Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Contestability without recourse: post-decision safeguards for disabled learners in AI-supported educational decision-making

Jonas A. Mandalunes

jonas.mandalunes@perpetual.edu.ph

University of Perpetual Help System DALTA

Systematic Review

Keywords: contestability, algorithmic recourse, disability, special educational needs, automated decision-making, educational technology

Posted Date: August 28th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10621652/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Artificial intelligence increasingly informs consequential decisions about learners with disabilities and special educational needs. Research has concentrated on whether such decisions are made fairly. This article argues that a prior question has been displaced. Fairness concerns how a decision is produced; contestability concerns what its subject can do about it afterwards. A decision may be demonstrably fair and still leave the person it concerns with no route to change it.

We define a post-decision safeguard as a mechanism operating after an AI-supported educational decision through which a learner or their representative can change or contest it, and distinguish it from three things with which it is routinely conflated: explainability, human involvement in making the decision, and pre-deployment audit. Human oversight in particular is a control held by the operator, not a route out of the decision for the person affected.

Data protection law, the EU AI Act and the Convention on the Rights of Persons with Disabilities presume some route to challenge, while none specifies a school-level procedure through which a learner could use one. Across 138 full texts obtained through a search combining artificial intelligence, disability and education terms, terms denoting bias or fairness appeared near a decision in 74.6% of documents, while terms denoting appeal or due process appeared in none. The result identifies a marked asymmetry between attention to decision quality and attention to post-decision recourse within the retrieved corpus.

1 Introduction

1.1 The decision has already been made

A fourteen-year-old is moved to a lower set in mathematics. The decision drew on a predictive model of attainment, one input to which was attendance, and her attendance record is wrong: a fortnight of hospital appointments was logged as unauthorised absence. Her mother notices the placement, not the model. When she asks how the decision was reached, she is told the school uses a system that flags students at risk. When she asks how to change it, nobody knows. The account is illustrative. It is a constructed example, not a report of an identifiable individual.

Nothing in that account requires the model to be biased. It may be well calibrated across every protected characteristic and validated on a representative sample. The failure is not in how the decision was produced but in what happens afterwards, and the literature that would help the school answer the first question has little to say about the second.

That asymmetry is the subject of this article.

1.2 What the field has built

Research on artificial intelligence in education has developed a substantial apparatus for evaluating how decisions are made. Algorithmic bias in education is now a well-mapped field (Baker and Hawn 2022),

with work extending to student progress monitoring specifically (Idowu et al. 2024) and to the framing of responsible practice (Fu and Weng 2024). Reviews of fairness, accountability, transparency and ethics in educational AI have appeared in quick succession (Memarian and Doleck 2023; Wiese et al. 2025), and national AI policy strategies have been examined for how they treat education and ethics (Schiff 2022). Concern about the position of students specifically has been raised in this journal (Smith 2020).

This literature has established important standards for evaluating algorithmic decisions. Contestability concerns a different stage of the decision process.

1.3 Contestability is a different thing

A parallel literature, largely outside education, has begun to distinguish contestability from explainability. Bayamlıoğlu (2022) argues that the right to contest under the General Data Protection Regulation is analytically prior to, and more demanding than, the much-discussed right to explanation. Henin and Le Métayer (2022) similarly place justifiability and contestability beyond explainability. Grant et al. (2025) ask what is owed to decision-subjects and conclude that transparency and explanation do not exhaust it. Karimi et al. (2023) survey algorithmic recourse as a technical problem, showing that the machinery for telling a person what would have to change is well developed.

Contestable AI by design has been proposed as a research programme, and human oversight has been examined critically: Laux (2024), in this journal, reads oversight requirements under the EU AI Act as institutionalised distrust, while Green (2022) argues that policies requiring human oversight of government algorithms are flawed in ways their proponents do not acknowledge.

Little of this work concerns education, and less concerns disabled learners. That population is subject to frequent algorithmic and psychometric assessment, because assessment is how support is allocated, and in many jurisdictions holds statutory rights of appeal over identification, provision and placement.

1.4 What this article does

The article defines post-decision safeguards, examines their legal and normative foundations, measures their presence in a corpus of research concerning AI-supported educational decisions affecting disabled learners, and searches separately for the literature on recourse for that population in order to establish where it sits.

First, we define post-decision safeguards precisely enough to be identified or found absent, and specify what does not count. The exclusions carry most of the analytical weight, because the concepts excluded — explainability, human involvement, pre-deployment audit — are precisely those a reader is likely to accept as substitutes.

Second, we set the construct against three legal instruments and show that each presumes some route to challenge while none specifies a school-level procedure through which a learner could use one.

Third, we report a measurement. Across 138 research reports retrieved through a structured search, we measured how often each of a set of pre-decision concepts and post-decision mechanisms appears in proximity to a decision. Terms denoting bias or fairness appeared in proximity to a decision in 74.6% of documents; terms denoting appeal or due process in none.

The comparison reported here is internal to the retrieved corpus. Its scope and methodological limits are addressed in Section 7.

1.5 What follows

Section 2 defines the construct and its exclusions. Section 3 sets out the legal position under the GDPR, the AI Act and the CRPD, and identifies the gap between what these instruments presume and what educational practice provides. Section 4 describes how the corpus was assembled and measured. Section 5 reports the measurement. Section 6 draws out four design requirements and considers what the pattern implies about the field's priorities. Section 7 states the limits of the empirical claim, and Section 8 concludes.

2 Defining post-decision safeguards

Because this article's contribution depends on the boundary of the construct, we state it before the methods.

The construct is not invented here. Contestability has been developed as a design property in its own right: Alfrink et al. (2023) set out a framework for contestable AI by design, and Ploug and Holm (2020) specify four dimensions along which a patient must be able to contest an AI-supported diagnosis — the data used, the system's reasoning, the professional's role and the availability of a second opinion. Both treat the capacity to object as something a system must be built to afford rather than a by-product of accuracy or transparency. Our contribution is to ask whether that property has been described in educational systems affecting learners with disabilities, and to define it tightly enough that the question can be answered from published reports.

A **post-decision safeguard** is a mechanism that operates after an AI-supported educational decision has been made and that enables the learner, or a person acting on their behalf, to change or contest that decision. Five forms qualify:

Human reconsideration. A named person or body can review the decision on request and substitute their own. Automatic human review that occurs whether or not anyone asks does not qualify unless the learner can also trigger it.

Data correction. The learner or representative can require that the inputs on which the decision rested be amended, and the decision revisited on the corrected basis. Correction without reconsideration is a data-protection right rather than a safeguard against the decision.

Appeal. A route exists to a decision-maker other than the original one, with some statement of what will be considered.

Redress or remedy. Where a decision has caused detriment, something can be done about the detriment — restoration, compensation or reversal.

Explanation provided in order to enable challenge. An account of the basis for the decision, furnished to the affected person, sufficient for them to identify a ground on which to dispute it.

2.1 What counts as a consequential educational decision

The construct applies only to decisions with material educational consequences, so that threshold is defined rather than illustrated.

A decision is **consequential** where it materially affects access, admission, placement, grouping or setting, grading, progression, discipline, accommodation, allocation of support, certification or any other outcome bearing on the learner's educational rights or trajectory. The test is whether a different decision would have produced a materially different educational position for that learner.

Excluded on this test are low-stakes personalisation and ordinary instructional recommendation: the order in which practice items are presented, the pacing of material within a lesson, or a suggested optional resource. These may be produced by the same system and even by the same model output, but they do not by themselves alter the learner's position. Where adaptive difficulty determines which curriculum or qualification a learner can access, it is consequential; where it varies the next exercise, it is not.

We record throughout the number of **reports** rather than studies, because a single study may be described across several reports. Because no instance of multiple reports of one study was identified in this corpus, the two counts coincide here at 138, but the term used is reports.

2.2 What is excluded, and why

Because the result is sensitive to where the construct's boundary falls, the exclusions are specified explicitly.

Accuracy and fairness metrics describe a system's aggregate behaviour. They are not available to an individual as a means of changing an individual outcome.

Pre-deployment bias audits occur before any decision exists.

Consent at the point of collection concerns whether data may be gathered, not what follows from conclusions drawn from it. Consent is also a weak instrument in compulsory education, where the practical alternative to consenting is often nothing.

Privacy and security controls govern who may see data.

Teacher involvement in making the decision is the exclusion most likely to be contested, and the one on which the finding reported here most depends. A teacher who reviews a model's output before acting on it is exercising professional judgement within the decision process. That is desirable, and it is not recourse. Recourse requires that someone outside the decision can be made to look again at the instance of the affected party. A teacher in the loop may be the person a parent complains to, but unless a mechanism obliges reconsideration, the parent's position rests on that teacher's goodwill.

There is now empirical reason to doubt that human involvement performs the protective work attributed to it. Alon-Barkat and Busuioc (2023) find that public-sector decision-makers exhibit selective adherence to algorithmic advice, following it where it aligns with prior expectations and discounting it otherwise — a pattern that leaves the affected person worse off precisely where the model errs in the direction of existing bias. Cavalcante Siebert et al. (2023) argue that meaningful human control requires specific system properties rather than the mere presence of a person. On both accounts, human involvement during decision-making does not by itself give an affected learner a route to challenge the outcome afterwards.

General explainability without a route to challenge is excluded for a related reason. Interpretability serves developers, auditors and researchers. An explanation becomes a safeguard only when it is given to the affected person for the purpose of enabling them to object — which is the distinction Ploug and Holm (2020) build their framework around.

Section 5.4 reports the permissive automated measure separately.

3 Legal and normative context

3.1 Article 22 GDPR and the right to contest

Article 22 of the General Data Protection Regulation establishes that a data subject has the right not to be subject to a decision based solely on automated processing which produces legal effects or similarly significantly affects them. Article 22(3) requires suitable measures to safeguard rights, expressly including the right to obtain human intervention, to express a point of view and to contest the decision, for decisions permitted under Article 22(2)(a) and (c). Decisions authorised by Union or Member State law under Article 22(2)(b) are governed instead by the safeguards laid down in that law.

The scholarly debate on Article 22 has concentrated on whether it confers a right to explanation. Bayamlıoğlu (2022) argues that this emphasis is misplaced: the operative protection is the right to contest, and explanation matters instrumentally, as a precondition of contesting rather than as an end. Hamon et al. (2022) examine how far explainability techniques can in practice discharge the Regulation's requirements, and Custers and Heijne (2022) show that even the narrower right of access under Article

15(1)(h) is inconsistently realised. Grant et al. (2025) generalise the point philosophically: what is owed to decision-subjects is a capacity to respond, of which transparency is only a component.

Two features of the provision matter for education. The first is the word "solely". Where a teacher acts on a model's recommendation, the decision is arguably not solely automated, and the Article 22 protections may not engage. Van Kolfschooten (2024), analysing the equivalent problem in healthcare, reads the provision purposively to resist exactly this narrowing. The second is that the right to contest presupposes a mechanism through which contesting occurs, which the Regulation does not specify.

3.2 The EU AI Act

The AI Act designates as high risk certain systems used in education, including those determining access to institutions, evaluating learning outcomes and assessing appropriate levels of education. For high-risk systems it establishes obligations of transparency and human oversight under Articles 13 and 14. Article 86 provides a right to explanation for a person affected by a decision taken on the basis of certain Annex III high-risk systems where that decision produces legal or similarly significant adverse effects.

Laux et al. (2024) argue that the Act conflates trustworthiness with the acceptability of risk, so that compliance can be achieved without the affected person gaining anything. Laux (2024) develops a related critique of human oversight specifically, arguing that oversight functions as institutionalised distrust of the system rather than as empowerment of the person the system decides about. Both critiques bear directly on the finding reported here.

3.3 The UNCRPD

Article 24 of the Convention on the Rights of Persons with Disabilities obliges states parties to ensure an inclusive education system at all levels, with reasonable accommodation of individual requirements. Article 13 provides for effective access to justice. Read together, these provisions may support an argument that decisions bearing on a disabled learner's education should be susceptible to challenge by or on behalf of that learner. That reading is an interpretation rather than explicit treaty language. Langford (2020) situates automated decision-making within the international human rights framework more broadly, and Nilholm (2021) notes that inclusive education research has tended to under-theorise the mechanisms through which inclusion is actually secured or denied in practice.

3.4 The gap this creates

Each instrument bears on recourse differently. The General Data Protection Regulation is directly applicable in member states and Article 22(3) expressly requires safeguards including human intervention and the right to contest, for decisions within its scope. The AI Act provides a complaint route and, for certain decisions, an explanation. The reading of the Convention offered above is interpretive rather than an explicit entitlement. What none of them supplies is a school-level procedure specifying how a learner or their representative initiates reconsideration and what follows from it.

Whether recourse exists in practice is therefore an empirical question about the systems and processes actually deployed — which is the question this article addresses to the research literature.

4 Corpus and measurement

The corpus analysis compares the presence of pre-decision and post-decision concepts across the retrieved reports.

4.1 What we needed the corpus to do

The claim we wanted to test is comparative: that a literature attending closely to how algorithmic educational decisions are made attends far less to what follows them. Testing that does not require a complete sample of the field. It requires a body of documents assembled without reference to either concept, so that neither is favoured by the way the corpus was built, and then the same instrument applied to both.

Our search terms concerned artificial intelligence, disability or special educational needs, and education. They contained no term for fairness, bias, appeal, redress or contestation. The study was designed as a within-corpus comparison rather than a prevalence estimate. The search did not directly select on fairness or recourse terminology.

4.2 Sources and search

Twelve bibliographic sources were searched: Crossref, OpenAlex, Europe PMC, Scopus, OpenAIRE, PubMed, DOAJ, ERIC, Semantic Scholar, CORE, DBLP and arXiv. Searches were executed on 4 and 5 August 2026, all sources within that window, with no subsequent update. A publication floor of 2015 was applied.

The ACM Digital Library and IEEE Xplore were not searched through their own interfaces: the former prohibits automated retrieval under its terms of use and the latter requires a developer key we did not hold. Records from both publishers were instead reached through Crossref, filtered by the digital object identifier prefixes those publishers register. This covers the metadata they deposit rather than their own search indexes, so coverage of these two publishers is partial.

The full query expression for every search, with the filters applied, the number of records returned and a checksum of each raw export, is given in Online Resource 6.

4.3 From records to documents read

After deduplication on normalised digital object identifier and normalised title, 1,690 unique candidate records were tiered by whether artificial intelligence, disability and education could each be verified in the record's own metadata. The 727 records in which all three were verified were screened at title and abstract by two reviewers working independently, with a third resolving disagreements. Agreement was 96.6%, Cohen's kappa 0.929.

Screening advanced 430 records to full text. Of those, 138 full texts were obtained and read. The remaining 292 could not be obtained through open-access routes, which was the retrieval condition we set ourselves. Figure 1 sets out the flow from the deduplicated candidate pool to the analysed corpus.

The remaining 963 records were not screened because their metadata did not permit verification of all three concepts, most often because no abstract was available.

Both reviewers' decisions for all 727 records, the adjudication of every disagreement and the retrieval outcome for the 430 records advanced are given in Online Resource 5.

4.4 What we counted, and how

Each of the 138 documents was searched for terms belonging to six pre-decision concepts and five post-decision mechanisms, and a term was recorded as occurring in a decision context when it fell within 240 characters of a decision term. The concept vocabularies and the list of decision terms were fixed before measurement and are given in full in Online Resource 4, which also contains the per-report measurement output.

The same 240-character window was applied to every concept. All passages flagged for post-decision mechanisms were then inspected manually to determine whether they described learner-invocable recourse.

Alongside the automated measurement, both reviewers recorded whether each report described a post-decision safeguard meeting the definition in Section 2, with a verbatim quotation and a locator required for any positive finding.

4.5 A second, targeted search

The corpus search described above contained no term for appeal, redress or contestation, so that neither side of the comparison reported in Section 5 was favoured by the way the corpus was assembled. That design leaves a question open: whether a literature on recourse for disabled learners exists elsewhere and was therefore outside the corpus by construction.

A second search was run to answer it, on 6 August 2026, and is reported separately. Crossref and OpenAlex were queried with 80 paired terms, combining sixteen expressions from the recourse family with five expressions denoting the population or the educational setting. The full query list and the per-query return counts are given in Online Resource 2.

The searches returned 2,456 records, 1,297 distinct by digital object identifier. Records were retained where the record's own title contained both a recourse term and a population or education term, which admitted 114. Of the remainder, 915 lacked a recourse term in the title and 268 lacked a population or education term.

Inspection of the retained titles showed that 12 used *remedy* or *repair* in a general rather than a procedural sense, as in education as a remedy for a misconception or open resources as a remedy for cost. These were set aside, leaving 102 records whose titles linked recourse terminology with disability or education.

This search selects on the concept the corpus analysis measures, and its results are reported separately.

4.6 Reliability of the full-text judgement

Both reviewers independently assessed all 138 full texts for the presence of a post-decision safeguard meeting the definition in Section 2. A positive judgement required identification of the safeguard type, a supporting quotation and a document locator.

Reviewer 1 recorded two positive judgements and reviewer 2 recorded three, with one report in common. Treating the judgement as binary, the reviewers agreed on 135 of 138 reports, giving observed agreement of 97.8%. Cohen's kappa was 0.389 and the prevalence- and bias-adjusted kappa was 0.957; the difference between the coefficients reflects the highly imbalanced distribution of ratings.

Four reports received a positive judgement from at least one reviewer. Two were referred to the third reviewer, who resolved both as not describing a safeguard. The remaining two record capability held by a teacher and are reported in Section 5.3 as adjacent material rather than as meeting the strict definition. The construct definition was fixed before coding and was not revised during assessment. The coding table is given in Online Resource 3.

4.7 Use of generative AI

The author used OpenAI ChatGPT 5.6 Sol for language editing the English-language text and to assist with revising and debugging the Python 3.12 scripts. The author reviewed and verified all code, analytical outputs and manuscript text and takes full responsibility for the final work.

5 What the corpus shows

5.1 The contrast

Table 1 **describes the corpus.** Table 2 **reports how often each concept appeared within 240 characters of a decision term across the 138 documents, and** Fig. 2 **shows the same measurement alongside the wider presence of the post-decision terms.**

Table 1
Publication year of the analysed corpus (n = 138). Document type and publication venue as recorded in the source metadata are given in Online Resource 1. Country or jurisdiction, educational setting, disability or special educational needs group, AI function, decision type and system maturity were not extracted, which limits subgroup interpretation

Publication year	Reports	%
2015–2019	4	2.9
2020–2022	9	6.5
2023–2024	22	15.9
2025–2026	103	74.6
Total	138	100.0

Table 2
Concepts appearing within 240 characters of a decision term across the analysed corpus (n = 138)

Concept	Documents	%
Bias or fairness	103	74.6
Ethics, general	98	71.0
Privacy or data protection	88	63.8
Accuracy or performance	67	48.6
Explainability	27	19.6
Consent	23	16.7
Human reconsideration	9	6.5
Redress or remedy	3	2.2
Data correction	2	1.4
Appeal or challenge	0	0.0
Due process or rights	0	0.0

Three quarters of a corpus assembled without any reference to fairness nonetheless discusses fairness in proximity to a decision. Appeal and due process do not appear in that position in a single document.

These frequencies describe the retrieved corpus rather than field-wide prevalence.

5.2 Vocabulary that appears and then detaches

The post-decision mechanisms are not uniformly absent from the documents. They are present and unattached. Table 3 sets out the difference between mentions anywhere in a document and mentions in a decision context.

Table 3
Post-decision mechanisms: mentions anywhere in the document compared with mentions in a decision context (n = 138)

Mechanism	Mentioned anywhere	Within 240 characters of a decision term	Elsewhere in the document
Redress or remedy	18	3	15
Human reconsideration	16	9	7
Data correction	5	2	3
Due process or rights	4	0	4
Appeal or challenge	3	0	3

The language of recourse is not unknown to this literature. It occurs, and then does not occur near a decision. For appeal and due process the attrition is total: every occurrence in the corpus is detached from any decision term.

This distinguishes two situations that a flat count of zero would conflate. A concept genuinely foreign to a field and a concept invoked in the abstract while never being connected to the decisions the field studies are different conditions, and the second is what we observe.

5.3 What reviewers found on reading

Across the 138 documents, no report described a mechanism that a learner or their representative could invoke to change or contest a decision.

Two reports contained material adjacent to the construct. Ochoa Barrera et al. (2025) record teachers requesting greater control over AI decisions, including manual override of content difficulty. Tsirantonaki and Vlachou (2025) recommend that teacher training address the critical assessment of AI tools. In both cases the capability described belongs to the teacher. Neither describes a route by which the learner or their family could act. On our definition these are operator-side oversight, and we count them separately rather than as safeguards.

The three coding thresholds produced the following results:

- **strict**: no report describes a mechanism a learner or representative can invoke — 0 of 138 (0.0%)

- **broad**, admitting operator override and general oversight — 2 of 138 (1.4%)
- **permissive**, automated term proximity for any post-decision mechanism — 12 of 138 (8.7%)

5.4 What the permissive count contains

The permissive measure flagged 12 reports. Reading the flagged passages, nine describe professional oversight of a system rather than a route available to a learner, in formulations such as the retention of meaningful human oversight or the requirement that a tool should not override pedagogical judgement. One describes teachers requesting manual override of content difficulty. In two, the matched term concerns the correction of bias during system development rather than correction of a learner's record. Ten flagged documents were re-inspected directly. The classifications of the remaining two were confirmed from the contemporaneous coding records because the full texts were no longer available.

None of the matches described learner-invocable recourse. The permissive measure therefore produced false positives and was not used as evidence for the strict count.

5.5 What the corpus discusses instead

The upper rows of Table 2 are not a foil. They describe a field doing serious work on a real problem. Bias in educational algorithms is well documented (Baker and Hawn 2022), extends to progress monitoring (Idowu et al. 2024), and has been the subject of repeated synthesis (Memarian and Doleck 2023; Fu and Weng 2024; Wiese et al. 2025). The concern is not that this attention is misplaced.

The concern is that it is being asked to do work it cannot do. A fair decision and a contestable decision are different achievements. Karimi et al. (2023) show that the technical machinery for telling a person what would need to change is well developed; the corpus we read does not connect that machinery to the position of a learner who disagrees with an outcome. Explainability, which appears in 19.6% of documents, is not a substitute either: Henin and Le Métayer (2022) argue that contestability lies beyond explainability, and Bayamlıoğlu (2022) that the right to contest is more demanding than the right to explanation.

Nor is human involvement a substitute. Human reconsideration appears in a decision context in 6.5% of documents, more often than any other post-decision term, and on inspection it consistently describes a capability held by a teacher or clinician. Laux (2024) and Green (2022) both give reason to doubt that oversight arrangements of this kind deliver what they are assumed to deliver. A human in the loop is a control held by the operator. It is not a route out of the loop for the person the decision is about.

5.6 A literature on recourse exists but does not signal algorithmic decision-making in its titles

The targeted search described in Section 4.5 identified 102 records whose titles linked recourse terminology with disability or education. Of these, 36 were published before 2010, 34 between 2010 and 2019 and 29 from 2020 onwards; publication year was not recorded for three.

The subject matter is specific and institutional. Records concern special education due process hearings, the First-tier Tribunal for special educational needs and disability, complaints and grievance procedures, statutory dispute resolution and the experience of disabled people pursuing complaints. Recurrent venues include Exceptional Children, the Journal of Disability Policy Studies, and proceedings of the American Educational Research Association.

No title among the 102 mentions artificial intelligence, algorithms, automated systems, machine learning or predictive analytics. Whether the underlying publications discuss such systems elsewhere in their text was not assessed.

This bears directly on how the corpus result should be read. The absence measured in Section 5.1 is not an absence of scholarly attention to recourse for disabled learners, which is an established and continuing field. It is an absence of connection between that field and the literature on AI-supported educational decision-making. The targeted search therefore suggests limited explicit connection at title level between scholarship on AI-supported educational decisions and scholarship on disability-related educational recourse.

6 What follows

6.1 Four requirements

Four procedural requirements follow from the analysis.

A route the learner can start. Contestability requires that the person the decision concerns, or someone acting for them, can initiate reconsideration. This is the requirement that operator-side oversight does not meet, however conscientiously it is exercised. A teacher who can override a recommendation holds a control. A parent who can trigger a review holds a right. Only the second survives the teacher's absence, disagreement or workload.

Disclosure that the decision was AI-supported. A route cannot be used by someone who does not know there is anything to contest. In the case with which this article opened, the mother noticed a placement, not a model. Contestability presupposes that the involvement of an algorithmic system in a consequential decision is disclosed to the person affected, at the time and in terms they can act on.

Access to and correction of the inputs. Where a decision draws on records, the person concerned must be able to see the records and correct them, and correction must trigger reconsideration rather than merely amend a file. Data correction appeared in a decision context in 1.4% of the documents we read. Attendance, behaviour and prior-attainment records may be particularly vulnerable to misinterpretation where disability has interrupted a learner's schooling.

An outcome that can change. A procedure that receives a challenge and cannot alter a decision is not recourse. This is the requirement that the AI Act's complaint and explanation provisions do not by

themselves meet, and it is where the distinction between explanation and contestation bites hardest: being told why is not being able to change it.

These requirements are primarily procedural rather than dependent on a new computational technique. Algorithmic recourse as a computational problem is well developed (Karimi et al. 2023). What is absent is not the capability but the procedure, and the research attention that would establish whether the procedure exists.

6.2 Why this population

The argument applies to any decision-subject, and we have framed it generally where we can. But learners with disabilities occupy a particular position, and it is worth stating why the absence matters more here.

Identification, support, placement and accommodation are decided through assessment, and in many jurisdictions those decisions carry formal rights of review or appeal. They are cumulative: a placement affects the next assessment, which affects the next placement.

That combination is unusual. A population with formal appeal rights, subject to frequent consequential assessment, whose research literature does not discuss appeal.

6.3 What the pattern suggests about the field

The concepts that dominate the upper rows of Table 2 share a feature: they are properties of a system, and they can be evaluated by the people who build and deploy it. Bias can be audited. Accuracy can be validated. Privacy can be engineered. Ethical governance can be documented. Each is a question a developer or institution can answer about itself.

Contestability is not like that. It requires ceding a decision to someone outside the institution — the learner, the parent — and accepting that the institution may be wrong in a way it did not detect. It is not a property to be verified but a power to be transferred.

A field will tend to develop the questions it can answer with the methods and authority it has. That is not a criticism of anyone's intentions. But it does suggest that the absence we measured is unlikely to close on its own, because nothing in the current research programme requires it to.

6.4 What would settle the empirical question

Our measurement establishes what the research literature discusses. It cannot establish what schools do. The two may diverge in either direction: a school may operate an effective reconsideration practice that nobody has written up, or may operate none while believing statutory appeal rights suffice.

Settling it requires work of a kind this corpus contains almost none of: asking parents and learners whether they knew an algorithmic system was involved, whether they attempted to contest an outcome, and what happened. The methods are ordinary. The question has simply not been put.

7 Scope of the evidence

The reported percentages describe the retrieved corpus rather than field-wide prevalence. Of 430 reports advanced to full text, 138 were obtained and read. A further 963 records with insufficient metadata were not screened, and the ACM Digital Library and IEEE Xplore were not searched directly. Relevant reports may therefore be absent from the corpus.

The proximity measure identifies vocabulary rather than meaning. Although a common window and decision-term vocabulary were applied across concepts, the concept dictionaries may differ in sensitivity and specificity. Manual inspection showed that the post-decision matches included false positives, and the recall of the automated measure was not established. The strict finding rests on the independent full-text assessments.

The analysis concerns safeguards described in research reports and does not establish whether undocumented procedures operate in schools. The legal discussion is limited to the GDPR, the EU AI Act and the CRPD.

Separate determinations against all four substantive eligibility criteria were not recorded for reports coded as containing no safeguard. The corpus is therefore defined as the 138 reports obtained and assessed in full rather than as a fully re-evaluated set satisfying each criterion individually.

8 Conclusion

Across 138 full texts obtained through a search combining artificial intelligence, disability and education terms, bias or fairness terminology appeared near a decision in three quarters of documents, while appeal or due-process terminology appeared in none. Two reports described capabilities adjacent to the construct; in both, the capability belonged to a teacher rather than to the learner.

The retrieved literature devotes substantial attention to how algorithmic decisions are produced but little to what affected learners can do after those decisions are made. The question is separable from the ones the literature does ask. Fairness concerns how a decision is produced. Contestability concerns what its subject can do about it. A field can become highly sophisticated about the first while saying little about the second. In the corpus we retrieved, that is the pattern we observe.

What would change it is neither technically difficult nor conceptually obscure. It requires treating the learner not only as someone about whom decisions should be made well, but as someone who may disagree.

Declarations

Author Contribution

J.A.M. conceived the study, designed the methodology and measurement approach, conducted the literature searches, screening and analysis, interpreted the findings, and wrote and revised the manuscript.

References

1. Alfrink K, Keller I, Kortuem G, Doorn N (2023) Contestable AI by design: towards a framework. Minds Mach 33:613–639. https://doi.org/10.1007/s11023-022-09611-z
2. Alon-Barkat S, Busuioc M (2023) Human–AI interactions in public sector decision making: automation bias and selective adherence to algorithmic advice. J Public Adm Res Theory 33:153–169. https://doi.org/10.1093/jopart/muac007
3. Baker RS, Hawn A (2022) Algorithmic bias in education. Int J Artif Intell Educ 32:1052–1092. https://doi.org/10.1007/s40593-021-00285-9
4. Bayamlıoğlu E (2022) The right to contest automated decisions under the General Data Protection Regulation: beyond the so-called right to explanation. Regul Gov 16:1058–1078. https://doi.org/10.1111/rego.12391
5. Cavalcante Siebert L, Lupetti ML, Aizenberg E, Beckers N, Zgonnikov A, Veluwenkamp H et al (2023) Meaningful human control: actionable properties for AI system development. AI Ethics 3:241–255. https://doi.org/10.1007/s43681-022-00167-3
6. Custers B, Heijne AS (2022) The right of access in automated decision-making: the scope of article 15(1)(h) GDPR in theory and practice. Comput Law Secur Rev 46:105727. https://doi.org/10.1016/j.clsr.2022.105727
7. Fu Y, Weng Z (2024) Navigating the ethical terrain of AI in education: a systematic review on framing responsible human-centered AI practices. Comput Educ Artif Intell 7:100306. https://doi.org/10.1016/j.caeai.2024.100306
8. Grant DG, Behrends J, Basl J (2025) What we owe to decision-subjects: beyond transparency and explanation in automated decision-making. Philos Stud 182:55–85. https://doi.org/10.1007/s11098-023-02013-6
9. Green B (2022) The flaws of policies requiring human oversight of government algorithms. Comput Law Secur Rev 45:105681. https://doi.org/10.1016/j.clsr.2022.105681
10. Hamon R, Junklewitz H, Sanchez I, Malgieri G, De Hert P (2022) Bridging the gap between AI and explainability in the GDPR: towards trustworthiness-by-design in automated decision-making. IEEE Comput Intell Mag 17:72–85. https://doi.org/10.1109/mci.2021.3129960
11. Henin C, Le Métayer D (2022) Beyond explainability: justifiability and contestability of algorithmic decision systems. AI Soc 37:1397–1410. https://doi.org/10.1007/s00146-021-01251-8
12. Idowu JA, Koshiyama AS, Treleaven P (2024) Investigating algorithmic bias in student progress monitoring. Comput Educ Artif Intell 7:100267. https://doi.org/10.1016/j.caeai.2024.100267

13. Karimi AH, Barthe G, Schölkopf B, Valera I (2023) A Survey of Algorithmic Recourse: Contrastive Explanations and Consequential Recommendations. ACM Comput Surv 55:1–29. https://doi.org/10.1145/3527848
14. Langford M (2020) Taming the digital leviathan: automated decision-making and international human rights. AJIL Unbound 114:141–146. https://doi.org/10.1017/aju.2020.31
15. Laux J (2024) Institutionalised distrust and human oversight of artificial intelligence: towards a democratic design of AI governance under the European Union AI Act. https://doi.org/10.1007/s00146-023-01777-z. AI Soc
16. Laux J, Wachter S, Mittelstadt B (2024) Trustworthy artificial intelligence and the European Union AI Act: on the conflation of trustworthiness and acceptability of risk. Regul Gov 18:3–32. https://doi.org/10.1111/rego.12512
17. Memarian B, Doleck T (2023) Fairness, accountability, transparency, and ethics (FATE) in artificial intelligence (AI) and higher education: a systematic review. Comput Educ Artif Intell 5:100152. https://doi.org/10.1016/j.caeai.2023.100152
18. Nilholm C (2021) Research about inclusive education in 2020 — how can we improve our theories in order to change practice? Eur J Spec Needs Educ 36:358–370. https://doi.org/10.1080/08856257.2020.1754547
19. Ochoa Barrera MN, Ramírez Apolo DI, Estacio Dávila MF, Ureña Garces CB, Chica Tomalá MP (2025) Artificial intelligence to strengthen pedagogical support for students with learning disorders. Prisma J 1:384–396. https://doi.org/10.63803/prisma.v1n4.33
20. Ploug T, Holm S (2020) The four dimensions of contestable AI diagnostics: a patient-centric approach to explainable AI. Artif Intell Med 107:101901. https://doi.org/10.1016/j.artmed.2020.101901
21. Schiff D (2022) Education for AI, not AI for education: the role of education and ethics in national AI policy strategies. Int J Artif Intell Educ 32:527–563. https://doi.org/10.1007/s40593-021-00270-2
22. Smith H (2020) Algorithmic bias: should students pay the price? AI Soc 35:1077–1078. https://doi.org/10.1007/s00146-020-01054-3
23. Tsirantonaki S, Vlachou A (2025) AI-driven inclusive practices: innovative approaches to differentiated teaching and educational inclusion. Eur J Incl Educ 4:134–146. https://doi.org/10.7146/ejie.v4i1.156613
24. van Kolfschooten HB (2024) A health-conformant reading of the GDPR's right not to be subject to automated decision-making. Med Law Rev 32:373–391. https://doi.org/10.1093/medlaw/fwae029
25. Wiese LJ, Patil I, Schiff DS, Magana AJ (2025) AI ethics education: a systematic literature review. Comput Educ Artif Intell 8:100405. https://doi.org/10.1016/j.caeai.2025.100405

Legal instruments

1. Cited as primary texts not as journal articles

2. Regulation (EU) 2016/679 (General Data Protection Regulation). Article 22
3. Regulation (EU) 2024/1689 (Artificial Intelligence Act), Articles 6, 13, 14, 85 and 86, and Annex III
4. United Nations Convention on the Rights of Persons with Disabilities (2006) Articles 13 and 24

Figures

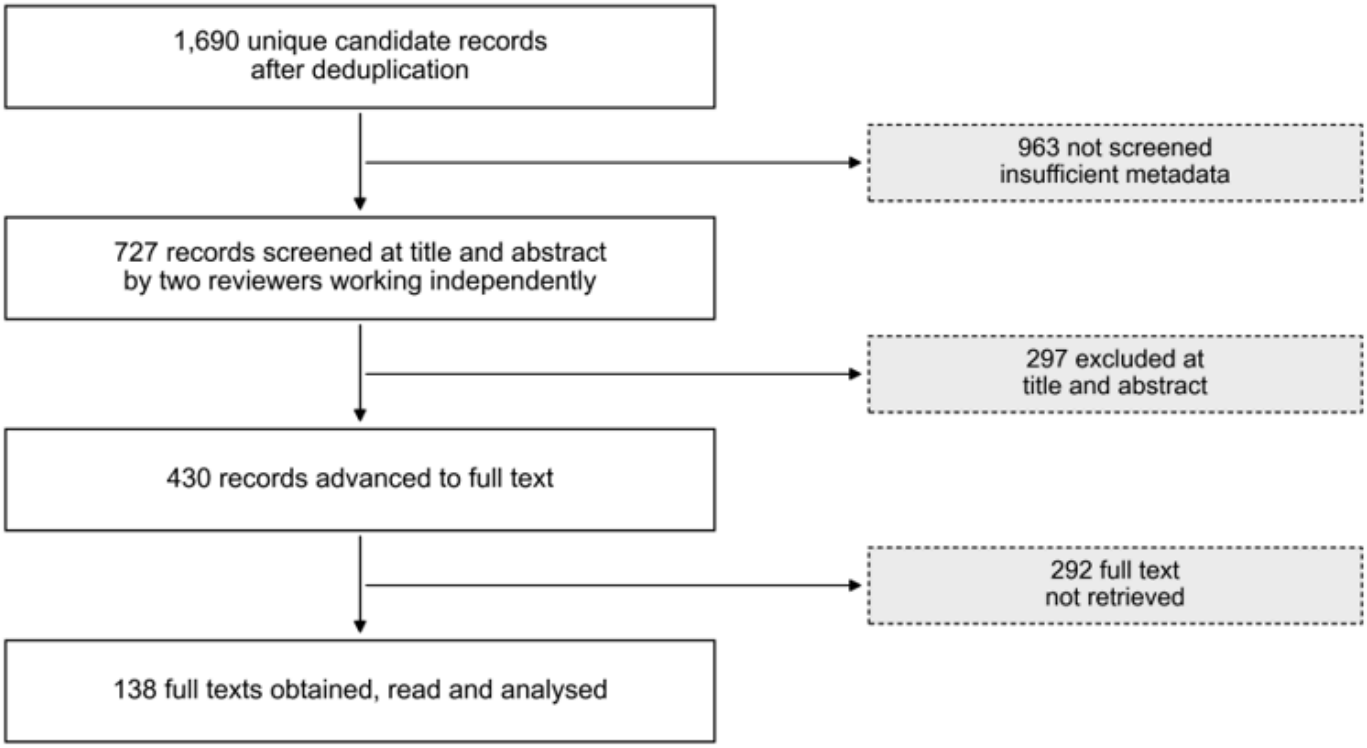


Figure 1

Flow of records from the deduplicated candidate pool to the analysed corpus. Solid boxes give the records carried forward at each stage and broken boxes those removed, with the reason for removal

The remaining 963 records were not screened because their metadata did not permit verification of all three concepts, most often because no abstract was available.

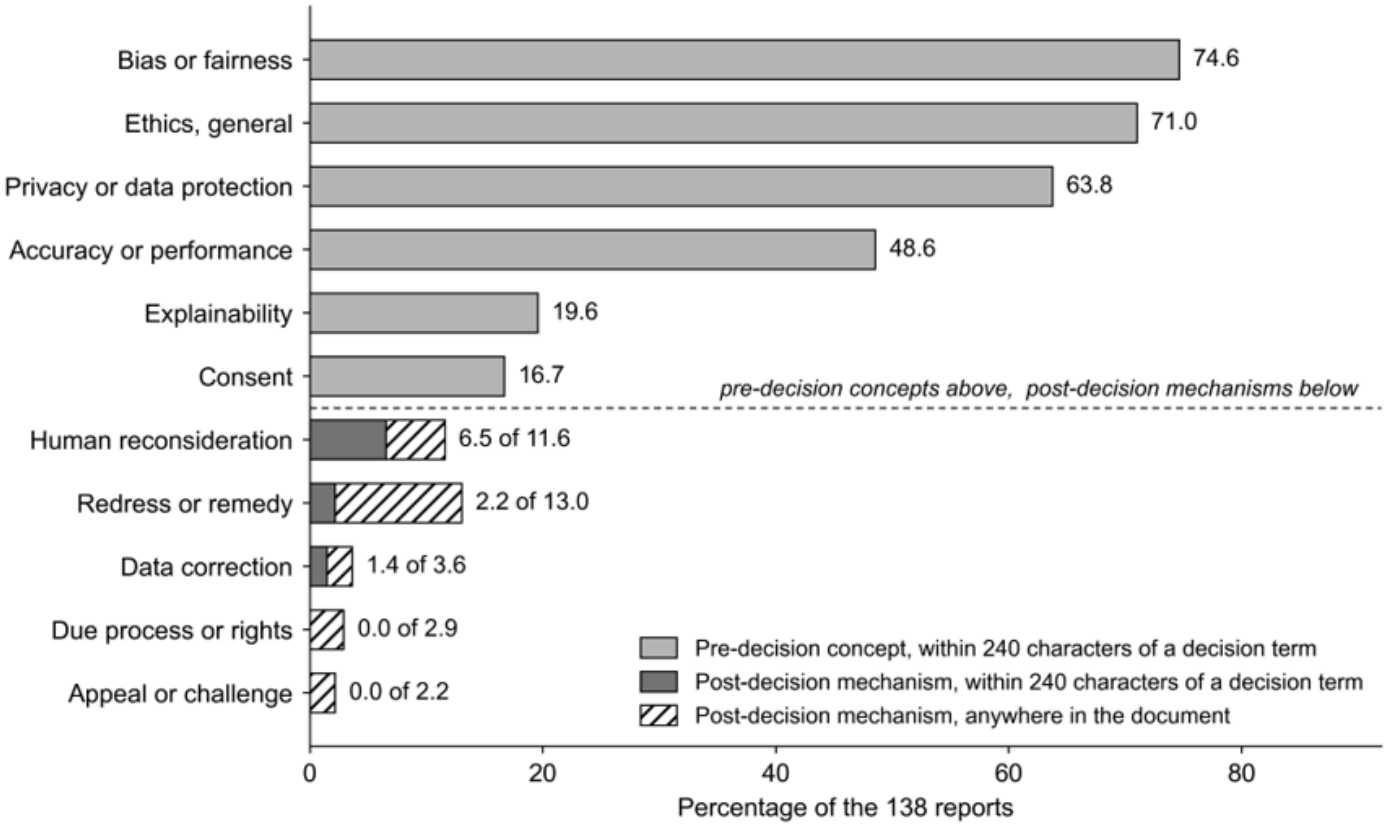


Figure 2

Presence of each concept across the 138 reports. Bars give the percentage of reports in which a term for the concept fell within 240 characters of a decision term, with the six pre-decision concepts above the broken line and the five post-decision mechanisms below it. For the post-decision mechanisms the hatched bar behind gives the percentage of reports mentioning the term anywhere in the document, so the distance between the two bars is the share of mentions detached from any decision. Labels give the value near a decision followed by the value anywhere

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- supplementary.zip
- Supplementaryinformation.docx