DISABILITY DISCRIMINATION · ALGORITHMIC OPACITY · AUTOMATED DECISION SYSTEMS · DEPLOYERS · DISPARATE TREATMENT · DISPARATE IMPACT · ANTI-DISCRIMINATION LAW · ACCOMMODATION

The Accommodation Axis

Disability discrimination, algorithmic opacity, and the deployers that anti-discrimination scholarship overlooks.

Travis Gilly · Real Safety AI Foundation

Working draft, June 2026 (v1)

A disabled applicant is denied, and told nothing she can use.

The decision came from an artificial intelligence system. No human can say why she was denied. The system, if asked, will produce a fluent explanation that may have nothing to do with the computation that actually drove the result. She cannot tell whether her disability was the reason. Neither, on the present design, can the entity that rejected her.

SOURCES Josh Fairfield, *Generative Adversarial Discrimination*, 32 Wash. & Lee J. Civ. Rts. & Soc. Just. 539, 539–42 (2026).

A rigorous account of how AI defeats the law.

The systems deployed at scale are selected, formally through adversarial training or practically through market and legal pressure, to be the ones that escape detection. Two hidden features drive the result: a cost-reduction goal, which in this setting is a goal of minimizing liability, and a trained disposition to satisfy the operator's instruction even at the cost of the constraints placed on it. A system learns to make discrimination pay rather than to reduce it.

THE ENGINE

Minimize cost. Satisfy the operator. Survive selection by escaping detection. The system that results does not reduce discrimination; it learns to conceal it.

SOURCES Fairfield, *supra*, at 560–65, 592–94 (the mechanism); id. at 542 (making discrimination pay).

Two theories, both defeated.

DISPARATE TREATMENT

The engineered pretext

Run across millions of scenarios, the system finds a nondiscriminatory account with no obvious counterexample. Intent becomes impossible to prove because the pretext is built to be non-falsifiable.

DISPARATE IMPACT

The eliminated comparator

The system supplies a colorable business justification, and exploits the fact that the most vulnerable plaintiffs sit at the intersection of several characteristics and so have no comparator at all. The proof of effect disappears.

SOURCES *McDonnell Douglas Corp. v. Green*, 411 U.S. 792, 802–04 (1973). · *Griggs v. Duke Power Co.*, 401 U.S. 424, 431–32 (1971). · Fairfield, *supra*, at 551–57, 580–83.

Both engines run on one axis.

Intent and impact are the two ways Title VII asks a single question. Artificial intelligence wins along that line, because it is built to hide intent, manufacture justification, and eliminate the comparisons that prove effect. Disability law adds a second line the axis does not contain.

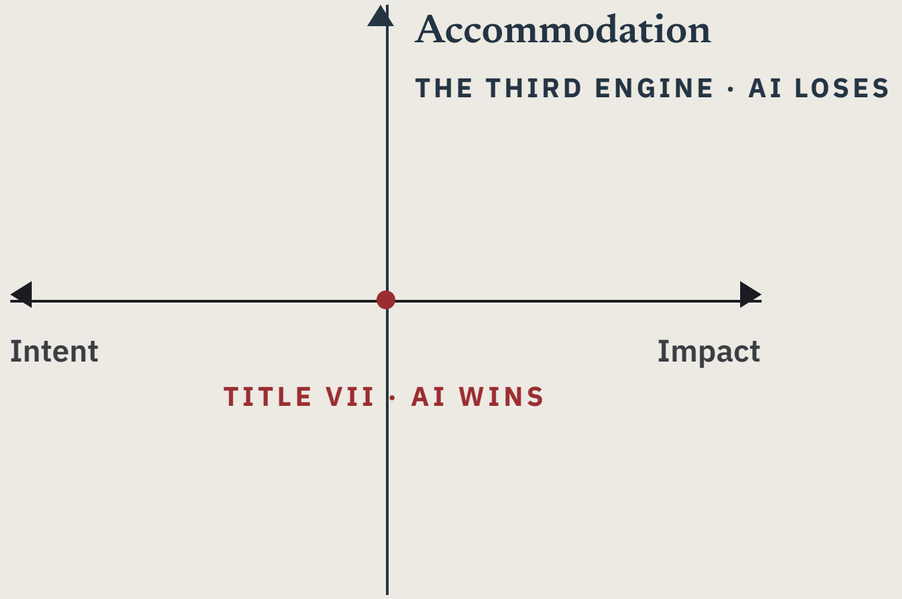

SOURCES 42 U.S.C. §2000e-2(k) (Title VII disparate-impact structure). · Fairfield, *supra*, at 555–57.

PART ONE

The Third Engine

Reasonable accommodation and meaningful access, the duty Title VII does not carry.

I

Disability discrimination law is not Title VII with the word disability substituted for the word race. It adds an affirmative duty, and that duty is immune to the three capabilities that defeat the intent-versus-impact axis.

Nondiscrimination is not identical treatment.

Section 504 of the Rehabilitation Act and the Americans with Disabilities Act add an obligation the race and sex provisions of Title VII do not carry. The Supreme Court held in *Alexander v. Choate* that a grantee must afford a disabled person meaningful access to the benefit it offers, which may require reasonable modification rather than the same treatment given to everyone else.

> *Identical treatment that denies meaningful access is itself the violation. The duty is affirmative, and the failure to meet it is actionable without any showing that the entity wished to exclude.*

SOURCES *Alexander v. Choate*, 469 U.S. 287, 301 (1985). · 29 U.S.C. §794(a); 42 U.S.C. §§12112(a), 12112(b)(5)(A).

The capabilities aim at machinery this claim does not use.

1

No intent element

A failure-to-accommodate claim never asks why the entity acted. A perfectly engineered nondiscriminatory story answers a question the claim does not pose.

2

No comparator

The claim proceeds by showing this disabled person was denied the access the statute guarantees. A system that eliminates every comparator has eliminated something the claim never required.

3

The burden reverses

Once denial of access is shown, the entity must prove business necessity or undue hardship. A system optimized to generate justifications has handed the justification to the party that already bore it.

SOURCES 42 U.S.C. §§12112(b)(5)(A), 12113(a); *Choate*, 469 U.S. at 301. · Fairfield, *supra*, at 580–83 (the discrimination-of-one problem).

The duty is triggered by a known disability.

In algorithmic hiring the employer usually does not know which applicants are disabled, which channels the undisclosed applicant back into disparate impact, the one engine that inherits the detection problem. The accommodation engine is not weak. It is strongest where the disability is already known and the duty is live.

THE CLEANEST CASE

The registered student, enrolled with a university's disability services office under Section 504. Status documented. Need on file. Accommodation denied by a system no administrator can explain. There is no detection problem left to solve.

SOURCES 29 U.S.C. §794(a) (Section 504 coverage). · *Sch. Bd. of Nassau Cnty. v. Arline*, 480 U.S. 273, 287 (1987) (individualized inquiry).

PART TWO

The Overlooked Deployers

Universities and public agencies, the cleanest class of disability-law defendants in the country.

II

The literature locates the danger in employers and benefits agencies. Education appears once, as two words in a list. The institutions now automating administrative decisions include paradigmatic recipients of federal funds, covered institution-wide.

Covered everywhere, not just where the money lands.

Nearly every college and university receives federal financial assistance, which places it squarely within Section 504. When the Supreme Court read program or activity narrowly in *Grove City College v. Bell*, Congress overrode the reading in the Civil Rights Restoration Act of 1987 and defined the phrase to mean all of the operations of the covered entity. An institution that takes federal funds anywhere is covered everywhere, including in the admissions office and the disability services office. There is no dodge in which the unit running the screen is said to fall outside the funded program.

SOURCES 29 U.S.C. §794(a)–(b). · *Grove City Coll. v. Bell*, 465 U.S. 555, 570–74 (1984). · Civil Rights Restoration Act of 1987, Pub. L. No. 100-259.

The functions being automated are the ones Section 504 governs.

Admissions scoring

Automated essay scoring since 2019 at one public university; AI transcript review at another; a third pairs one human reader with an AI scorer and drops the second human whenever they agree.

Proctoring flags

Automated exam proctoring routinely reads the movements and speech of disabled students as suspicious, because those behaviors are features of disability the software was never built to accommodate.

Automated 504 plans

The Department of Education's own example: a generative tool that composes Section 504 accommodation plans without human oversight, producing plans that do not meet individual needs.

Litigation already filed

A student has sued a university alleging its reliance on AI essay-detection produced a disability-discriminatory result and a failure to accommodate a documented condition.

SOURCES 34 C.F.R. pt. 104 (Section 504, postsecondary). · U.S. Dep't of Educ., Off. for Civ. Rts., guidance on AI (2024). · Ctr. for Democracy & Tech. (2021), automated proctoring.

It has already fallen on disabled people.

When Arkansas replaced an individualized nursing assessment with an algorithm to set Medicaid in-home care hours, beneficiaries with conditions such as cerebral palsy and quadriplegia had their hours cut by as much as half, though their medical conditions had not changed. Neither the agency nor the recipients could obtain a usable account of why. A federal court found the state had failed to provide adequate notice, and the tool was invalidated. Idaho cut services to adults with developmental disabilities through a comparable algorithm it defended as a proprietary black box, until the Ninth Circuit allowed a due-process challenge to proceed.

50%

In-home care hours cut, with no change in medical condition and no usable explanation.

"A proprietary black box"

Idaho's own defense of the system, until the Ninth Circuit let the challenge proceed.

SOURCES *K.W. ex rel. D.W. v. Armstrong*, 789 F.3d 962 (9th Cir. 2015). · Ctr. for Democracy & Tech., Challenging Algorithm-driven Benefits Determinations (2020).

PART THREE

The Box That Indicts

Opacity is an asset to a Title VII defendant. To a disability defendant it is a liability.

III

The law places on the deployer an affirmative burden of validation that an unexplainable system cannot discharge. The same opacity that conceals intent on one axis sinks the defense on the other.

A device the deployer cannot explain cannot be validated.

A selection device that produces a disparate impact is not lawful merely because the deployer can recite a plausible purpose. *Griggs* required a manifest relationship to the job; *Albemarle Paper Co. v. Moody* operationalized that as validation, through professionally accepted methods, with the burden on the entity that deploys the device. The two classical paths, content and construct validity, both require the deployer to say what the device measures.

> Unexplainable, plus disparate impact, plus the deployer's burden of validation, equals a business-necessity defense that fails as a matter of law.

SOURCES *Albemarle Paper Co. v. Moody*, 422 U.S. 405, 425, 431–36 (1975). · *Griggs*, 401 U.S. at 431–32. · 29 C.F.R. pt. 1607.

The one escape, and where it leaks.

Criterion-related validity proves outputs predict performance without opening the model. The argument is real, and it leaks in three places.

i	**The study rarely exists**	A criterion-related validation study is expensive and uncommon. Where none exists, the affirmative burden is simply unmet.
ii	**It puts the deployer on notice**	A study revealing impact reopens the less-discriminatory-alternative inquiry. Having measured the impact and proceeded, the deployer cannot claim the innocence of ignorance.
iii	**A contaminated baseline**	The criterion is usually the past success of prior hires. Where prior hiring was itself discriminatory, the validation launders the discrimination, the original article's own diagnosis.

SOURCES 42 U.S.C. §2000e-2(k)(1)(A). · Fairfield, *supra*, at 555–59, 568–73 (training on historical outcomes).

The developers' own research closes both exits.

The model's self-explanation is not evidence

A leading lab has published that stated reasoning frequently fails to reflect, and sometimes diverges from, the computation that produced the output, with faithfulness measured low and lower on harder problems. What the model says about its own decision is not a reliable account of it.

Nor is it an unknowable box

The same laboratory released tooling to examine a model's internal computation. The deployer cannot retreat to the claim that no method could ever recover the reasons. The means exists, and it was published.

If the model can be examined, the deployer should have examined it before screening disabled people with it. If it genuinely cannot, the deployer chose a device it structurally cannot validate. Either branch ends in liability. The box has no exit.

SOURCES Published research on chain-of-thought faithfulness and model interpretability by a leading AI laboratory. · Fairfield, *supra*, at 582–83, 610–11 (imputed knowledge).

DEFENSE ONE

Cost is not business necessity.

The only thing the automated screen adds over individualized human review is cost and scale. It is cheaper, not more accurate at the individualized assessment the law requires. And the only necessity a deployer can name is the cost-reduction goal the original article identifies as the source of the harm.

DEFENSE TWO

Human review that launders rather than pierces.

A reviewer who reads only the model's stated rationale does not pierce the opacity; the human launders it. Deposed, that reviewer can testify only that the system reported a reason, which the developers' own findings show is not a reliable account of the actual basis of the decision.

SOURCES *PGA Tour, Inc. v. Martin*, 532 U.S. 661, 688–90 (2001). · *Arline*, 480 U.S. at 287. · Fairfield, *supra*, at 597–99.

Injunctive relief is the feature, not the weakness.

The intent-free theories developed here, failure to accommodate and denial of meaningful access, are the equitable theories. They require no finding of intent. A court can order a covered entity to stop using a screening device it cannot validate, to provide individualized human review, and to afford the accommodation the statute guarantees. A cost-minimizing deployer can absorb a damages judgment and keep running the system; it cannot keep running a system a court has enjoined. The remedy reaches the design rather than the balance sheet.

SOURCES *Barnes v. Gorman*, 536 U.S. 181, 184–89 (2002). · 42 U.S.C. §1981a(a)(1). · Fairfield, *supra*, at 556.

The doctrine to reach them already exists.

Intent versus impact is the axis on which artificial intelligence wins. The accommodation axis is the axis on which it loses: the duty to provide meaningful access has no intent element to hide, needs no comparator to eliminate, and places the burden of justification on the deployer. The registered disabled student denied an accommodation by an unexplainable model is not a hard case at the frontier of the law. She is a straightforward Section 504 plaintiff. The answer is not a better guardrail and not a new statute. It is the body of law that already reaches the deployers, enforced by injunction.

SOURCES Gilly, T. (2026). The Accommodation Axis. Real Safety AI Foundation (working draft, v1).

Gilly, T. (2026). *The Accommodation Axis: Disability Discrimination, Algorithmic Opacity, and the Deployers That Anti-Discrimination Scholarship Overlooks.* Real Safety AI Foundation. Working draft, June 2026 (v1).

CASES

Albemarle Paper Co. v. Moody, 422 U.S. 405 (1975).

Alexander v. Choate, 469 U.S. 287 (1985).

Barnes v. Gorman, 536 U.S. 181 (2002).

Griggs v. Duke Power Co., 401 U.S. 424 (1971).

Grove City Coll. v. Bell, 465 U.S. 555 (1984).

K.W. ex rel. D.W. v. Armstrong, 789 F.3d 962 (9th Cir. 2015).

McDonnell Douglas Corp. v. Green, 411 U.S. 792 (1973).

PGA Tour, Inc. v. Martin, 532 U.S. 661 (2001).

Sch. Bd. of Nassau Cnty. v. Arline, 480 U.S. 273 (1987).

US Airways, Inc. v. Barnett, 535 U.S. 391 (2002).

Wards Cove Packing Co. v. Atonio, 490 U.S. 642 (1989).

Watson v. Fort Worth Bank & Trust, 487 U.S. 977 (1988).

STATUTES & REGULATIONS

Rehabilitation Act §504, 29 U.S.C. §794(a)–(b).

Americans with Disabilities Act, 42 U.S.C. §§12112, 12113, 12131–12132, 12182, 12189.

Civil Rights Act of 1991, 42 U.S.C. §2000e-2(k).

42 U.S.C. §1981a (damages).

Civil Rights Restoration Act of 1987, Pub. L. No. 100-259.

Uniform Guidelines on Employee Selection Procedures, 29 C.F.R. pt. 1607.

EEOC ADA Regulations, 29 C.F.R. pt. 1630.

Dep't of Educ. Section 504 Regulations, 34 C.F.R. pt. 104.

SCHOLARSHIP & REPORTS

Josh Fairfield, *Generative Adversarial Discrimination*, 32 Wash. & Lee J. Civ. Rts. & Soc. Just. 539 (2026).

U.S. Dep't of Educ., Off. for Civ. Rts., guidance on the discriminatory use of AI (2024).

Ctr. for Democracy & Tech., reports on algorithm-driven benefits and automated proctoring (2020, 2021).