of weekly action counts f_w^l for each learner $l \in \mathcal{L}$. We train Logistic Regression (LR), Naive Bayes (NB), and Stochastic Gradient Boosting (SGB) on these features, selected respectively for classification performance, fairness, and accuracy–fairness trade-off [13]. We apply learner-level 5-fold stratified cross-validation with feature selection. The full dataset description and ML pipeline are reported in [13], and here we focus on the fairness analysis.

2.2 Bias Analyses

As stated above, ML models trained on the GDP-MOOC dataset exhibited significant unfairness toward African learners [13]. To better characterize these biases, we conduct a subgroup analysis based on the country of African learners (N=460). Their distribution (shown in Table 1) is highly uneven, with the top six countries accounting for over two-thirds of learners and all others having fewer than 20 each. Consequently, we opt to address our research question in two complementary ways: first by focusing on the six countries with over 20 learners (Ivory Coast, DR Congo, Benin, Cameroon, Burkina Faso and Senegal); and second by studying larger groups of countries based on UN-based socio-economic indicators. This second approach highlights higher-level biases that span similar countries. In addition, it allows us to include all of African learners into meaningful and relevant larger analytical groups. The country characteristics (listed in Table 2) are based on five country-level UN-based indicators.

UN Regions: The five sub-regions of Africa defined by the UN Geoscheme, a statistical way to group countries by main geographical areas. This is to understand at the high level whether specific parts of the continent receive more unfair predictions. The data are from the UN Statistical Division[1].

HDI: The four Human Development Index groups as defined by the UN Development Program, a widely accepted way to quantify human development in terms of life expectancy, quality of education, and standard of living. This allows assessing whether ML models perform equitably across contexts with unequal structural opportunities. The HDI groups are based on the Human Development Report for 2022[2].

Literacy: The levels of literacy among adults, as compiled by the UNESCO, which is used as a proxy for the access to fundamental education. The 2022 data are extracted from the IUS Bulk Data[3].

Education: The quality of education (from elementary to university), which is compiled by the UNESCO across a large battery of indicators as part of its Sustainable Development Goals (SDG-4). This is to evaluate whether lower quality of education in a country hinders the overall usage of the MOOCs, and in turn, the fairness of the models. The data are extracted from the IUS Bulk Data[3].

GDI: We consider the Gender Development Index as measured by the UN Development Program, since lower levels of gender development may influence how men and women engage with online learning resources such as MOOCs, potentially affecting ML fairness. Using the raw GDI values, we construct two categories, "low" and "high" GDI, by applying the African median as the cutoff threshold. The data are extracted from the Human Development Report for 2022[2].

We also considered the countries' GDP and access to the Internet, but discarded them as both metrics were highly correlated to the HDI (Pearson $r > .7$).

Fairness metrics. We measure the potential biases across the different groups described in Table 2 through two complementary metrics, namely the group-level AUC, and the MADD. For the first one, we compute the AUC of the models in the test sets for each of the groups, by selecting only learners from a specific group in the test sets (e.g., "West Africa") and computing the AUC of the models' outputs for those learners. This is a simple and commonly used approach to gauge how a model's accuracy varies across groups [10, 21, 23]. For a model to be fair, its AUC should be similar across the different groups of learners. As for the MADD, it leverages the histograms of probabilities output by a model for two groups to compute the areas where the histograms do not overlap [27]. This approach allows to measure how a given characteristics influence the models' outputs, regardless of its performances, as well as toward whom the output probabilities are inflated. To be fair, a models' outputs should not be related to any countries or characteristics it does not know about. Because the MADD involves comparing two groups in a pairwise manner, we use it by comparing each target group against all of the other learners (e.g., "Western Africa" learners vs. everybody else). AUC ranges from 0 (worst performance) to 1 (best performance) and MADD ranges from 0 (fairest) to 2 (most unfair). Both AUC and MADD are independent from

[1] `https://unstats.un.org/unsd/methodology/m49/`
[2] `https://hdr.undp.org/system/files/documents/global-report-document/hdr2023-24reporten.pdf`
[3] `https://databrowser.uis.unesco.org/resources/bulk`

Table 1: **Number of learners (n) and passing rates (pr) per African country in the GDP-MOOC dataset.**

Country	n (pr)	Country	n (pr)	Country	n (pr)	Country	n (pr)
Ivory Coast	82 (22%)	Togo	16 (31%)	Niger	10 (20%)	Burundi	3 (0%)
DR Congo	74 (36%)	Guinea	14 (14%)	Tchad	10 (3%)	Mauritania	2 (50%)
Benin	60 (40%)	Congo Republic	13 (23%)	Tunisia	9 (22%)	Central African Rep.	2 (0%)
Cameroon	42 (36%)	Mali	13 (8%)	Algeria	8 (38%)	Rwanda	1 (0%)
Burkina Faso	40 (28%)	Gabon	11 (27%)	Madagascar	4 (0%)	Mozambique	1 (0%)
Senegal	33 (27%)	Morocco	8 (38%)	Djibouti	4 (0%)		

Table 2: Country characteristics, sample size per group within Africa and passing rates. We report all groups for completeness, however, only groups with ≥ 20 learners are included in the analysis.

Characteristics	Groups (#learners [passing rate])
Sub-regions of Africa	Northern (25 [32%]), Western (270 [27%]), Eastern (13 [0%]) , Middle (152 [34%]), Southern (0)
HDI	Low (369 [28%]), Medium (63 [33%]), High (28 [29%]), Very high (0)
Literacy	Low (231 [26%]), Medium (155 [33%]), High (41 [27%]), Very high (33 [27%])
Quality of education	Low (79 [22%]), Lower middle (292 [30%]), Upper middle (65 [31%]), High (24 [33%])
Gender Gap Index	Low (281 [33%]), High (179 [22%])

the decision threshold of the ML models, i.e., the probability cutoff above which a binary classifier labels an instance as positive. Thus, they can provide a more robust and stable assessment than metrics tied to a single threshold. In addition, AUC and MADD are complementary, as AUC deals with the likelihood of correct predictions across groups and MADD with the independence of the models' decisions toward group membership. Both metrics have been used in EDM fairness studies, e.g., [19, 23, 27, 29].

3. RESULTS

Table 3 reports on the evaluation of the ML models' fairness over the countries and characteristics. For the MADD, we indicate with a +/- sign whether the model's behaviour is in favour/biased (i.e., tends to give higher/lower probabilities of success in the MOOC) toward the target group. To formally compare these results, we apply a separate Friedman test for each unique combination of metric (group-level AUC and MADD), ML model (LR, SGB, and NB) and factor (countries or characteristics). Each test assesses differences in the metric across the levels of a single factor within one ML model, with repeated measures defined by the AUC/MADD values computed at each week and fold of cross-validation. To account for multiple comparisons, we apply the Benjamini–Hochberg method (BH) to adjust all p-values simultaneously across the full set of Friedman tests, making the correction stricter. We then examine significant main effects using pairwise Wilcoxon signed-rank tests with adjustment, and indicate in bold in Table 3 the values for which the tests show significant unfairness.

Sub-regions. Table 3 shows contrasting results. LR favors North Africa and is more biased against Middle Africa, while NB favors Western Africa and is more biased against Northern Africa. The Friedman tests show a significant main effect of sub-regions on AUC for RF ($p = 0.008, \eta^2 = 0.30$), and on MADD for LR, SGB and NB ($p = 0.0001, \eta^2 = 0.84$, $p = 0.0001, \eta^2 = 0.75$, $p = 0.001, \eta^2 = 0.67$, respectively). The post-hoc pairwise tests indicate that LR is significantly more biased toward Middle Africa on both metrics ($p < 0.03$). For SGB and NB, there is no significant difference in AUC, but the main effects and pairwise tests for MADD also show significantly more biases toward Middle Africa. These results provide evidence that the Middle Africa sub-region countries would experience some levels of unfairness with all models, especially LR, even though its passing rate is the highest (34%) among the three regions.

HDI. Table 3 shows that the AUC for High HDI countries is substantially higher than for the other groups with LR and SGB, but not NB. The Friedman tests only shows a significant main effect of HDI on AUC for LR ($p = 0.049, \eta^2 = 0.19$) and NB ($p = 0.038, \eta^2 = 0.20$), confirming the trends that LR is significantly less accurate for Low HDI and NB for high HDI. For all models, there is also a significant main effect in MADD ($p < 0.0001$, $\eta^2 > 0.66$ in all cases), consistently showing more biases for Low and Medium HDI groups, according to the post-hoc tests. These results show that LR exhibits clear biases for low HDI countries.

Literacy. The Friedman tests show no significant main effect of literacy groups on AUC, but significant main effects for all models on MADD ($p < 0.0001, \eta^2 > 0.51$). For all of these effects, the pairwise comparisons always show significant biases toward Low and Medium levels of Literacy, as compared to High (and surprisingly Very high for LR). These results show that literacy can capture ML unfairness in this dataset, but only in terms of MADD.

Quality of education. The Friedman tests show no main effects on AUC, but on MADD for LR and SGB ($p < 0.0001$, $\eta^2 > 0.52$). This suggests that the quality of education impacts to some extent the models' behaviors, but not their performance. The post-hoc tests for MADD indicate more biases toward the Low education groups ($p < 0.05$) across the models. More surprisingly, the same effect is found for the High education group (as compared to Lower middle and Upper middle), possibly due to its small size.

GDI. The Friedman tests reveal no main effect of GDI on

Table 3: Results per group and ML model. Bold means significantly more biased toward that group.

Factor	Values	LR		SGB		NB	
		AUC	MADD	AUC	MADD	AUC	MADD
Sub-regions	Northern Africa	0.82	1.03+	0.73	0.95+	0.69	0.28+
	Western Africa	0.75	**0.81-**	0.73	**0.59-**	0.75	0.64+
	Middle Africa	**0.73**	**0.47-**	0.74	**0.40-**	0.70	**0.22-**
HDI	Low	**0.74**	**0.75-**	0.73	**0.54-**	0.74	**0.66-**
	Medium	0.75	**0.73-**	0.75	**0.65-**	0.73	**0.22-**
	High	0.87	0.94+	0.78	0.86+	**0.63**	0.19+
Literacy	Low	0.76	**0.78-**	0.73	**0.58-**	0.74	**0.63-**
	Medium	0.73	**0.49-**	0.72	**0.44-**	0.73	**0.29-**
	High	0.84	0.82+	0.78	0.74+	0.69	**0.18-**
	Very high	0.71	**0.90-**	0.74	0.68-	0.77	0.45+
Quality of education	Low	0.72	**0.79-**	0.73	**0.62-**	0.72	0.52-
	Lower middle	0.76	0.62+	0.74	0.48+	0.74	0.52+
	Upper middle	0.78	0.78+	0.77	0.66+	0.68	0.27+
	High	0.71	**1.01-**	0.66	**0.96-**	0.83	0.40+
GDI	Low	0.75	0.53-	0.74	0.41-	0.73	0.42-
	High	0.75	0.85-	0.72	0.68-	0.75	0.57-
Countries	Benin	0.78	**0.75-**	0.76	0.61-	0.72	0.37-
	Burkina Faso	**0.71**	**0.84-**	0.71	**0.79-**	0.75	0.35-
	Cameroon	0.76	**0.89-**	0.76	**0.75-**	**0.69**	**0.23-**
	DR Congo	0.72	**0.57-**	0.71	0.56-	0.72	0.16-
	Ivory Coast	0.75	**0.66-**	**0.69**	**0.77-**	0.76	0.66-
	Senegal	**0.71**	**0.90-**	0.74	**0.88-**	0.77	0.45+
	All Africa	**0.75**	0.86-	**0.73**	0.60-	0.73	0.70-
	All OECD	0.81	0.86+	0.77	0.60+	0.70	0.70+

AUC and MADD, suggesting that no unfair treatments was observed in the models' outputs for this factor. This could stem from the fact that gender gaps are on the high side on the UN gender gap scales, meaning that this factor does not capture in sufficient details the nuances in gender gap across African countries, or conversely that all model can be trained fairly across GDI groups.

Countries. Table 3 shows that none of the African countries reach an AUC above 0.8, and that performances vary substantially across models. Friedman tests reveal a significant main effect of country on AUC for LR ($p = 0.006, \eta^2 = 0.25$), SGB ($p = 0.008, \eta^2 = 0.24$) and NB ($p = 0.002, \eta^2 = 0.29$). Post-hoc tests further indicate that LR is significantly less accurate for all African countries as compared to OECD ones, SGB for learners from Ivory Coast, and NB for learners from Cameroon. This shows that no model can perform reliably across African countries. There are also main effects of country on MADD for LR, SGB and NB. Post-hoc tests show that LR is more biased against all African countries as compared to OECD ones, SGB against all countries but Benin and DR Congo, and NB against Cameroon. Lastly, similarly to [13] when comparing African versus OECD countries (last two rows of Table 3), significant biases are observed for LR ($p = 0.002, \eta^2 = 0.19$) and SGB ($p = 0.039, \eta^2 = 0.11$), confirming that these models are more biased against Africa, even though African learners constitute the majority of the learners in this dataset.

4. DISCUSSION

Our results do show several patterns of unfairness in our ML predictors, thus answering our research question affirmatively. Three main trends of unfairness emerge in our results. First, we found evidence of biases toward learners from Middle Africa, a region that, to our knowledge, has received no prior attention in the ML fairness literature. We thus contribute to this work by providing strong evidence of biases in ML predictors trained on an international educational dataset for this part of Africa. Within Middle Africa, learners from Cameroon are exposed to biases in all models, especially NB. This calls for investigating further the possible reasons that could shape these learners' data in future work. In particular, Cameroon is characterized by strong multilingualism, severe infrastructural deficiencies in rural schools, and high dropout rates in primary and secondary education among girls and students from lower-income households [20]. Second, we uncovered biases against Western Africa, at least with LR and SGB. In related work on AI fairness in education, solely one study [6] looked at Nigerian learners (in Western Africa), without finding evidence of unfairness, possibly due to their small number of Nigerian data points. Within this region of Africa, we found in particular that learners from Burkina Faso, Senegal and Ivory Coast are exposed to varied levels of biases in the ML models, further reinforcing the per-African country analysis we propose in this paper. In a recent institutional report on deploying a MOOC at scale in Western Africa, educators from these countries reported specific issues that could explain different usage patterns of MOOCs resources [25].

Specifically, educators reported Internet connectivity issues leading to student dropouts in Burkina Faso; a clear lack of experience with remote learning in Ivory Coast; and the need for offline access to the MOOC in Senegal to enable participation for all learners. For Ivory Coast, we also found very low levels of interaction across the board, further showing that students from this country may have struggled with the MOOC's activities, which could have in turn make it more difficult for ML models to learn from these data. Third, we also found evidence of unfairness that appears to be matching existing struggles, namely countries with low HDI and low levels of adult literacy. To our knowledge, no prior studies have uncovered these trends, which, worth noticing, are not explained by sample size, since many biases affect large groups (e.g., Low HDI, Middle Africa). Overall, our findings highlight the need to broaden fairness research in ML to more diverse contexts, including replicating existing fairness studies in African settings.

Furthermore, as for the ML models, we found that LR generates substantial unfairness towards several groups, and in general is more favourable to OECD learners, which is consistent with the findings in [13] that LR is the most accurate but the most biased model. While this previous work also found that NB is the fairest model, we actually found that NB can be biased as well (e.g., toward Cameroon), showing that examining in depth unfairness is crucial to actually deploy an ML model. This confirms that different models trained on the same data can generate their own algorithmic biases.

Overall, the absence of correlation between the observed unfairness and subgroup pass rates, proportional representation, or absolute numbers of successful learners suggests that the sources of bias in our setting differ from typical representational and measurement biases, typically identified in Western educational contexts [4]. This observation complements the regional patterns discussed above and supports the hypothesis that commonly used modelling features, largely designed around assumptions and behavioural norms characteristic of Western learning environments, may not transfer well to African contexts. In particular, prior work has shown that while infrastructural constraints constitute a major barrier to MOOC adoption (often exceeding challenges related to self-discipline [9]), course contextualization, including the extent to which content reflects learners' lived realities, also plays a critical role in shaping engagement [5]. Alongside these factors, variations in digital literacy, constrained access to reliable infrastructure, and educational trajectories may shape learner behaviours in ways that are not adequately captured by current feature representations. As a result, existing modelling approaches may fail to encode meaningful indicators of learning and engagement, thereby contributing to systematic disparities in predictions. These findings highlight the need for context-sensitive feature design and fairness interventions explicitly accounting for linguistic, cultural, and infrastructural differences.

In addition to potential biases introduced by the ML models, the diversity of learning contexts and modalities for accessing the MOOC in African countries likely contributed to greater feature heterogeneity, which in turn affected model training. In the GDP-MOOC dataset, for instance, African learners interacted significantly less with videos (fewer videos played, and fewer actions such as pausing, seeking, or enabling transcripts) than OECD learners, according to a t-test with amount of video interaction as the dependent measure ($p < .05$). Within African subgroups, learners from the DR Congo showed the lowest levels of video interaction. These findings are consistent with prior work (see Introduction) emphasizing diverse modalities and challenges for accessing remote courses in Africa (e.g., mobile learning, access through digital open spaces), which may help explain some of the patterns observed here.

5. CONCLUSION

In this work, we explored the biases in ML predictors of MOOC completion, focusing on African learners, by leveraging the publicly available GDP-MOOC dataset, which includes logged events, outcomes, and learners' countries, with about half of them being African. To the best of our knowledge, our study is the first that focuses on eliciting complex patterns of unfairness within Africa in educational data.

Overall, our findings show diverse forms of unfairness that vary by ML model, highlighting the importance of investigating ML unfairness for subgroups of African learners. As these findings show no association with subgroup pass rates, proportional representation, or absolute numbers of successful learners, they suggest that the observed biases cannot be attributed solely to representation effects.

These findings provide strong evidence for the need to control for the fairness of ML predictors in education before deploying them in Africa, and provide insights about the specific groups needing more attention. However, due to the GDP-MOOC being in French, most of the African students came from countries with a sizable French-speaking population, typically Western, Middle, and Northern Africa, and future work should consider other parts of Africa and multilingual MOOCs. Another limitation is that we focused on one offering of a single MOOC, hence future work should aim at replicating the results over a larger and more balanced pool of MOOC learners. Lastly, while we provide possible reasons for the observed biases in the ML models, including structural and economical challenges documented in African studies about MOOCs, future work is still needed to link algorithmic unfairness with actual underlying causes.

With the rise of ML-driven tools in education, it is crucial to carefully consider their robustness for African learners, a continent with millions of learners who have received little attention so far in existing ML fairness studies in education. If we are to fulfill the promise that knowledge can reach a wider audience, we must pay close attention to cultural diversity at every stage of the educational and ML pipeline.

Acknowledgements. We thank Rémi Bachelet and the MOOC GdP's team for the data collection.

6. REFERENCES

[1] E. R. Aikins and J. Cilliers. Education. https://futures.issafrica.org/thematic/06-education/, 2026. Published online at futures.issafrica.org. Updated October 11, 2024. Accessed February 9, 2026.

[2] S. Alghamdi, B. Soh, and A. Li. A comprehensive review of dropout prediction methods based on multivariate analysed features of mooc platforms. *Multimodal Technologies and Interaction*, 9(1):3, 2025.

[3] H. Anderson, A. Boodhwani, and R. S. Baker. Assessing the fairness of graduation predictions. In *Proceedings of The 12th International Conference on Educational Data Mining*, 2019.

[4] R. S. Baker and A. Hawn. Algorithmic bias in education. *International Journal of Artificial Intelligence in Education*, pages 1052–1092, 2021.

[5] S. Boga and R. McGreal. Introducing moocs to africa: New economy skills for africa program. 2014.

[6] B. Bridgeman, C. Trapani, and Y. Attali. Considering fairness and validity in evaluating automated scoring. In *Proceedings of the Annual meeting of the National Council on Measurement in Education*, 2009.

[7] J. Buolamwini and T. Gebru. Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency*, pages 77–91, 2018.

[8] S. V. Chinta, Z. Wang, Z. Yin, N. Hoang, M. Gonzalez, T. L. Quy, and W. Zhang. Fairaied: Navigating fairness, bias, and ethics in educational ai applications. *arXiv preprint arXiv:2407.18745*, 2024.

[9] C. E. Franco Yáñez, D. Nigmonova, and W. Panichpathom. Demoocratization of education?: Massive open online courses (moocs) and the opportunities and challenges for developing countries. 2014.

[10] J. Gardner, C. Brooks, and R. 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.

[11] Q. Hu and H. Rangwala. Towards fair educational data mining: A case study on detecting at-risk students. In *Proceedings of The 13th International Conference on Educational Data Mining (EDM)*, 2020.

[12] A. I. Jouicha, K. Berrada, R. Bendaoud, S. Machwate, A. Miraoui, and D. Burgos. Starting moocs in african university: The experience of cadi ayyad university, process, review, recommendations, and prospects. *IEEE Access*, 8:17477–17488, 2020.

[13] S. Lallé, F. Bouchet, M. Verger, and V. Luengo. Fairness of mooc completion predictions across demographics and contextual variables. In *International Conference on Artificial Intelligence in Education*, pages 379–393. Springer, 2024.

[14] R. Lembani, A. Gunter, M. Breines, and M. T. B. Dalu. The same course, different access: the digital divide between urban and rural distance education students in south africa. *Journal of Geography in Higher Education*, 44(1):70–84, 2020.

[15] X. Li, D. Song, M. Han, Y. Zhang, and R. F. Kizilcec. On the limits of algorithmic prediction across the globe. *arXiv preprint arXiv:2103.15212*, 2021.

[16] V. Maphosa and M. Maphosa. Opportunities and challenges of adopting moocs in africa: A systematic literature review. *Massive Open Online Courses-Current Practice and Future Trends*, 2023.

[17] B. Memarian and T. Doleck. Fairness, accountability, transparency, and ethics (fate) in artificial intelligence (ai) and higher education: A systematic review. *Computers and Education: Artificial Intelligence*, 5:100152, 2023.

[18] P. M. Moreno-Marcos, C. Alario-Hoyos, P. J. Muñoz-Merino, and C. D. Kloos. Prediction in moocs: A review and future research directions. *IEEE transactions on Learning Technologies*, 12(3):384–401, 2018.

[19] N. Nasiar, C. Belitz, H. Lee, F. Stinar, R. Baker, J. Ocumpaugh, S. Fancsali, S. Ritter, and N. Bosch. Investigating algorithmic bias in affect detectors with constructed categories of student identity. In *Proceedings of the International Conference on Computers in Education*, 2025.

[20] A. Ndindeng. Factors influencing under-education in cameroon: A comprehensive analysis. *African Educational Research Journal*, 12(3):182–201, 2024.

[21] N. Pham, H. P. Ngoc, and A. Nguyen-Duc. Fairness for machine learning software in education: A systematic mapping study. *Journal of Systems and Software*, 219:112244, 2025.

[22] L. Sha, M. Raković, A. Das, D. Gašević, and G. Chen. Leveraging class balancing techniques to alleviate algorithmic bias for predictive tasks in education. *IEEE Transactions on Learning Technologies*, 15(4):481–492, 2022.

[23] V. Švábenskỳ, M. Verger, M. Rodrigo, C. Monterozo, R. Baker, M. L. Saavedra, S. Lallé, and A. Shimada. Evaluating algorithmic bias in models for predicting academic performance of filipino students. In *Proceedings of the 17th Educational Data Mining Conference*, 2024.

[24] K. Sylla, G. M. M. Nkwetchoua, and F. Bouchet. How does the use of open digital spaces impact students success and dropout in a virtual university?. *International Association for Development of the Information Society*, 2022.

[25] C. Tkaczuk, L. Jeannin, and A. Jaillet. Rapport du mooc certicescol – session 2 – année 2021 – cy cergy paris université, inspé versailles, cy bonheurs, 2025.

[26] M. Verger, C. Fan, S. Lallé, F. Bouchet, and V. Luengo. A comprehensive study on evaluating and mitigating algorithmic unfairness with the madd metric. *Journal of Educational Data Mining*, 16(1):365–409, 2024.

[27] M. Verger, S. Lallé, F. Bouchet, and V. Luengo. Is your model" madd"? a novel metric to evaluate algorithmic fairness for predictive student models. In *Proceedings of the International Conference on Educational Data Mining*, pages 91–102, 2023.

[28] A. A. Yunusa, I. N. Umar, and B. Bervell. Massive open online courses (moocs) in sub-saharan african higher education landscape: a bibliometric review. *MOOC*, pages 1–25, 2021.

[29] A. F. Zambrano, J. Zhang, and R. S. Baker. Investigating algorithmic bias on bayesian knowledge tracing and carelessness detectors. In *Proceedings of the 14th International Conference on Learning Analytics & Knowledge*, 2024.