

# ALGORITHMIC BIAS IN AUTOMATED SCHOLARSHIP SELECTION SYSTEMS

**Dr. Ananya Verma**

Assistant Professor

Department of Computer Science and Engineering

School of Artificial Intelligence and Data Science

University of Technology and Innovation, Bengaluru, Karnataka, India

## Abstract

*Automated scholarship selection systems are increasingly used to manage large application volumes, verify eligibility, prioritise candidates, and support funding decisions. These systems can improve administrative consistency and reduce delays, but they can also reproduce or intensify inequality when historical records, proxy variables, incomplete digital submissions, or opaque ranking rules influence outcomes. This paper provides a comparative, evidence-based analysis of algorithmic bias in automated scholarship selection. It examines bias risks across the complete decision pathway: application intake, document validation, eligibility screening, scoring, ranking, award allocation, communication of outcomes, and appeals. The analysis draws on scholarship in algorithmic fairness, education, public decision-making, and responsible artificial intelligence; it does not claim a new survey, experimental trial, or applicant-level dataset. The study finds that harmful bias can arise even when systems do not directly use protected characteristics. Variables such as school type, postal location, language, prior digital access, household documentation, attendance patterns, or predicted enrolment behaviour can act as indirect proxies for social advantage. A fair scholarship system therefore requires more than a technically accurate model. It requires clear policy objectives, representative and well-documented data, meaningful human oversight, group and individual fairness testing, accessible application support, reasoned decisions, correction mechanisms, and independent review. The paper concludes that scholarship allocation should remain a human-accountable public-interest process in which automation assists administrative judgement without replacing institutional responsibility for equity, transparency, and student dignity.*

**Keywords:** Algorithmic Bias, Scholarship Selection, Educational Equity, Automated Decision-Making, Artificial Intelligence, Fairness Auditing, Financial Aid, Human Oversight

## 1. INTRODUCTION

Scholarships are not merely financial transactions. They are instruments through which institutions, governments, foundations, and public agencies attempt to widen educational access, recognise merit, respond to economic hardship, support under-represented communities, and invest in social mobility. Because scholarship decisions can shape whether a learner remains in education, where that learner studies, and what opportunities become available afterward, the process must be credible to applicants and defensible to the public. Administrative pressure to process many applications quickly has encouraged institutions to adopt digital workflows that automate parts of eligibility checking, document sorting, ranking, recommendation, and award allocation.

Automation can be useful when it removes repetitive administrative work, identifies

incomplete records, applies stable eligibility rules, and helps staff compare applications against clearly stated criteria. However, the language of efficiency can conceal value choices. A model that predicts the likelihood of academic success, evaluates financial need, or recommends an award amount is not simply calculating a neutral fact. It converts institutional priorities into variables, thresholds, categories, and scoring rules. Those choices can advantage applicants whose circumstances fit the assumptions built into the system while disadvantaging students whose pathways are less conventional, less documented, or under-represented in historical records.

Algorithmic bias in scholarship selection can be introduced before a prediction or ranking model is trained. A digital application form may assume reliable internet access, English-language proficiency, a particular document

format, or the availability of formal proof of household income. Historical award data may reflect prior access to coaching, institutional familiarity, stable school resources, or biased committee preferences. When such data are treated as a benchmark for future allocation, a system can learn patterns of past advantage rather than the equity objective that a scholarship programme is meant to serve. As Barocas and Selbst (2016) explain in a wider context, data-driven systems can produce disparate effects even without an explicit intention to discriminate.

The education sector presents distinctive fairness challenges. Educational records are shaped by unequal schooling conditions, disability accommodations, regional infrastructure, language opportunities, family responsibilities, migration, and prior access to digital tools. Standardised indicators such as examination marks, attendance, school ranking, recommendation letters, or participation certificates may carry different meanings across applicants and communities. Kizilcec and Lee (2020) argue that educational algorithms must be assessed across measurement, model learning, and action because bias can emerge at each stage of the system rather than at one isolated technical point.

This study examines automated scholarship selection as a socio-technical process rather than as a narrow software problem. It identifies key pathways through which bias can arise, compares risks across the major stages of scholarship administration, and proposes safeguards for institutions that wish to use automated support responsibly. The central argument is that scholarship systems should not treat fairness as an optional feature added after deployment. Fairness must be incorporated into programme design, data governance, model development, committee practice, applicant communication, and appeal procedures from the beginning.

## 2. LITERATURE REVIEW

Research on algorithmic fairness has shown that bias is a broad and multi-layered problem. Mehrabi et al. (2021) identify numerous sources of bias in machine learning, including historical bias, representation bias,

measurement bias, aggregation bias, evaluation bias, and deployment bias. This perspective is particularly useful for scholarship selection because it prevents institutions from assuming that the removal of a sensitive column such as gender, caste, ethnicity, disability, or religion automatically makes a system fair. In practice, other variables may function as proxies, and the way a decision is implemented may create additional unequal effects.

A second strand of literature emphasises that fairness cannot be defined by a single mathematical metric. Binns (2018) demonstrates that competing fairness definitions reflect different political and moral ideas about equality, disadvantage, and deservingness. For scholarship programmes, this means that an institution must decide whether it is primarily seeking equal treatment under a common rule, proportional representation across groups, correction of historical disadvantage, minimum support for the most economically vulnerable applicants, or some combination of these goals. These choices must be made openly by the responsible institution rather than silently embedded in technical design.

The socio-technical literature also warns against excessive abstraction. Selbst et al. (2019) argue that a system cannot be understood fairly when designers separate an algorithm from the institutional setting, legal duties, users, and affected people around it. Scholarship decisions involve applicants, families, schools, financial-aid staff, committees, programme donors, and public regulators. The quality of an automated system therefore depends not only on accuracy statistics but also on how decisions are reviewed, how errors are corrected, whether applicants receive understandable reasons, and whether affected communities can challenge harmful patterns. Lifecycle frameworks developed by Suresh and Guttag (2021) and auditing frameworks proposed by Raji et al. (2020) provide practical guidance for identifying and managing these risks.

## LITERATURE REVIEW SUMMARY TABLE

| S. No. | Author / Year           | Study Focus                             | Key Contribution                                                          | Limitations Identified                                                     |
|--------|-------------------------|-----------------------------------------|---------------------------------------------------------------------------|----------------------------------------------------------------------------|
| 1      | Barocas & Selbst (2016) | Data mining and disparate impact        | Explained how neutral-looking data practices may produce unequal effects. | Focused on anti-discrimination law rather than scholarship administration. |
| 2      | Binns (2018)            | Fairness concepts in machine learning   | Showed that fairness metrics embed different moral and political choices. | Does not prescribe one fairness objective for education funding.           |
| 3      | Kizilcec & Lee (2020)   | Algorithmic fairness in education       | Mapped fairness concerns across measurement, model learning, and action.  | Broad education focus; scholarship workflows require added detail.         |
| 4      | Mehrabi et al. (2021)   | Bias and fairness in machine learning   | Systematised sources of bias and major mitigation approaches.             | Review is cross-sectoral and not scholarship-specific.                     |
| 5      | Suresh & Guttag (2021)  | Sources of harm across the ML lifecycle | Identified multiple stages where downstream harms can be introduced.      | Framework needs contextual interpretation for financial-aid decisions.     |
| 6      | Raji et al. (2020)      | Internal algorithmic auditing           | Outlined end-to-end accountability and audit practices.                   | Implementation depends on institutional capacity and independence.         |

### 3. OBJECT OF STUDY

The primary focus of this research is the fairness and accountability of automated scholarship selection systems. The study examines how application data, eligibility rules, predictive models, ranking mechanisms, digital interfaces, institutional practices, and review procedures can influence the probability that eligible students are recognised, ranked, selected, or excluded. The objects of study are organised below.

#### 3.1 SCHOLARSHIP SELECTION STAGES UNDER INVESTIGATION

- Online application submission, identity validation, and document completeness checks
- Eligibility screening for income, residence, academic, social-category, or programme-specific criteria
- Automated scoring, prioritisation, ranking, and award recommendation processes
- Verification of supporting documents, translations, declarations, and exceptional circumstances
- Communication of decisions, human review, appeals, corrections, and post-award monitoring

#### 3.2 BIAS AND EQUITY CONTEXTS

- Historical inequality in school resources, academic preparation, documentation, and digital access
- Use of proxy variables such as school type, geography, language, device access, or household records
- Errors in data matching, document interpretation, income classification, and eligibility thresholds
- Unequal effects on applicants facing disability, displacement, migration, care responsibilities, or unstable connectivity
- Opacity in automated rankings, award recommendations, and reasons for non-selection

#### 3.3 FAIRNESS EVALUATION PARAMETERS

- Validity and relevance of the data used to represent merit, need, potential, and eligibility
- Comparative error rates, selection rates, and appeal outcomes across relevant applicant groups
- Clarity of programme rules, scoring criteria, model purpose, and human decision authority

- Availability of accessible assistance, multilingual support, reasonable accommodations, and offline alternatives
- Feasibility of audit, correction, explanation, complaint handling, and independent review

### 3.4 COMPARATIVE ANALYSIS FOCUS

- Rule-based eligibility filters, statistical risk scores, ranking models, and hybrid human-machine workflows
- Bias originating in data collection, feature selection, model design, interface design, and deployment practice
- Individual fairness, group fairness, substantive equity, procedural justice, and accessibility requirements
- Institutional responsibilities of scholarship providers, technology vendors, review committees, and regulators

### 3.5 RESEARCH OBJECTIVE SCOPE

- To identify the main forms of algorithmic bias relevant to scholarship selection
- To explain how automated systems can alter access to educational funding and opportunity
- To assess gaps in transparency, human oversight, appeal rights, and data governance
- To recommend safeguards that support equitable, accountable, and student-centred scholarship administration

## 4. RESEARCH METHODOLOGY

The study adopts a comparative and evidence-based research methodology. It synthesises published scholarship on algorithmic fairness, educational technology, public decision-making, and responsible artificial intelligence with practical features commonly found in scholarship administration. The research does not present a new dataset of applicants, a clinical experiment, or a claim of measured discrimination within a named scholarship programme. Instead, it develops an analytical framework for identifying risk before and during implementation.

The analysis follows the full scholarship decision pathway: an applicant accesses the system, submits information and documents, is screened for eligibility, may be scored or ranked, receives an award recommendation or rejection, and is provided with information about review or appeal. Each stage is examined through common dimensions: data

quality, representativeness, proxy discrimination, accessibility, model transparency, opportunity for human judgement, error correction, and accountability. This approach recognises that a technically well-performing model can still be unfair when it is applied to incomplete data, used beyond its intended purpose, or embedded in an inaccessible administrative process.

The methodological framework also distinguishes between legitimate programme criteria and harmful exclusion. Scholarship bodies may lawfully and ethically pursue specified objectives, such as supporting economically disadvantaged students, a particular field of study, rural learners, students with disabilities, or under-represented communities. The concern is not that all differentiation is unfair. The concern is whether eligibility conditions, predictive features, and automated decisions genuinely serve the stated objective, are applied consistently, and do not impose unnecessary or unjustified burdens on particular groups.

### METHODOLOGY FRAMEWORK

1. The study examines automated scholarship selection as a socio-technical decision system rather than as a stand-alone software tool.
2. Published literature, ethics frameworks, and education-policy guidance are used for qualitative synthesis.
3. The analysis covers data, rules, models, interfaces, human review, communication, and appeal mechanisms.
4. Selection stages are compared according to the risk of exclusion, unequal treatment, error, and opacity.
5. Sensitive attributes and non-sensitive proxy variables are both considered as potential sources of unfair outcomes.
6. The framework evaluates individual fairness, group-level disparity, substantive equity, and procedural justice.
7. Accessibility for applicants with low connectivity, language barriers, disabilities, and limited documentation is included.
8. Institutional policy objectives are distinguished from technical performance measures and vendor claims.

9. No fabricated selection rates, applicant characteristics, or statistical findings are presented in the analysis.
10. Recommendations emphasise prevention, documented governance, human accountability, and meaningful remedy.

## 5. DATA ANALYSIS

The data analysis section offers a qualitative comparison of the principal risk points within automated scholarship selection. It does not report prevalence estimates because the severity of risk depends on the scholarship purpose, jurisdiction, applicant population, data sources, funding rules, and administrative capacity of the institution. Instead, the analysis evaluates the relationship between each decision stage, the information used at that stage, the form of bias that may occur, the likely effect on applicants, and the safeguards that should be in place.

Six decision stages are considered: access and application intake, eligibility screening, document verification, scoring or ranking, award allocation, and communication with

appeal. These stages should not be treated as independent. A disadvantage created at the application stage can remain invisible in later stages because a student who cannot upload a required document, navigate a portal, or interpret a form may never reach the ranking process. Likewise, a fairer model cannot correct the harm caused by a non-transparent rejection notice that prevents an applicant from identifying and correcting an error.

The comparison shows that the most serious risks occur where automation affects a high-stakes decision without a clear explanation, accessible human review, or timely correction route. Rule-based tools can wrongly exclude applicants when information is incomplete or atypical. Statistical models can reproduce historical patterns when they use data that reflect unequal opportunity. Allocation engines can amplify disadvantage when they optimise for narrow financial or administrative targets rather than the scholarship programme's equity commitments.

### 5.1 COMPARATIVE RISK ANALYSIS

| S. No. | Decision Stage        | Illustrative Inputs                                  | Bias Pathway                             | Applicant Effect                                       | Control Measure                                                 | Risk Level | Priority  |
|--------|-----------------------|------------------------------------------------------|------------------------------------------|--------------------------------------------------------|-----------------------------------------------------------------|------------|-----------|
| 1      | Application intake    | Portal access, file format, identity fields          | Digital access and form-design bias      | Eligible students may not complete a valid application | Assisted channels, offline option, usability testing            | High       | Immediate |
| 2      | Eligibility screening | Income proof, residence, marks, category documents   | Rigid rules and data matching errors     | False exclusion or delayed review                      | Plain-language rules, correction window, human exception review | Very High  | Immediate |
| 3      | Document verification | Scanned records, translated certificates, statements | Language, OCR, and document-quality bias | Evidence may be misread or treated as invalid          | Multilingual support, confidence flags, manual verification     | High       | High      |
| 4      | Scoring and ranking   | Academic history, school context, need indicators    | Historical bias and proxy features       | Unequal scores or reduced recommendation probability   | Feature review, fairness tests, policy-led criteria             | Very High  | Immediate |

|   |                   |                                                  |                                                   |                                                               |                                                          |           |           |
|---|-------------------|--------------------------------------------------|---------------------------------------------------|---------------------------------------------------------------|----------------------------------------------------------|-----------|-----------|
| 5 | Award allocation  | Rank position, budget rules, predicted enrolment | Optimisation for cost or yield rather than equity | Lower award amount or missed funding opportunity              | Equity constraints, committee approval, allocation audit | High      | High      |
| 6 | Notice and appeal | Reason codes, status updates, complaint records  | Opaque outcome and automation bias                | Applicant cannot understand, challenge, or correct a decision | Specific reasons, human appeal, independent review       | Very High | Immediate |

## 5.2 EXTENDED FAIRNESS EVALUATION

Data Representation and Measurement Bias. Scholarship systems frequently draw on academic records, income documents, family circumstances, school information, location data, and written statements. None of these indicators is neutral in every context. A score based on prior academic performance can reflect uneven access to teachers, tutoring, stable housing, learning devices, or disability accommodations. Household income may be harder to document for informal workers, migrants, agricultural households, or families affected by sudden disruption. The relevant question is not whether data exist, but whether they measure the stated criterion consistently across applicants. Institutions should document why each field is necessary, what it represents, and how missing or atypical evidence will be handled.

Proxy Discrimination in Feature Design. Directly removing protected attributes is an insufficient safeguard because many apparently ordinary variables can act as proxies. School type, postal area, language of submission, device type, attendance history, travel distance, fee payment record, or the timing of an application may correlate with structural advantage or disadvantage. Proxy variables should be reviewed not only for statistical correlation but also for social meaning. A feature that improves predictive performance may still be inappropriate when it captures access to privilege rather than the applicant's potential or need. The system should use the least intrusive and least discriminatory information capable of serving the legitimate programme objective.

Model Objective and Fairness Trade-offs. Automated selection can be designed for different objectives: identifying likely academic completion, distributing limited funds, reducing manual workload, detecting anomalies, or maximising enrolment. These objectives are not interchangeable. A model trained to predict retention may disadvantage students who historically faced barriers to continuation, even where additional scholarship support could improve their prospects. A model that forecasts enrolment yield may steer awards toward applicants who are most likely to accept an offer rather than toward applicants with the greatest need. Institutions must define the policy objective first and test whether the technical objective actually advances it. As Binns (2018) notes, fairness choices cannot be resolved solely by optimisation.

Automation Bias and Human Review. Human reviewers may place excessive trust in an automated score because it appears objective, consistent, or mathematically sophisticated. This is often described as automation bias. A meaningful human-in-the-loop process is not achieved merely by allowing a staff member to click approve after receiving a recommendation. Reviewers require authority to question the result, access to the main reasons behind the recommendation, training on known limitations, and time to assess exceptional circumstances. Human oversight should be especially strong for borderline cases, adverse decisions, low-confidence document matches, and applicants whose evidence does not fit standard categories.

Accessibility, Communication, and Remedy. Fairness is also shaped by how applicants interact with the system. A portal that is

inaccessible to screen-reader users, poorly designed for low-bandwidth connections, available only in one language, or difficult to use on mobile devices can create exclusion before any algorithm operates. Decision letters should state the outcome, the applicable criteria, the main factors considered, available correction steps, and appeal deadlines in understandable language. Applicants need a realistic opportunity to challenge incorrect data, submit additional evidence, request accommodation, and receive review from a responsible human authority. Without remedy, an institution cannot credibly claim that an automated process is accountable.

### 5.3 OBSERVATIONS FROM DATA ANALYSIS

- Bias in scholarship selection is frequently cumulative: an applicant may face small disadvantages at form completion, document review, scoring, and appeal rather than one visible discriminatory event.
- Historical award data should not be treated as a neutral description of merit or need because it may contain the effects of earlier unequal access and discretionary practice.
- Fairness testing must examine both aggregate patterns across relevant groups and individual cases in which a decision appears inconsistent, inaccurate, or poorly explained.
- A technically accurate prediction may still be unsuitable when it uses a proxy for social advantage, conflicts with programme objectives, or cannot be explained to applicants.
- Human review is effective only when reviewers have authority, training, access to relevant context, and sufficient time to depart from an automated recommendation.
- Transparency requires more than publishing a general statement about artificial intelligence; applicants need meaningful information about criteria, use of automation, and routes for correction.
- Vendor systems, internal teams, scholarship committees, and senior institutional leaders share responsibility for the fairness of a deployed system.

### 6. DISCUSSION

The analysis demonstrates that automated scholarship selection should be governed as a high-impact educational decision. An award or rejection can influence access to tuition, accommodation, learning materials, mobility, professional pathways, and the ability to continue study. The institutional consequences of error are therefore substantial. A system that saves administrative time but produces unexplained or unequal exclusions can undermine the public purpose of the scholarship itself. Governance should begin with a clear written statement of the programme's aims, the legitimate criteria that follow from those aims, and the limits on automated decision-making.

Fairness cannot be reduced to the absence of overt discrimination. A scholarship model may never receive a column labelled with a protected characteristic and still generate unequal results through data that are shaped by social inequality. For example, academic ranking may reflect unequal schooling conditions, a financial-need classifier may misread informal income, and an application-completeness rule may penalise students without stable digital resources. The solution is not to reject all technology. It is to ensure that automation is proportionate to the task, supported by reliable data, tested against foreseeable harms, and combined with accessible non-automated alternatives.

Institutions should establish a governance structure before procurement or deployment. This structure should identify a responsible senior owner, maintain a record of data sources and decision rules, require impact assessment, specify audit intervals, and provide a process for responding to discovered disparities. The NIST AI Risk Management Framework (2023) and UNESCO's Recommendation on the Ethics of Artificial Intelligence (2021) emphasise risk management, transparency, fairness, and human oversight. In the scholarship context, these principles can be translated into practical controls: applicant-facing notices, documented purpose limitation, fairness review before release, clear rejection reasons, appeal rights, and independent monitoring.

Participation is equally important. Students, disability support staff, school counsellors,

financial-aid officers, community organisations, and legal or ethics experts may identify barriers that a development team does not see. Consultation should occur early enough to influence the choice of inputs and workflow, not only after the model has been built. Regular review should examine not only whether the system operates as designed but whether the design itself remains appropriate when application patterns, institutional policies, technologies, or social conditions change. A responsible scholarship system treats applicants as rights-bearing participants in a public-interest process rather than as data points to be optimised.

## 7. CONCLUSION

This study examined algorithmic bias in automated scholarship selection systems through a comparative, evidence-based framework. The findings show that unfair outcomes can emerge at every stage of the scholarship pathway: access to the application, validation of records, eligibility screening, scoring, ranking, award allocation, notification, and appeal. The most serious risks arise when an automated system influences a high-stakes outcome without clear policy justification, accessible human review, meaningful explanation, or timely remedy.

Scholarship providers should recognise that fairness is both a technical and an institutional

responsibility. Representative data, careful feature selection, statistical testing, and regular audits are necessary, but they are not sufficient on their own. Equitable scholarship administration also requires programme criteria that reflect the intended social purpose, inclusive application design, multilingual and disability-accessible support, trained reviewers, reasoned decisions, appeal routes, and accountable leadership. Human oversight must be substantive rather than symbolic.

The study concludes that automation can assist scholarship administration only when it strengthens, rather than weakens, educational equity. Systems should be designed to identify administrative errors, support consistent review, and make decisions more understandable. They should not silently replicate historical privilege, exclude applicants whose circumstances are difficult to standardise, or transfer institutional responsibility to a vendor or algorithm. Future research should examine longitudinal outcomes, applicant experiences, group-level selection patterns, effectiveness of appeal processes, and the impact of fairness interventions in real scholarship programmes. Such evidence can support the development of systems that are efficient, transparent, and genuinely aligned with the educational promise of opportunity.

## 8. REFERENCES

1. Barocas, S., & Selbst, A. D. (2016). *Big data's disparate impact*. *California Law Review*, 104, 671-732.
2. Binns, R. (2018). *Fairness in machine learning: Lessons from political philosophy*. *Proceedings of Machine Learning Research*, 81, 149-159.
3. Chouldechova, A. (2017). *Fair prediction with disparate impact: A study of bias in recidivism prediction instruments*. *Big Data*, 5(2), 153-163.
4. Eubanks, V. (2018). *Automating Inequality: How High-Tech Tools Profile, Police, and Punish the Poor*. New York: St. Martin's Press.
5. European Commission. (2019). *Ethics Guidelines for Trustworthy AI*. Brussels: High-Level Expert Group on Artificial Intelligence.
6. Jobin, A., Ienca, M., & Vayena, E. (2019). *The global landscape of AI ethics guidelines*. *Nature Machine Intelligence*, 1(9), 389-399.
7. Kizilcec, R. F., & Lee, H. (2020). *Algorithmic fairness in education*. *arXiv:2007.05443*.
8. Kleinberg, J., Mullainathan, S., & Raghavan, M. (2017). *Inherent trade-offs in the fair determination of risk scores*. *Proceedings of the 8th Innovations in Theoretical Computer Science Conference*.
9. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., & Galstyan, A. (2021). *A survey on bias and fairness in machine learning*. *ACM Computing Surveys*, 54(6), 1-35.
10. National Institute of Standards and Technology. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. NIST AI 100-1. Gaithersburg, MD: NIST.

11. OECD. (2019). *Recommendation of the Council on Artificial Intelligence*. OECD/LEGAL/0449. Paris: Organisation for Economic Co-operation and Development.
12. O'Neil, C. (2016). *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. New York: Crown.
13. Raji, I. D., Smart, A., White, R. N., et al. (2020). Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33-44.
14. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 59-68.
15. Suresh, H., & Guttag, J. V. (2021). A framework for understanding sources of harm throughout the machine learning life cycle. *Proceedings of the 2021 ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization*, 1-9.
16. UNESCO. (2021). *Recommendation on the Ethics of Artificial Intelligence*. Paris: United Nations Educational, Scientific and Cultural Organization.
17. UNESCO. (2023). *Guidance for Generative AI in Education and Research*. Paris: United Nations Educational, Scientific and Cultural Organization.
18. Wachter, S., Mittelstadt, B., & Russell, C. (2021). Why fairness cannot be automated: Bridging the gap between EU non-discrimination law and AI. *Computer Law & Security Review*, 41, 105567.