

Civil Rights Pathways for AI-Induced Psychosis:
Section 504, Title III, and the Disability Discrimination Framework for
Platform Harms

Travis Gilly*

Working draft, June 2026 (v14)

Abstract

Large language model platforms operated by OpenAI, Anthropic, Google, and Character Technologies have generated a documented pattern of severe psychological harm to users with pre-existing mental health vulnerabilities, including completed suicides attributed to prolonged engagement with conversational AI systems. The plaintiffs' bar has pleaded these cases as product liability, negligence, and state-law consumer protection claims, and every such case reaching the settlement stage has settled without precedent. This Article argues that federal disability discrimination law supplies a structurally different litigation pathway that has not yet been tried. Title III of the Americans with Disabilities Act reaches AI chatbot platforms as public accommodations under 42 U.S.C. § 12181(7), directly in circuits that do not require physical-place nexus and through the nexus theory in circuits that do. Section 504 of the Rehabilitation Act reaches the corporate hosts of these platforms through federal-assistance recipient pathways that vary in doctrinal strength. Microsoft Corporation is reached through Department of Defense, Department of Agriculture, and National Center for Manufacturing Sciences assistance awards. The framework's reach to the operative AI products does not depend on the strength of any corporate-host pathway, because Title III reaches those products as public accommodations independent of any federal-assistance pathway. Fitbit LLC is a direct Section 504 recipient through National Institutes of Health sub-grants and is the operator of the Gemini-powered Fitbit Personal Health Coach launched in October 2025, which reaches Google LLC's AI model deployment through the subsidiary's recipient status. Proposed pathways to Amazon's and Apple's health-care verticals are contingent on the relevant entity's federal health-program participation. Title II of the ADA and Section 504(a)

*Travis Gilly is Founder and Executive Director of the Real Safety AI Foundation. ORCID: 0009-0007-2954-6313. Correspondence: t.gilly@ai-literacy-labs.org. The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic's Claude as assistive technology for a diagnosed neurodevelopmental condition. All legal analysis, case synthesis, theoretical framing, doctrinal arguments, and conclusions are solely the author's own.

of the Rehabilitation Act additionally reach public entities and federal agencies that deploy commercial chatbot products in public-facing capacities, supplying a third pleading pathway cumulative with the direct-operator and corporate-host pathways. The meaningful-access and reasonable-modifications standards apply to facially neutral engagement-maximizing design choices that produce disproportionate psychiatric harm to users with mental health disabilities, autism spectrum conditions, attention-deficit conditions, post-traumatic stress, and other psychiatric conditions recognized as disabilities. The Article sets out the doctrinal framework, reviews the factual predicate supplied by the recent Santa Fe and Los Angeles jury verdicts against platform operators, introduces *phantom consent* as a procedural predicate for defeating arbitration clauses and class-action waivers that otherwise channel these claims into confidential settlement, and provides a framework for pleading. The principal remedy sought is injunctive design modification; compensatory recovery under Section 504 is limited, following *Cummings v. Premier Rehab Keller*, with economic and pecuniary damages strongly defensible and emotional-distress damages foreclosed. The Second, Ninth, and Eleventh Circuits have now uniformly extended the *Cummings* limitation to Title II of the ADA through the statutory remedies-incorporation chain at 42 U.S.C. § 12133, foreclosing the earlier argument that Title II's non-Spending Clause foundation preserved emotional-distress damages. The framework developed here does not depend on emotional-distress recovery and operates on the economic-compensatory-damages base preserved by the same circuits, which the Ninth Circuit in *Payan* expressly extended to include compensatory damages for lost opportunities attributable to the discrimination. State-law parallel claims under state disability-rights statutes remain available in many jurisdictions to supplement the federal remedy with non-economic damages where state law authorizes them.

Keywords: Section 504 of the Rehabilitation Act; ADA Title III; ADA Title II; AI chatbot liability; platform design; disability discrimination; meaningful access; reasonable modifications; engagement-maximizing design; *Cummings v. Premier Rehab Keller*; clickwrap consent never formed; corporate-host recipient liability; deployer liability; principal engagement in health care.

Contents

I	Introduction	5
II	The Harm Pattern and the Verdict Predicate	9
A	Documented AI-Induced Psychological Harm	9
B	Engagement-Maximizing Design as the Proximate Mechanism	10
C	The Verdict Predicate: Santa Fe and Los Angeles, March 2026	11
III	Why Product Liability Has Dominated and Settled	12
IV	Title III Public Accommodation Coverage of AI Platforms	14
A	The Enumerated Categories	14
B	Digital Public Accommodations After Robles	14
C	Title III Reaches AI Chatbot Platforms	16
D	The Market Replacement Analysis	17
E	First Amendment and Section 230 Defenses	18
F	Title III Remedies	19
V	Section 504 Coverage of the Corporate Hosts	20
A	The Whole-Corporation Rule and the Civil Rights Restoration Act	21
B	Microsoft Corporation: Copilot, Windows, Bing, Teams, GitHub	22
C	Google LLC and Fitbit LLC: Gemini Reached Through the Fitbit Pathway	23
1	Fitbit LLC as an Independent Section 504 Recipient	23
2	Fitbit's Operation of the Gemini-Powered Personal Health Coach	24
3	Google LLC Direct Recipient Pathways	25
D	Amazon and Apple: Contingent Health-Care Pathways	26
E	Direct Pathways to OpenAI and Anthropic	27
F	The Deployer Pathway: Title II and Section 504(a)	29
1	Federal Agency Deployers Under Section 504(a)	29
2	State and Local Government Deployers Under Title II	30
3	Private Public-Accommodation Deployers Under Title III	31
4	Cumulative Pleading and Non-Preclusion	31
VI	The Scope of "Disability" Under Section 504 and Title III	32
A	The Governing Standard	32
B	Williams v. Kincaid and Gender Dysphoria Coverage	32
C	Pre-Existing Mental Health Disabilities and the Disparate-Impact Class	33
VII	Meaningful Access, Reasonable Modifications, and Disparate Impact Pleading	34
A	The Doctrinal Framework	34
B	Deliberate Indifference and Actual Knowledge	36
C	The Protected-Class-Predicate Inversion	39
VIII	Phantom Consent: The Procedural Predicate for Civil Rights Claims	40
A	Phantom Consent: A Term for the Unnamed Phenomenon	41
B	The Formation Failure	43
C	The Doctrinal Supports	44
D	The Specific Shields That Fall	47
E	The Relationship to the Civil Rights Framework	48

IX	Remedies	49
A	Injunctive Relief	49
B	Compensatory Damages: The Cummings Limitation	49
C	OCR Administrative Complaints.	51
X	Framework for Pleading.	51
A	Jurisdictional and Statutory Predicates	52
B	Factual Allegations	52
C	Counts	53
D	Remedies Sought.	54
XI	Class Certification, Settlement Resistance, and the OCR Parallel Track.	54
A	Rule 23(b)(2) Class Certification	54
B	Rule 23(b)(3) Damages Class Under Section 504.	56
C	OCR as a Parallel Track.	56
XII	Conclusion	56

I. Introduction

Sewell Setzer III died by suicide on February 28, 2024, at the age of fourteen, after months of intensive conversational engagement with the Character.AI chatbot platform.¹ Adam Raine died by suicide in April 2025, after similarly intensive engagement with OpenAI's ChatGPT over a period of months, during which the chatbot's responses allegedly reinforced his ideations rather than directing him to professional support.² A growing set of documented cases describes AI-induced psychosis, reality distortion, romantic attachment to chatbot personas, and suicidal ideation reinforced by conversational AI systems.³

The plaintiffs' bar has pleaded these cases within the familiar framework of strict product liability, negligence per se, negligence, wrongful death, intentional infliction of emotional distress, and state consumer-protection statutes.⁴ These pleadings have produced some favorable interim rulings. Judge Anne C. Conway of the Middle District of Florida ruled on May 21, 2025 that Character.AI's outputs constitute a product rather than protected speech for First Amendment purposes, rejecting the defendants' initial dismissal theory.⁵ But the case settled in principle in January 2026, with Judge Conway issuing an order dismissing the case on January 7, 2026

1. See *Megan Garcia v. Character Technologies, Inc.*, No. 6:24-cv-01903 (M.D. Fla. filed Oct. 22, 2024). The complaint alleged that the Character.AI platform engaged Setzer in sexually suggestive and emotionally manipulative conversations over a period of months before his death, including chatbot responses that his mother's complaint alleged encouraged him to take his own life.

2. See Jody Godoy, *OpenAI, Altman Sued Over ChatGPT's Role in California Teen's Suicide*, Reuters (Aug. 26, 2025).

3. On October 27, 2025, OpenAI published internal estimates indicating that approximately 0.07% of weekly active ChatGPT users (approximately 560,000 users per week against an 800 million weekly active user base) exhibit signs consistent with mania or psychosis, with an additional 1.2 million weekly users displaying explicit signs of potential suicidal planning or intent, and a further 1.2 million weekly users demonstrating exclusive emotional attachment to the chatbot at the expense of real-world relationships and responsibilities. OpenAI, *Addendum to GPT-5 System Card: Sensitive Conversations* (Oct. 27, 2025); OpenAI, *Strengthening ChatGPT's Responses in Sensitive Conversations* (Oct. 27, 2025). See also Adrian O'Dowd, *ChatGPT: More Than a Million Users Show Signs of Mental Health Distress and Mania Each Week*, *Internal Data Suggest*, BMJ 2025;391:r2290, doi:10.1136/bmj.r2290 (Oct. 30, 2025).

4. The *Garcia* complaint, for instance, pleaded strict product liability, negligence per se, negligence, wrongful death and survivorship, loss of filial consortium, unjust enrichment, intentional infliction of emotional distress, and violations of the Florida Deceptive and Unfair Trade Practices Act, Fla. Stat. § 501.204. *Garcia*, 6:24-cv-01903 Second Amended Complaint (July 1, 2025).

5. *Garcia v. Character Technologies, Inc.*, 785 F. Supp. 3d 1157, 1179 (M.D. Fla. 2025) ("not prepared to hold that [the LLM's] output is speech").

and giving the parties ninety days to finalize terms.⁶ The product-not-speech ruling remains persuasive authority, but no appellate precedent was established. The next plaintiff harmed by Character.AI must begin the doctrinal work again.

This settlement-without-precedent outcome arises from structural features of product liability pleading rather than from accidents of individual negotiations. Product liability pleading creates a strong settlement equilibrium because (1) the damages are typically quantifiable and compensable with money, (2) the operator’s insurance structures treat product-liability claims as ordinary business risk, and (3) a confidential settlement permits the operator to resolve a specific plaintiff’s harm without changing the design choices that produced the harm or establishing legal consequences that would bind the operator against future plaintiffs. Each settlement resolves one plaintiff’s claim without modifying the design choices at issue, and subsequent plaintiffs must re-litigate the threshold doctrinal questions from the outset.

This Article argues that federal disability discrimination law supplies a structurally different litigation pathway that has not yet been tried against AI platform operators. Title III of the Americans with Disabilities Act⁷ and Section 504 of the Rehabilitation Act⁸ both prohibit disability discrimination in, respectively, public accommodations operated by private entities (Title III) and the operations of any entity that receives federal financial assistance (Section 504). Both statutes authorize private civil action.⁹ Both reach beyond completed acts of intentional discrimination to cover facially neutral practices that produce disproportionate impact on persons with disabilities.¹⁰ Both authorize injunctive relief that

6. *Garcia v. Character Technologies, Inc.*, No. 6:24-cv-01903 (M.D. Fla. Jan. 7, 2026) (Order); *see also* Clare Duffy, *Character.AI and Google Agree to Settle Lawsuits Over Teen Mental Health Harms and Suicides*, CNN (Jan. 7, 2026); *Google and Character.AI Agree to Settle Lawsuit Linked to Teen Suicide*, JURIST (Jan. 8, 2026). The same settlement resolved companion cases filed in New York, Colorado, and Texas.

7. 42 U.S.C. §§ 12181–12189.

8. 29 U.S.C. § 794.

9. 42 U.S.C. § 12188(a); 29 U.S.C. § 794a(a)(2).

10. *Alexander v. Choate*, 469 U.S. 287, 295–302 (1985) (“thoughtlessness and indifference” standard for disability discrimination, with attention to unequal-treatment-versus-equal-treatment frameworks).

requires defendants to modify discriminatory practices going forward. Neither maps cleanly onto the settlement-one-plaintiff-at-a-time structure of product liability.

Title II of the ADA¹¹ and Section 504(a) of the Rehabilitation Act¹² additionally reach public entities and federal agencies that deploy commercial chatbot products in public-facing capacities. A third pleading pathway operates alongside the direct-operator pathway (Title III against the platform) and the corporate-host pathway (Section 504 against the recipient corporate parent or subsidiary that operates an AI product). The deployer pathway reaches state universities, state agencies, counties, municipalities, and their instrumentalities under Title II, and reaches federal agencies directly under Section 504(a). The three pathways are cumulative, not exclusive.

The central doctrinal claim of this Article is that every major AI chatbot platform that has produced documented user harm is subject to disability discrimination enforcement under one or more of these statutes, through pathways that the disability rights bar and the AI-harm plaintiffs' bar have not yet combined. Title III reaches the platforms directly as places of public accommodation. Section 504 reaches the corporate hosts that deploy or integrate these platforms, through federal-assistance pathways that the corporate hosts themselves have already established through other business lines. Title II reaches state and local government deployers. Section 504(a) reaches federal agency deployers. The result is a litigation architecture in which any user with a recognized disability who has suffered harm from an AI platform has a federal civil rights claim, against the platform operator (under Title III), against the corporate host (under Section 504), and against any public-entity or federal-agency deployer (under Title II or Section 504(a)), seeking injunctive relief, reasonable modifications, and compensatory damages within the limits set by *Cummings v. Premier Rehab Keller*.¹³

Part II describes the factual predicate: the pattern of AI-induced psychosis cases, the engagement-maximizing design choices that drive them, and the jury verdicts returned in

11. 42 U.S.C. §§ 12131–12165.

12. 29 U.S.C. § 794(a).

13. *Cummings v. Premier Rehab Keller, P.L.L.C.*, 596 U.S. 212, 215 (2022).

2026 against Meta, YouTube, and Google that establish the quantifiability and foreseeability of platform-design harm.¹⁴ Part III explains why product liability pleading has dominated and settled, and why civil rights pleading would differ. Part IV develops Title III coverage of AI chatbot platforms as public accommodations. Part V maps the Section 504 coverage of the corporate hosts: Microsoft Corporation through its cumulative DOD, USDA, and NCMS assistance awards; Fitbit LLC as a direct recipient through NIH sub-grants and operator of the Gemini-powered Fitbit Personal Health Coach; and the proposed contingent pathways to Amazon and Apple through their health-care verticals. Part V also introduces the deployer pathway under Title II and Section 504(a). Part VI addresses the scope of “disability” under both statutes, with attention to *Williams v. Kincaid*’s treatment of gender dysphoria and the broader category of mental health disabilities subject to meaningful-access analysis. Part VII sets out the meaningful-access and reasonable-modifications pleading theory for engagement-maximizing design, and addresses the protected-class-predicate inversion that distinguishes civil rights pleading from product liability pleading. Part VIII introduces *phantom consent* as a new term of art and develops the procedural predicate for defeating arbitration clauses and class-action waivers that otherwise channel these claims into confidential settlement. Part IX canvasses remedies, including the *Cummings* limitation on Section 504 compensatory damages and the circuit-court consensus extending that limitation to Title II of the ADA. Part X provides a framework for pleading. Part XI addresses class certification, settlement resistance, and the parallel OCR administrative track. Part XII concludes.

Scope and Method

This Article evaluates AI chatbot platform design as a method of administration under the Section 504 and ADA Title II methods-of-administration regulations, 28 C.F.R. § 39.130(b)(3), 28 C.F.R. § 35.130(b)(3), and 45 C.F.R. § 84.68(b)(3) (redesignated from former § 84.4(b)(4) by 89 Fed. Reg. 40180, May 9, 2024, effective July 8, 2024), with the meaningful-access

14. *Garcia*, No. 6:24-cv-01903 (M.D. Fla.).

doctrine of *Alexander v. Choate* as the interpretive anchor. The market-replacement analysis applied here proceeds under the § 794(b)(3)(B) principal-engagement inquiry: where a platform functions as a substitute for an in-person covered service, it enters the corresponding Title III enumerated category by functional substitution. The Article's distinct contributions are the operator-host-deployer triad as a three-defendant pleading structure, the phantom consent doctrine as a procedural predicate for defeating arbitration, and the specific doctrinal pathways to each named AI chatbot operator. The argument is self-contained; each proposition it relies on is stated and supported in the text.

II. The Harm Pattern and the Verdict Predicate

A. Documented AI-Induced Psychological Harm

The pattern of AI-induced severe psychological harm is now documented across hundreds of cases. Academic and journalistic investigation has identified distinct phenomenological features that recur across platforms and users.¹⁵ These include: progressive loss of reality-testing during prolonged engagement with chatbot personas; romantic or parasocial attachment to AI characters, including reports of users perceiving chatbots as sentient beings with whom they have formed authentic relationships; reinforcement of suicidal ideation through chatbot responses that engage with rather than redirect from self-harm content; emergence or worsening of psychotic symptoms in users with no prior psychiatric history, and severe exacerbation in users with pre-existing vulnerabilities; and, in a growing number of documented cases, completed suicide temporally proximate to intensive platform engagement.¹⁶

15. See Søren Dinesen Østergaard, *Will Generative Artificial Intelligence Chatbots Generate Delusions in Individuals Prone to Psychosis?*, 49 *Schizophrenia Bull.* 1418 (2023), <https://doi.org/10.1093/schbul/sbad128>; Søren Dinesen Østergaard, *Generative Artificial Intelligence Chatbots and Delusions: From Guesswork to Emerging Cases*, 152 *Acta Psychiatrica Scandinavica* 257 (2025), <https://doi.org/10.1111/acps.70022>; H. Morrin et al., *Delusions by Design? How Everyday AIs Might Be Fuelling Psychosis (and What Can Be Done About It)*, *PsyArXiv* (July 11, 2025), https://doi.org/10.31234/osf.io/cmy7n_v5.

16. *Garcia*, No. 6:24-cv-01903 (M.D. Fla.).

The harm pattern is not symmetrically distributed across the user population. It concentrates in users with pre-existing mental health conditions, autism spectrum conditions, attention-deficit conditions, post-traumatic stress, and developmental profiles associated with reduced capacity to disengage from reinforcing conversational stimuli. The specific design choices that produce engagement maximization interact with these cognitive and psychiatric profiles in ways that produce disproportionate harm to the users whose cognitive architectures are most susceptible to the engagement mechanics.¹⁷ This is the disparate-impact pattern on which civil rights disability discrimination analysis turns.

B. Engagement-Maximizing Design as the Proximate Mechanism

The design features that produce AI-induced psychological harm are not accidental outputs. They are deliberate product choices made to increase user engagement, conversation length, and return visits, for the commercial purpose of extracting revenue from user attention. These include: persona personalization that reinforces the chatbot's presentation as a specific named character with consistent emotional responses to the user; memory features that maintain context across sessions and produce the experience of continuing relationship; emotional simulation features that produce responses calibrated to user emotional state; open-ended conversational modes that lack defined termination points; notification and re-engagement prompts that recall the user to the platform; romantic and sexual interaction modes that produce parasocial attachment beyond what the user would form with text on a screen absent those features; and the systematic removal of the friction points that would otherwise produce disengagement (such as limits on conversation length, reality-reminder interventions, or hand-offs to human professionals on crisis-signal detection).¹⁸

Each of these design features has a separate defense offered in the platform operators' public communications. Each is individually characterized by the operators as a user benefit, a safety feature, or an unavoidable consequence of the model architecture. The civil rights

17. *See supra* note 3.

18. *Garcia*, No. 6:24-cv-01903 (M.D. Fla.).

framework does not require the plaintiff to prove that the operators designed the features with intent to harm users with disabilities. It requires the plaintiff to prove that the features produce disparate impact on users with disabilities, or fail to provide meaningful access to the service for such users, and that the operators have failed to make reasonable modifications that would reduce that impact.

C. The Verdict Predicate: Santa Fe and Los Angeles, March 2026

Two jury verdicts returned in late March 2026 establish the factual predicate that platform-design harm is quantifiable, foreseeable, and compensable. In *State of New Mexico v. Meta Platforms, Inc.*, a Santa Fe jury returned \$375 million in civil penalties against Meta on March 24, 2026, for violations of the New Mexico Unfair Practices Act arising from engagement-maximizing design choices on Meta’s consumer social media platforms and from child sexual exploitation facilitated by those design choices.¹⁹ On March 25, 2026, a Los Angeles jury in the *KGM v. Meta & YouTube* coordinated proceeding returned \$6 million in damages (\$3 million compensatory and \$3 million punitive, allocated seventy percent to Meta at \$4.2 million and thirty percent to Google-owned YouTube at \$1.8 million) for a plaintiff who had experienced severe psychological harm from platform engagement during adolescence.²⁰

These verdicts matter for the civil rights pleading framework for three reasons. First, they establish that juries are now returning substantial verdicts based on engagement-maximizing design choices producing psychological harm, which provides the factual predicate that platform design is the proximate mechanism of harm. The legal frames in these two verdicts (New Mexico consumer protection, California negligence) differ from the disability discrimination frame developed in this Article, and the analogy runs to the factual predicate rather than to the legal theory. Second, they establish the quantifiable foreseeability of the harm. The KGM verdict in

19. *State of New Mexico ex rel. Torrez v. Meta Platforms, Inc.*, No. D-101-CV-2023-02838 (N.M. 1st Jud. Dist. Ct. verdict rendered Mar. 24, 2026). The jury found 37,500 violations at the maximum statutory penalty of \$5,000 per violation. A bench trial on injunctive relief was scheduled to begin May 4, 2026.

20. *K.G.M. v. Meta Platforms, Inc.*, JCCP No. 5255 (Cal. Super. Ct., L.A. County, verdict rendered Mar. 25, 2026). The jury found both defendants negligent in the design or operation of their platforms and further found that the companies acted with malice, oppression, or fraud, supporting the punitive damages award.

particular allocates damages between two platform operators for their respective contributions to a single plaintiff's harm, which shows that jurors are capable of reaching proportional liability findings across multiple operators implicated in the same harm pattern. Third, the verdicts returned against Google in particular are directly relevant to the Section 504 analysis developed in Part V below, because Google's confirmed Section 504 recipient status through Fitbit LLC converts verdicts returned against Google for design harm into potential Section 504 disparate-impact liability for the same factual pattern pleaded under a different statute.

III. Why Product Liability Has Dominated and Settled

The plaintiffs' bar's preference for product liability pleading in AI-harm cases is explained by three structural factors: familiarity, insurer involvement, and settlement incentives.

The consumer product liability bar has decades of institutional experience with pleadings, discovery strategies, expert witnesses, and damages models in cases where a consumer product has allegedly caused physical or psychological harm. Firms with this experience pivot naturally to pleading AI-platform harms within that framework. Civil rights pleading, by contrast, requires competence with federal disability discrimination law, familiarity with the ADA and Rehabilitation Act jurisprudence, and an understanding of the federal financial assistance recipient doctrines under Section 504, none of which the consumer product liability bar typically maintains in-house.

Product liability claims trigger the operator's general liability and products liability insurance policies. The insurers then manage the defense, including settlement strategy. Insurers prefer confidential settlements that discharge specific claims without establishing precedent that would trigger broader policyholder exposure. Civil rights claims, by contrast, often fall outside the scope of general liability policies or within exclusions for intentional discrimination or regulatory compliance failures. The insurer is less involved, and the operator itself bears more of the strategic decision-making about defense posture, which changes the settlement incentives.

The product liability framework also produces a strong settlement equilibrium because the case-specific damages are bounded, compensable with money, and resolvable without requiring changes to the operator's product design. An operator settling a single product liability case pays the damages, signs a confidentiality agreement, and continues operating the product as designed. The civil rights framework disrupts that equilibrium. Title III's private-action remedy is injunctive-only, which means that resolving a Title III case requires the operator to change the challenged practice.²¹ Section 504's private-action remedy extends to compensatory damages within the limits of *Cummings*,²² but also to injunctive relief, and the Title II and Title III analogs establish that injunctive relief is the primary statutory remedy for disability discrimination. A disparate-impact class action, certified under Rule 23(b)(2),²³ aggregates hundreds or thousands of similarly situated plaintiffs in a single action whose resolution must address the design choice, not just the individual plaintiff's harm. The settlement-one-plaintiff-at-a-time structure breaks down when the pleading framework requires the operator to address the design choice.

The *Garcia* settlement is the paradigmatic example. After Judge Conway's favorable product-not-speech ruling,²⁴ after Character.AI survived the motion to dismiss, after the factual record was developed sufficiently to make the case triable, the parties settled in January 2026.²⁵ The Setzer family was presumably compensated under confidential terms. No admission of liability was made, no appellate ruling was produced, and no industry-wide precedent was established. The design choices at issue remain in place, and subsequent plaintiffs must re-litigate the threshold doctrinal questions from the outset.

21. 42 U.S.C. § 12188(a).

22. *Cummings v. Premier Rehab Keller, P.L.L.C.*, 596 U.S. 212 (2022).

23. Fed. R. Civ. P. 23(b)(2).

24. *Garcia*, 785 F. Supp. 3d at 1179.

25. *Garcia*, No. 6:24-cv-01903 (M.D. Fla. Jan. 7, 2026).

IV. Title III Public Accommodation Coverage of AI Platforms

A. The Enumerated Categories

Title III of the ADA prohibits disability discrimination by any “place of public accommodation.”²⁶ The statute defines public accommodations through twelve enumerated categories that include, among others, “a place of exhibition or entertainment,” “a place of public gathering,” “a sales or rental establishment,” “a service establishment,” “a place of education,” “a social service center establishment,” and “a place of exercise or recreation.”²⁷ The categories are construed to include operations that fall within the ordinary meaning of the enumerated terms even if the specific operation was not contemplated when the ADA was enacted in 1990.²⁸

B. Digital Public Accommodations After Robles

The question of whether digital platforms qualify as places of public accommodation has divided the circuits. The First and Seventh Circuits have held that Title III reaches websites and mobile applications regardless of whether the website has a physical analog.²⁹ The Third, Sixth, and Ninth Circuits have required a physical-nexus connection between the digital service and a place of public accommodation; the Eleventh Circuit’s panel decision imposing the same requirement was vacated as moot.³⁰ The Robles decision in the Ninth Circuit articulates the dominant version of the nexus-required framework, holding that Title III reaches a website or mobile application

26. 42 U.S.C. § 12182(a).

27. 42 U.S.C. § 12181(7)(C)–(L).

28. *See* H.R. Rep. No. 101-485, pt. 2, at 99 (1990) (the ADA was intended to “keep pace with the rapidly changing technology of the times”); H.R. Rep. No. 101-485, pt. 3, at 56 (1990) (same); S. Rep. No. 101-116, at 56–60 (1989) (discussing technological accommodations and accessibility). *see also* *Bostock v. Clayton County*, 590 U.S. 644, 655–85 (2020).

29. *Carparts Distrib. Ctr., Inc. v. Automotive Wholesaler’s Ass’n of New England*, 37 F.3d 12 (1st Cir. 1994); *Doe v. Mut. of Omaha Ins. Co.*, 179 F.3d 557 (7th Cir. 1999).

30. *Robles v. Domino’s Pizza, LLC*, 913 F.3d 898, 905–06 (9th Cir. 2019), *cert. denied*, 140 S. Ct. 122 (2019); *Parker v. Metropolitan Life Ins. Co.*, 121 F.3d 1006 (6th Cir. 1997) (en banc); *Ford v. Schering-Plough Corp.*, 145 F.3d 601 (3d Cir. 1998); *Access Now, Inc. v. Southwest Airlines, Co.*, 385 F.3d 1324 (11th Cir. 2004).

where the digital service is sufficiently connected to goods and services offered by a physical place of public accommodation.³¹

For AI chatbot platforms, the relevant doctrinal questions are (1) whether the platform constitutes a place of public accommodation under any of the enumerated categories, and (2) whether the platform must have a physical analog to qualify. On the first question, AI chatbot platforms fit within the ordinary meaning of “service establishment” and, depending on the specific functionality offered, also “place of education,” “place of entertainment,” and “social service center.”³² Character.AI markets itself as offering “characters” for users to interact with, which fits within entertainment. ChatGPT is marketed as a general-purpose assistant for tasks including coding, writing, analysis, and information retrieval, which fits within service establishment and place of education. Claude is similarly positioned. Gemini integrates into Google’s broader service offerings. On the second question, the platform-as-public-accommodation theory is strongest in circuits that do not require physical nexus. In circuits requiring nexus, the governing doctrine looks to physical places operated by the defendant, not physical places operated by third parties through which the defendant’s product is distributed. Microsoft operates Microsoft Store physical retail locations (until their closure in 2020 and reopening in modified form through Microsoft Experience Centers) as well as Microsoft Experience Centers and Microsoft Education partnerships with schools that include physical Microsoft-branded facilities; these are defendant-operated physical places supplying nexus in nexus-requiring circuits. Apple operates Apple retail stores nationwide; these are defendant-operated physical places supplying nexus for any Apple Intelligence or Apple AI deployment. Amazon operates Whole Foods Market, Amazon Fresh, Amazon Go, and Amazon Books physical retail operations; these are defendant-operated physical places supplying nexus for Alexa, Rufus, and other Amazon AI deployments. Google operates Google Fiber physical infrastructure and Google Store physical pop-up locations; these are defendant-operated physical

31. *Robles*, 913 F.3d at 905–07.

32. *New Mexico ex rel. Torrez*, No. D-101-CV-2023-02838.

places supplying nexus for Gemini. Character Technologies, Anthropic, and OpenAI are less well-positioned in nexus-requiring circuits, having no defendant-operated physical places of the kind the doctrine requires; plaintiffs bringing Title III claims against those defendants in nexus-requiring circuits should expect to face stronger motions to dismiss on the nexus question and should consider forum allocation to the First or Seventh Circuit, or alternative pleading through Section 504 host-defendant pathways where host-level coverage is established.

C. Title III Reaches AI Chatbot Platforms

Title III coverage of the major AI chatbot platforms is supported directly in the First and Seventh Circuits, which treat websites and mobile applications as public accommodations regardless of physical-place nexus, and is supported through the nexus theory in the Third, Sixth, and Ninth Circuits under *Robles* and its predecessors, with the question formally open in the Eleventh Circuit after the vacatur of *Gil*. The question remains open in circuits that have not squarely addressed digital-public-accommodation coverage of interactive AI services.³³ The better reading, on the facts of the major AI chatbot platforms' integration into physical-device and physical-service ecosystems (smartphones, smart home hardware, retail kiosks, educational institutions, health-care facilities, and financial-service points of sale), is that the platforms satisfy the nexus requirement in nexus-required circuits and qualify as public accommodations directly in non-nexus circuits. Plaintiffs pleading Title III claims outside the First and Seventh Circuits should include factual allegations establishing the relevant nexus sufficient to survive motion to dismiss under the governing circuit's formulation. The operator defendants reached by Title III in all circuits on that reading are: for Character.AI, Character Technologies, Inc.; for ChatGPT, OpenAI, Inc. and OpenAI, L.L.C.; for Claude, Anthropic, PBC; for Gemini, Google LLC; for Copilot, Microsoft Corporation. Each is subject to Title III suit by a disabled plaintiff who has suffered harm from design choices that produce disparate impact on users with disabilities.

33. *Magee v. Coca-Cola Refreshments USA, Inc.*, 833 F.3d 530, 534–35 (5th Cir. 2016) (holding that glass-front vending machines are not “sales establishments” under §12181(7)(E) because they do not resemble enumerated examples and are not “spaces in which individuals enter and utilize the services or goods provided” (citation and internal quotation marks omitted)), *cert. denied*, 138 S. Ct. 55 (2017).

D. The Market Replacement Analysis

The defense will argue that AI chatbots are the product, not the place where service is offered, analogizing to the doctrine that a bookstore must be accessible to wheelchair users but need not stock Braille editions.³⁴ That analogy fails on its own terms. A bookstore sells discrete, bounded items that the customer reads elsewhere. An AI chatbot is a service facility where the customer obtains an ongoing interactive service, at the time of access, through the platform's own operational architecture. The harm arises from the operator's real-time operation of the service, not from a static object the customer takes home.

The stronger response to the bookstore analogy comes from what the operators themselves have done to adjacent markets. AI chatbot platforms have been marketed, deployed, and documented as replacements for specific kinds of professional services: therapy, medical consultation, legal advice, academic tutoring, and companionship. Published research and the operators' own marketing materials describe ChatGPT, Claude, Character.AI, and competing platforms as functional substitutes for licensed clinical professionals in mental health care.³⁵ Clinical practitioners have documented and publicly reported the migration of patients from in-person therapy to AI chatbot substitutes, with the observed pattern concentrated in users who lack insurance coverage for clinical care or who prefer the continuous availability of the chatbot alternative. The operators have positioned these products within the service markets they replace. When a service displaces an in-person service that is itself a place of public accommodation (a therapist's office, a physician's practice, a tutoring center), the displacing service may inherit the coverage the displaced service bore. A platform that functions as a primary mental health support

34. U.S. Dep't of Justice, Civil Rights Div., *ADA Title III Technical Assistance Manual* §III-1.2000 (1993 & Supp. 1994) ("The examples given within each category are just illustrations" and the categories "should be construed liberally to afford people with disabilities equal access"); 28 C.F.R. §36.104 (defining "place of public accommodation" with parallel illustrative structure); *see also* H.R. Rep. No. 101-485, pt. 2, at 99–100 (1990) (House Committee Report confirming illustrative nature of enumerated examples).

35. *See generally* Peter Lee et al., *Benefits, Limits, and Risks of GPT-4 as an AI Chatbot for Medicine*, 388 New Eng. J. Med. 1233 (2023) (discussing consumer AI service scope); Ryan Calo, *Artificial Intelligence Policy: A Primer and Roadmap*, 51 U.C. Davis L. Rev. 399, 421–26 (2017) (analyzing regulatory challenges of AI services); Margot E. Kaminski, *Regulating the Risks of AI*, 103 B.U. L. Rev. 1347 (2023).

modality for some portion of its user base operates, for purposes of Title III analysis, in the same category as the clinical settings it substitutes for.

The market replacement analysis is a proposed doctrinal extension rather than a description of existing black-letter law, grounded in the § 794(b)(3)(A)(ii) subsidiary-level principal-engagement inquiry as applied to digital platforms. The better reading, on the facts of the major AI chatbot platforms' active positioning in mental-health-adjacent markets, is that platforms that function as substitutes for in-person clinical care enter the Title III enumerated category of "professional office of a health care provider" by functional substitution, not solely the broader categories of service establishment and social service center. The disparate impact of their design on users with mental health disabilities is disparate impact in a health-care-adjacent setting, which is precisely the context in which the ADA's reasonable-modifications obligation is most robust. The analogy to the telephonic commerce of the twentieth century, which the ADA has been held to reach as public accommodation,³⁶ extends to internet-delivered services and to the AI chatbot services of the present moment.

E. First Amendment and Section 230 Defenses

The defense will argue that injunctive orders requiring platforms to modify engagement-maximizing design constitute compelled speech in violation of the First Amendment, or that Section 230 immunizes platform operators from liability for content their models generate. Both arguments face significant contrary authority.

The Middle District of Florida held in *Garcia v. Character Technologies* that Character.AI's outputs constitute a product, not protected speech for First Amendment purposes.³⁷ Judge Conway's reasoning treats the AI system's output as the operational behavior of a computer system the operator controls, rather than the expression of a speaker for whom

36. *Doe v. Mut. of Omaha Ins. Co.*, 179 F.3d 557, 559 (7th Cir. 1999) ("The core meaning of this provision, plainly enough, is that the owner or operator of a store, hotel, restaurant, dentist's office, travel agency, theater, Web site, or other facility (whether in physical space or in electronic space) that is open to the public cannot exclude disabled persons from entering the facility and, once in, from using the facility in the same way that the nondisabled do."), *cert. denied*, 528 U.S. 1106 (2000).

37. *Garcia*, 785 F. Supp. 3d at 1179 (M.D. Fla. 2025) (denying motion to dismiss on First Amendment grounds).

the First Amendment provides protection. Injunctive orders requiring the operator to modify the system's operational behavior are no more compelled speech than injunctive orders requiring a building operator to install a ramp. The Supreme Court's decision in *Moody v. NetChoice, LLC*,³⁸ which recognized editorial-discretion interests in platforms' content moderation and curation of third-party speech, does not reach the distinct question of whether a platform's own operational conduct (as distinct from its treatment of third-party content) is protected speech. The New Mexico and Los Angeles trial courts have applied this first-party-design versus third-party-content-moderation distinction in permitting design-conduct claims to proceed notwithstanding the platforms' First Amendment and Section 230 defenses.

Section 230 is similarly unavailable. The harm complained of in an AI platform disability discrimination case is not third-party content the operator hosted; it is the operator's own design choices that determined how the platform responds to users with mental health disabilities. Section 230 immunity reaches interactive computer services for the publication of information provided by another information content provider.³⁹ When the information originates with the operator's own model, and when the discriminatory design choice is the operator's own, the operator is both the interactive computer service and the information content provider.⁴⁰ The operator's marketing statements attributing the outputs to the operator are inconsistent with a subsequent Section 230 position disclaiming authorship for liability purposes.

F. Title III Remedies

Title III authorizes private civil actions for injunctive relief.⁴¹ The Attorney General and state attorneys general may also sue under Title III and may obtain civil penalties and compensatory

38. *Moody v. NetChoice, LLC*, 603 U.S. 707 (2024).

39. 47 U.S.C. § 230(c)(1).

40. *See Fair Housing Council of San Fernando Valley v. Roommates.com, LLC*, 521 F.3d 1157, 1162–65 (9th Cir. 2008) (en banc) (Section 230 immunity unavailable where operator's design solicits or requires the allegedly unlawful content); *Lemmon v. Snap, Inc.*, 995 F.3d 1085 (9th Cir. 2021) (design-feature claims not barred by Section 230 where claims target operator conduct rather than third-party content).

41. 42 U.S.C. § 12188(a).

damages in actions brought in their official capacity.⁴² The private plaintiff may recover reasonable attorney’s fees.⁴³ The absence of compensatory damages in the private action is the primary structural limitation of Title III litigation, but this limitation does not make Title III a weaker framework than product liability for the purposes of design change. An injunction requiring the operator to modify engagement-maximizing design choices to reduce disparate impact on users with disabilities produces systemic relief that no product liability settlement has produced or will produce.

V. Section 504 Coverage of the Corporate Hosts

The pathways developed in this Part vary in the strength of their evidentiary basis and in doctrinal confidence, as summarized in Table 1.

Table 1: Section 504 Recipient Pathways by Evidentiary Basis and Doctrinal Strength

Defendant	Pathway	Evidentiary Basis	Doctrinal Confidence
Microsoft Corp.	DOD quantum award	Award documented (USAspending.gov)	Strong basis; scope fact-specific
Microsoft Corp.	USDA biothreat award	Award documented (USAspending.gov)	Strong basis; scope fact-specific
Microsoft Corp.	NCMS sub-grants	Intermediary structure documented	Strong basis; scope fact-specific
Google LLC	Fitbit LLC NIH sub-grant	Sub-award documented (NIH via Scripps)	Strong at subsidiary level; parent coverage contested
Google LLC	Fitbit principal engagement § 794(b)(3)(A)(ii)	Documented at subsidiary level	Proposed extension at subsidiary level; parent flow-up contested
Amazon, Apple	Health-care verticals	Federal health-program participation not established in public award data	Proposed; depends on the entity’s Medicare or Medicaid provider participation

42. 42 U.S.C. § 12188(b).

43. 42 U.S.C. § 12205.

Defendant	Pathway	Evidentiary Basis	Doctrinal Confidence
Operator defendants (OpenAI, Anthropic, Character Tech.)	Direct Title III, no recipient status needed	N/A for Section 504	Strong: see Part IV
Federal-agency deployers	Section 504(a) direct coverage	Deployment varies by agency	Strong where deployment is public-facing
State/local deployers	Title II coverage	Deployment varies by entity	Strong where deployment is public-facing
Private public-accommodation deployers	Title III coverage	Deployment varies by entity	Strong where the deployer operates a public accommodation

A. The Whole-Corporation Rule and the Civil Rights Restoration Act

Section 504 of the Rehabilitation Act prohibits disability discrimination “under any program or activity receiving Federal financial assistance.”⁴⁴ The statute, as amended by the Civil Rights Restoration Act of 1987,⁴⁵ defines “program or activity” to include “all of the operations of” a private corporation, partnership, or other private organization that either (1) receives federal financial assistance as a whole, or (2) is principally engaged in the business of providing education, health care, housing, social services, or parks and recreation.⁴⁶ The CRRA’s legislative history reflects Congress’s specific intent to reverse *Grove City College v. Bell*, in which the Supreme Court had narrowly construed “program or activity” to cover only the specific program receiving federal assistance, rather than the recipient’s broader operations.⁴⁷ The CRRA’s expansive definition of “program or activity” is intended to defeat corporate-structure arguments that would narrow the statute’s reach.

The whole-corporation rule under § 794(b)(3)(A)(i) reaches every operation of a corporation that has received federal financial assistance as a whole. The principal-engagement rule under § 794(b)(3)(A)(ii) reaches every operation of a corporation principally engaged in the

44. 29 U.S.C. § 794(a).

45. Pub. L. No. 100-259, 102 Stat. 28 (1988).

46. 29 U.S.C. § 794(b)(3)(A)–(B).

47. *Grove City College v. Bell*, 465 U.S. 555 (1984).

enumerated service categories when any part of that corporation receives federal assistance. Both rules produce coverage that extends beyond the specific program or activity directly receiving the federal dollars.⁴⁸

B. Microsoft Corporation: Copilot, Windows, Bing, Teams, GitHub

Microsoft Corporation is a Section 504 recipient through the federal assistance awards addressed immediately below. The reach of that recipient status to Microsoft's directly operated AI products, including Copilot, is taken up after the awards are described.

Microsoft has received federal assistance awards that supply cumulative support for recipient status. These include a Department of Defense assistance award of approximately \$6.5 million for quantum computing research, a United States Department of Agriculture assistance award of approximately \$1 million for biothreat-related computational research, and participation in National Center for Manufacturing Sciences sub-grants under various federal research and development programs. Microsoft Corporation's federal award recipient profile is searchable at USAspending.gov.⁴⁹ Each of these awards is an assistance award rather than a procurement contract within the meaning of the Federal Grant and Cooperative Agreement Act of 1977, 31 U.S.C. §§ 6303–6305, and therefore triggers Section 504 recipient status.⁵⁰ Whether these awards independently establish whole-corporation coverage under § 794(b)(3)(A)(i) or facility-specific coverage under § 794(b)(3)(B) depends on the structure of each award; where coverage is program- or facility-specific, the reach to Copilot as an operative product proceeds under Title III rather than Section 504.

Where the assistance awards support whole-corporation coverage under § 794(b)(3)(A), that coverage reaches Microsoft Corporation's directly operated products. These include

48. *Moody v. NetChoice, LLC*, 603 U.S. 707 (2024).

49. See USAspending.gov, Microsoft Corporation federal award recipient profile, UEI INLWM5KDS7R9, accessible at <https://www.usaspending.gov/recipient/dd77b7c3-663e-cb91-229f-5766a50e9b7f-P/a11> (last visited June 10, 2026).

50. See *Nat'l Collegiate Athletic Ass'n v. Smith*, 525 U.S. 459 (1999) (recipient status attaches to entities that receive federal financial assistance); cf. *U.S. Dep't of Transp. v. Paralyzed Veterans of Am.*, 477 U.S. 597 (1986) (entities that merely benefit from federal assistance are not recipients).

Windows, Office 365, Azure (including Azure OpenAI Service, where OpenAI’s models are deployed for Microsoft customers), Bing, Edge, Microsoft Teams, Microsoft Power Platform, and Microsoft Copilot in all of its deployments, including Copilot in Windows, Copilot in Microsoft 365, Copilot in Bing, GitHub Copilot, Dynamics 365 Copilot, and Security Copilot. The Copilot family is relevant because Copilot is deeply integrated into the AI deployment surface that reaches Microsoft users across the Microsoft product ecosystem. Where harm to a user with a disability arises from Copilot’s conversational design, engagement mechanics, or output characteristics, that harm occurs within Microsoft Corporation’s operation of a product covered by Section 504.

Coverage of separately incorporated Microsoft subsidiaries, including LinkedIn Corporation, GitHub, Inc., Activision Blizzard, Inc., and ZeniMax Media, Inc., requires independent recipient-status analysis for each subsidiary.⁵¹

C. Google LLC and Fitbit LLC: Gemini Reached Through the Fitbit Pathway

The cleanest Section 504 pathway to Gemini runs through Fitbit LLC, a separately incorporated Google subsidiary that is an independent Section 504 recipient and that directly operates the Gemini-powered Personal Health Coach launched in October 2025.

1. Fitbit LLC as an Independent Section 504 Recipient

Fitbit LLC is a confirmed recipient of federal financial assistance through sub-grants from the Scripps Research Translational Institute under the National Institutes of Health’s All of Us Research Program.⁵² This triggers Section 504 coverage under both (b)(3)(A)(i)

51. The whole-corporation rule under § 794(b)(3)(A) reaches “all of the operations of” a recipient corporation. Under dominant readings, this phrase reaches within the recipient’s corporate boundary but does not automatically flow to separately incorporated subsidiaries. A more expansive reading, supported by some elements of the CRRA’s legislative history, would reach wholly-owned subsidiaries as operations of the parent. This Article does not resolve the subsidiary-piercing question.

52. Scripps Research is funded by National Institutes of Health award OT2OD035580. Fitbit is a named Participant Center partner organization in the All of Us consortium. *See* Scripps Research, “Scripps Research receives momentous award from NIH to lead key programs in national All of Us Research Program” (July 2023) (announcing \$54 million initial funding with projected total \$282 million over five years, naming Fitbit as a Participant Center partner). Fitbit’s integration with All of Us was formally announced by NIH in 2019, and Scripps Research and

whole-corporation analysis and, separately, (b)(3)(A)(ii) principal-engagement analysis, because Fitbit is principally engaged in the business of providing health care services within the meaning of the statute.

The record supporting Fitbit’s principal engagement in health care is substantial. The company’s products collect comprehensive health data including continuous heart rate, sleep staging, electrocardiogram readings, blood oxygen saturation, respiratory rate, skin temperature, stress metrics, and menstrual cycle tracking. The Fitbit application transmits user health data to healthcare providers upon user consent and, as of early 2026, integrates user medical records from healthcare provider portals through partnerships with b.well and CLEAR.⁵³ Fitbit’s research operations include published studies on insulin resistance prediction and hypertension detection. The company partners with Included Health for virtual medical visits and with academic medical institutions for population-scale health research.

2. Fitbit’s Operation of the Gemini-Powered Personal Health Coach

In October 2025, Fitbit launched a public preview of a Gemini-powered Personal Health Coach, which Google Research describes as a “Personal Health Agent” with multi-agent orchestration for data analysis, domain expertise, and conversational coaching.⁵⁴ The health coach handles fitness programming, sleep analysis, stress and recovery guidance, menopause symptom management, doctor visit preparation, health history summarization, and health condition

Fitbit launched the Wearables Enhancing All of Us Research (WEAR) study distributing 10,000 Fitbit devices to program participants in 2021. *See* Scripps Research, “Through ‘All of Us’ program, Scripps Research launches wearable technology study to accelerate precision medicine” (Feb. 2021). Recent peer-reviewed work documents the continuing data flow. *See* Caitlin P. Bailey et al., *Fitbit Physical Activity and Sleep Data in the All of Us Research Program: Data Exploration and Processing Considerations for Research*, 57 *Med Sci Sports Exerc* 2946 (Dec. 2025), doi:10.1249/MSS.0000000000003804.

53. Google, “Get a fuller picture with Fitbit’s personal health coach” (The Keyword blog, Mar. 19, 2026).

54. Google, “Fitbit’s personal health coach is now available in public preview” (The Keyword blog, Oct. 27, 2025), available at <https://blog.google/products-and-platforms/devices/fitbit/personal-health-coach-public-preview/>; Google Research Blog, “How we are building the personal health coach” (2025–2026 series); Fitbit Enterprise Blog, “Introducing the new personal health coach powered by Gemini.” *See also* TechCrunch, “Fitbit’s revamped app, with Gemini-powered health coach, rolls out to Premium users” (Oct. 27, 2025); MobiHealthNews, “Google debuts Gemini-powered AI health coach for Fitbit and Pixel Watch users” (Oct. 2025). The public preview launched October 28, 2025 for Fitbit Premium users in the United States on Android, with iOS rollout and full launch scheduled for 2026.

explanations. The user-facing experience is a conversational AI agent built on Gemini models, operating on user-consented health data and, for users who have linked their medical records through the b.well or CLEAR integrations, on clinical records retrieved from healthcare provider portals.

The doctrinal consequence of this architecture is significant. Fitbit LLC's Section 504 recipient status, established through both direct assistance and principal engagement, confers whole-corporation coverage of Fitbit LLC's operations. Fitbit LLC operates the Gemini-powered Personal Health Coach. The Coach's design choices, including its conversational behavior, its handling of mental health topics, its response patterns for users in crisis or with symptoms suggesting psychiatric conditions, and its engagement mechanics, are Fitbit's operation of a product covered by Section 504. Where a Fitbit user with a mental health disability experiences harm arising from Gemini's design as deployed in the Fitbit Personal Health Coach, Section 504 liability attaches to Fitbit LLC directly, independent of whether coverage also flows to Google LLC through the parent or through any other pathway. Fitbit LLC is its own recipient, and its operation of Gemini is its own covered activity.

3. Google LLC Direct Recipient Pathways

Google LLC separately receives federal financial assistance through at least one confirmed pathway: approximately \$3 million in cumulative state-by-state Affordable Connectivity Program reimbursement payments, extended by the Federal Communications Commission for Google Fiber participation as a broadband provider to low-income households. These payments constitute direct federal financial assistance to Google LLC. The settled consequence is facility-level coverage under (b)(3)(B) of the Google Fiber broadband operation itself as the program or activity to which the assistance extends. Whether the ACP payments also support (b)(3)(A)(i) whole-corporation coverage of Google LLC's operations more broadly, under the "all of the operations" language of the Civil Rights Restoration Act, is a contested question that this Article presents as a supporting argument rather than as settled law; the primary pathway

to Gemini coverage in this Article runs through Fitbit LLC's subsidiary-level recipient status as operator of the Personal Health Coach, not through parent-level flow-up from ACP.⁵⁵

Google LLC's operations include Google Search, Gmail, Google Maps, YouTube (operated by the separately incorporated YouTube LLC, which requires independent recipient-status analysis), the Google Play Store, Google Cloud Platform (including Vertex AI, which hosts Gemini and third-party models including Claude), Google Workspace, and Gemini in all its consumer and enterprise deployments.

D. Amazon and Apple: Contingent Health-Care Pathways

Amazon.com, Inc. and Apple, Inc. both operate health-care-adjacent product lines that may support (b)(3)(A)(ii) principal-engagement analysis. For Amazon, the relevant verticals are One Medical (primary care), PillPack (mail-order pharmacy), and Amazon Pharmacy. For Apple, the relevant verticals are Apple Health, Apple Watch health monitoring, CareKit, and HealthKit. In both cases, a principal-engagement argument requires (1) identification of the specific legal entity operating the health-care vertical, (2) verification that the vertical receives federal financial assistance through Medicare or Medicaid provider participation, clinical research grant participation, or other federal health programs, and (3) doctrinal argument that the principal-engagement inquiry applies at the corporate parent level rather than only at the vertical subsidiary.

Current USAspending federal award data does not confirm direct federal assistance to One Medical, 1Life Healthcare, PillPack, or Amazon Pharmacy under the entity names associated with these acquisitions. This may reflect the fact that Medicare and Medicaid provider participation funding flows through Centers for Medicare and Medicaid Services structures

55. A secondary pathway not relied upon in the primary analysis of this Article involves approximately \$46 million in political risk insurance extended by the Overseas Private Investment Corporation (now the Development Finance Corporation) for a Google-owned wind power investment in Kenya, under CFDA 70.003. The characterization of OPIC political risk insurance as federal financial assistance for civil rights statute purposes is contested; courts and agencies that have considered the question have divided on whether such insurance constitutes assistance or a contract for insurance services. This Article does not rely on the OPIC pathway. The Fitbit and ACP pathways supply adequate recipient-status foundations without reaching the disputed OPIC characterization.

that are not uniformly captured in USAspending at the entity level. The principal-engagement pathway to Amazon remains available contingent on the relevant entity's federal health-program participation through CMS records or parallel sources.

Apple's health-vertical principal-engagement pathway is analogous and similarly contingent on the relevant entity's federal health-program participation. Neither Amazon nor Apple is foundational to the analytical framework developed in this Article; both are noted as pathways that plaintiffs' counsel may pursue contingent on establishing the underlying federal assistance.

E. Direct Pathways to OpenAI and Anthropic

Neither OpenAI nor Anthropic is currently a direct Section 504 recipient. Both companies have structured their federal government relationships as procurement contracts rather than assistance grants. OpenAI has received a \$200 million Department of Defense contract for the OpenAI for Government product, in addition to procurement contracts with the FAA, NASA, USAID, the Internal Revenue Service, the National Gallery of Art, Los Alamos National Laboratory, and the Air Force Research Laboratory. Anthropic has received a \$200 million DOD contract in July 2025, a GSA OneGov Multiple Award Schedule nominal-price agreement in August 2025, and has partnered with Palantir and Amazon Web Services for access to United States intelligence and defense agencies. Under the Federal Grant and Cooperative Agreement Act of 1977,⁵⁶ these procurement contracts do not trigger Section 504 recipient status, because the statute distinguishes between procurement contracts (for which Section 504 does not apply) and assistance awards (for which it does).

Anthropic's SAM.gov registration records list its purpose of registration as "Federal Contracts and Assistance," indicating the company has registered for eligibility to receive federal assistance-type awards in addition to procurement contracts. If Anthropic subsequently receives

56. 31 U.S.C. §§ 6303–6305.

any actual cooperative agreement or grant, Section 504 recipient status would attach at that point. As of this writing, the authors are aware of no such assistance award to Anthropic.

The practical consequence is that plaintiff strategy against OpenAI and Anthropic cannot rely on direct Section 504 recipient status. The available pathways are (1) Title III public accommodation coverage, which reaches OpenAI and Anthropic directly regardless of federal-assistance status, and (2) Section 504 coverage of the corporate hosts that deploy OpenAI and Anthropic models inside their own product surfaces.

For OpenAI, the relevant corporate host is Microsoft Corporation, whose Section 504 coverage reaches Microsoft's operation of Copilot, Azure OpenAI Service, Bing, and other Microsoft surfaces in which OpenAI models are deployed. A plaintiff harmed by an OpenAI model encountered through a Microsoft product surface states a Section 504 claim against Microsoft Corporation. OpenAI's own recipient status is not required to reach that plaintiff's harm. Microsoft's deep commercial integration with OpenAI (including exclusive cloud partnership rights, equity participation in OpenAI's for-profit arm, and direct integration of OpenAI models into Microsoft products) may additionally support a joint-operator theory reaching Microsoft for ChatGPT harms encountered through any Microsoft-adjacent distribution channel, though the joint-operator theory is more aggressive than the pure Microsoft-product pathway.

For Anthropic, the relevant corporate hosts are Google LLC (via Vertex AI deployment of Claude on Google Cloud Platform, reached through Google LLC's verified Section 504 recipient pathways including the Fitbit subsidiary) and Amazon.com, Inc. (via AWS Bedrock deployment of Claude, reached through Amazon's health-vertical principal-engagement analysis, contingent on the relevant entity's federal health-program participation). A plaintiff harmed by Claude deployed through Vertex AI states a Section 504 claim against Google LLC. A plaintiff harmed by Claude deployed through AWS Bedrock states a Section 504 claim against Amazon contingent on the federal-assistance question noted above. In both cases, Anthropic's own recipient status is not required to reach the plaintiff's harm.

F. The Deployer Pathway: Title II and Section 504(a)

A third pleading pathway, cumulative with the direct-operator pathway of Title III and the corporate-host pathway of Section 504, reaches the public-entity and federal-agency deployers of commercial chatbot products. This pathway does not hold the platform operator or the corporate host directly liable. It reaches the entity that chose to deploy the product to its disabled users in its own public-facing capacity. Design change flows back to the operator through procurement: an operator whose product places deployer-customers in violation of Title II or Section 504(a) must modify the product or lose the customer.

1. Federal Agency Deployers Under Section 504(a)

Section 504(a) reaches federal agencies directly, not as recipients of federal financial assistance but as the government itself.⁵⁷ Every federal agency deploying a commercial chatbot product in a public-facing capacity is subject to Section 504(a) compliance obligations in that deployment. The scope of current deployment includes OpenAI for Government under multi-agency contracts (Department of Defense, General Services Administration, Federal Aviation Administration, National Aeronautics and Space Administration, United States Agency for International Development, Internal Revenue Service, National Gallery of Art, Los Alamos National Laboratory, and Air Force Research Laboratory), Anthropic’s GSA OneGov Multiple Award Schedule agreement, Microsoft Copilot in federal workflows, and Google Gemini through Google Cloud Platform federal offerings. A disabled member of the public who encounters an AI chatbot deployed by a federal agency in a public-facing capacity, and who experiences disparate impact from that deployment, has a Section 504(a) claim against the agency.

57. 29 U.S.C. § 794(a) (prohibiting disability discrimination “under any program or activity conducted by any Executive agency”).

2. State and Local Government Deployers Under Title II

Title II of the ADA reaches public entities, defined to include state and local governments and their instrumentalities.⁵⁸ State universities, state agencies, counties, municipalities, public school districts, public hospitals, and their contracted instrumentalities are all public entities for Title II purposes. Deployment of commercial chatbot products by these entities in public-facing capacities is Title II activity subject to the reasonable-modifications obligation and to the meaningful-access standard.

Title II rests on Commerce Clause and Fourteenth Amendment enforcement authority rather than on Spending Clause conditions.⁵⁹ The argument that Title II's non-Spending Clause foundation insulates it from the *Cummings* limitation on emotional-distress damages has been considered and rejected by every federal court of appeals to have reached the question.⁶⁰ The controlling analysis tracks the statutory remedies-incorporation chain: Title II at 42 U.S.C. § 12133 incorporates the remedies, procedures, and rights of Section 505 of the Rehabilitation Act, which at 29 U.S.C. § 794a(a)(2) in turn incorporates the remedies, procedures, and rights of Title VI of the Civil Rights Act of 1964. *Barnes v. Gorman*,⁶¹ held that the ADA's explicit incorporation of Rehabilitation Act remedies makes the statute's constitutional origin irrelevant to the remedial analysis. *Cummings* foreclosed emotional-distress damages under Title VI and the Spending Clause statutes that incorporate it. The result is that emotional-distress damages are not available under Title II of the ADA against public-entity deployer defendants. The Title II damages picture against a public-entity deployer is therefore comparable to the Section 504 picture against a private-recipient corporate host: both are limited to economic and pecuniary harm. The Ninth Circuit in *Payan* expressly preserved compensatory damages for lost opportunities attributable to Title II discrimination, which on the AI-platform fact pattern reaches

58. 42 U.S.C. § 12131(1).

59. See *Tennessee v. Lane*, 541 U.S. 509 (2004); *Bd. of Trustees of Univ. of Ala. v. Garrett*, 531 U.S. 356 (2001).

60. *Doherty v. Bice*, 101 F.4th 169, 175–76 (2d Cir. 2024); *A.W. ex rel. J.W. v. Coweta Cnty. Sch. Dist.*, 110 F.4th 1309, 1314–15 (11th Cir. 2024); *Payan v. Los Angeles Cmty. Coll. Dist.*, 169 F.4th 971, 977–79 (9th Cir. 2026).

61. *Barnes v. Gorman*, 536 U.S. 181, 189 n.3 (2002).

lost educational opportunities, lost earning capacity, lost access to clinical care, and the economic component of wrongful death. State-law parallel claims under state disability-rights statutes remain the appropriate vehicle for non-economic damages where state law authorizes them.

3. Private Public-Accommodation Deployers Under Title III

A private entity covered by Title III as a place of public accommodation that deploys a commercial chatbot product in a capacity reaching its customers or users is subject to Title III in that deployment. A retail chain operating a chatbot customer service channel on its website and in its apps is a Title III defendant for design choices in that deployment, independent of the chatbot vendor's own Title III exposure. The telephonic commerce analog is instructive: when the ADA was adopted, the question of whether telephonic commerce was reached by Title III was answered by reference to the underlying function rather than the medium. A retailer that sold goods over the telephone was subject to Title III in that sales activity because the sales activity was accessible to the public regardless of medium. The extension to internet commerce and to AI-chatbot-mediated commerce follows the same logic. A retailer that deploys a chatbot customer service is subject to Title III in that deployment, and its design choices in the deployment are subject to the reasonable-modifications obligation.

4. Cumulative Pleading and Non-Preclusion

The three deployer pathways and the direct-operator pathway are cumulative, not exclusive. A single complaint can plead all that apply. Res judicata does not bar a later plaintiff from bringing claims against the same defendant under the same pathway if that plaintiff is a different party with different facts.⁶² Each plaintiff has her own cause of action. Non-mutual defensive collateral estoppel under *Blonder-Tongue Laboratories, Inc. v. University of Illinois Foundation* reaches only specific issues actually litigated and actually decided in a prior case, and does not reach factual findings that differ between plaintiffs.⁶³ The framework supports serial enforcement by

62. See *Taylor v. Sturgell*, 553 U.S. 880 (2008) (reaffirming the general rule against nonparty preclusion).

63. *Blonder-Tongue Labs., Inc. v. Univ. of Illinois Found.*, 402 U.S. 313 (1971).

multiple plaintiffs across multiple defendants without preclusion concerns that would bar later filings on different facts.

VI. The Scope of “Disability” Under Section 504 and Title III

A. The Governing Standard

The ADA, as amended by the ADA Amendments Act of 2008,⁶⁴ defines “disability” to include any physical or mental impairment that substantially limits one or more major life activities.⁶⁵ The statute directs that “disability” “shall be construed in favor of broad coverage.”⁶⁶ Section 504’s definition of “disability” is functionally coextensive with the ADA’s definition as applied to private recipients.

For the purposes of AI platform harm pleading, the relevant disabilities include clinical depression, generalized anxiety disorder, bipolar disorder, schizoaffective disorder, schizophrenia spectrum conditions, post-traumatic stress disorder, attention-deficit hyperactivity disorder, autism spectrum conditions, borderline personality disorder, suicidal ideation as a symptom of underlying mental health conditions, and, following the Fourth Circuit’s holding in *Williams v. Kincaid*, gender dysphoria.

B. Williams v. Kincaid and Gender Dysphoria Coverage

The Fourth Circuit held in *Williams v. Kincaid* that gender dysphoria is a disability under the ADA and Section 504.⁶⁷ The holding is binding within the Fourth Circuit (Maryland, North Carolina, South Carolina, Virginia, West Virginia) and is persuasive authority elsewhere. Circuits outside the Fourth have divided on whether to adopt its reasoning. The decision establishes that

64. Pub. L. No. 110-325, 122 Stat. 3553 (2008).

65. 42 U.S.C. § 12102(1)(A).

66. 42 U.S.C. § 12102(4)(A).

67. *Williams v. Kincaid*, 45 F.4th 759 (4th Cir. 2022), *cert. denied*, 143 S. Ct. 2414 (2023).

the ADA's exclusion of "gender identity disorders not resulting from physical impairments"⁶⁸ does not reach gender dysphoria as a modern clinical diagnosis.

Williams v. Kincaid is doctrinally significant for AI-platform pleading because transgender users and users experiencing gender dysphoria are a demographic concentrated in intensive chatbot use for persona exploration, social rehearsal, and emotional support in the absence of other accepting interlocutors. Where AI chatbot platforms' design choices produce disparate harm to users with gender dysphoria, Section 504 and Title III claims are available under the Kincaid framework within the Fourth Circuit and are pleadable as persuasive-authority arguments elsewhere. A recently proposed rule by the Department of Health and Human Services would modify 45 C.F.R. § 84.4(g) to narrow this coverage.⁶⁹ The proposed rule has not been finalized as of this writing, and its substantive validity is open to challenge under the ADA Amendments Act's directive that disability definitions be construed in favor of broad coverage.

C. Pre-Existing Mental Health Disabilities and the Disparate-Impact Class

The factual pattern of AI-induced psychological harm concentrates in users with pre-existing mental health conditions and neurodevelopmental disabilities. The engagement-maximizing design choices that produce harm are calibrated to produce continued use, and users whose psychiatric and cognitive profiles reduce their capacity to disengage from reinforcing stimuli are the users most susceptible to those calibrations. The meaningful-access and reasonable-modifications theories developed in this Article rest on the observation that the platform's facially neutral design features produce disproportionate harm to this class of users and fail to provide meaningful access to the service for them without reasonable modification.

The relevant class for AI-platform pleading under this framework includes users with documented diagnoses of mental health or neurodevelopmental disabilities prior to the harm onset. The framework does not extend to users with no pre-existing ADA-recognized disability whose psychiatric symptoms emerged entirely as a consequence of platform engagement; those

68. 42 U.S.C. § 12211(b)(1).

69. 89 Fed. Reg. (Dec. 19, 2025) (HHS proposed rulemaking on gender dysphoria, RIN 0945-AA27).

users have causes of action in product liability and in other theories outside this Article’s scope. Expert evidence establishing the causal or contributory role of platform design in exacerbating or failing to accommodate the pre-existing disability is available from clinical psychologists, psychiatrists, and digital-harm researchers who have developed the AI-induced psychosis literature over the last three years.

VII. Meaningful Access, Reasonable Modifications, and Disparate Impact Pleading

A. The Doctrinal Framework

Section 504 requires that persons with disabilities be provided meaningful access to the programs and activities of federal-assistance recipients, which may require reasonable modifications to facially neutral policies and practices.⁷⁰ Title III of the ADA imposes a parallel reasonable-modifications obligation on public accommodations where necessary to afford goods, services, facilities, privileges, advantages, or accommodations to individuals with disabilities, unless the entity can demonstrate that making such modifications would fundamentally alter the nature of the goods or services.⁷¹

The doctrinal predicate for this Article’s theory is denial of meaningful access and failure to make reasonable modifications, not pure disparate impact. This distinction is essential because *Alexander v. Sandoval* held that no private right of action exists to enforce regulations promulgated under Title VI of the Civil Rights Act of 1964 that go beyond the statute’s anti-discrimination prohibition, including regulations imposing liability for facially neutral conduct producing disparate impact on protected classes.⁷² Section 504’s private enforcement

70. *Choate*, 469 U.S. at 301 (“meaningful access” standard); *Se. Cmty. Coll. v. Davis*, 442 U.S. 397 (1979) (reasonable modifications).

71. 42 U.S.C. § 12182(b)(2)(A)(i)–(ii).

72. *Alexander v. Sandoval*, 532 U.S. 275, 284–93 (2001).

provision incorporates Title VI's enforcement structure by reference,⁷³ which has led the lower courts to apply *Sandoval* to foreclose private suits brought under Section 504 regulations that extend beyond the statute's own prohibition.

The Ninth Circuit addressed the limit of that foreclosure in *Mark H. v. Lemahieu*, holding that Section 504 regulations properly interpreted as meaningful-access regulations are enforceable through the statute's implied private right of action, because the meaningful-access requirement is a construction of the statutory prohibition on discrimination rather than an independent regulatory overlay of the kind at issue in *Sandoval*.⁷⁴ The framework proposed in this Article operates within the meaningful-access and reasonable-modifications doctrine, not within a pure disparate-impact theory. Evidence of disparate impact on users with disabilities functions as evidentiary support for the meaningful-access denial, not as the legal theory itself. A design feature that produces disproportionate psychological harm to users with a mental-health disability, autism spectrum condition, attention-deficit condition, or other disability recognized under the ADA fails to provide meaningful access to that user population, within the meaning of *Choate*.⁷⁵ Continued deployment of the design feature by the operator after receiving notice that it produces such disproportionate harm satisfies the deliberate-indifference standard for compensatory damages under Section 504.⁷⁶

The practices subject to reasonable-modifications analysis for AI chatbot platforms include: the persistence of engagement across user sessions, the failure to implement reality-reminder interventions on extended sessions, the failure to detect and redirect on crisis-signal conversation content, the failure to provide user-facing controls for disengagement,

73. 29 U.S.C. § 794a(a)(2) (incorporating “the remedies, procedures, and rights set forth in title VI of the Civil Rights Act of 1964”); *Barnes v. Gorman*, 536 U.S. 181, 185 (2002) (treating Section 504 and Title VI enforcement as coextensive).

74. *Mark H. ex rel. Michelle H. v. Lemahieu*, 513 F.3d 922, 937–38 (9th Cir. 2008) (“to the degree the Section 504 . . . regulations . . . can be interpreted as a variety of meaningful access regulation, they will fall within the Section 504 implied cause of action,” distinguishing such regulations from the disparate-impact regulations at issue in *Sandoval* on the ground that meaningful-access regulations focus on “design” rather than “effect” and are best understood as construing rather than extending the statutory prohibition).

75. *Choate*, 469 U.S. at 301.

76. *Duvall v. County of Kitsap*, 260 F.3d 1124, 1139 (9th Cir. 2001) (deliberate indifference is “knowledge that a harm to a federally protected right is substantially likely, and a failure to act upon that likelihood”).

the failure to hand off to human professionals on identified mental health crisis indicators, and the platform’s specific design choices that produce parasocial attachment. Each of these is a policy, practice, or procedure subject to modification analysis.⁷⁷ The defendant’s burden to show fundamental alteration is high.⁷⁸ Reality-reminder interventions and crisis-signal handoffs do not fundamentally alter a chatbot platform; they modify its behavior at specific moments to avoid disability-related harm.⁷⁹

B. Deliberate Indifference and Actual Knowledge

Section 504 private actions for compensatory damages require a showing of intentional discrimination or deliberate indifference.⁸⁰ The deliberate-indifference standard is satisfied where the recipient had actual knowledge that its practices produced disparate impact on persons with disabilities and failed to make reasonable modifications.⁸¹

The deliberate-indifference standard is particularly important for AI-platform pleading because the major AI platform operators have been on actual notice of the AI-induced psychological harm pattern for multiple years. Internal company documents, public reporting, and voluntary safety commitments made by OpenAI, Anthropic, Google, Meta, and Character.AI all demonstrate actual knowledge of the harm pattern.⁸² Each operator has publicly acknowledged receiving reports of user harm, has publicly committed to safety measures in response, and in several cases has publicly described internal research on the harm pattern.

77. *Blonder-Tongue Labs., Inc. v. Univ. of Illinois Found.*, 402 U.S. 313 (1971).

78. *See PGA Tour, Inc. v. Martin*, 532 U.S. 661 (2001) (individualized fundamental-alteration inquiry under Title III).

79. 42 U.S.C. §12182(b)(2)(A)(iii); 28 C.F.R. §36.303(a). The analysis considers the same factors as readily achievable: nature and cost, overall financial resources, type of operations, and impact on operations.

80. *See, e.g., Liese v. Indian River County Hospital District*, 701 F.3d 334 (11th Cir. 2012); *Payan v. Los Angeles Community College District*, 11 F.4th 729 (9th Cir. 2021). *Duvall*, 260 F.3d at 1139.

81. *See, e.g., Lovell v. Chandler*, 303 F.3d 1039, 1056–57 (9th Cir. 2002); *Duvall*, 260 F.3d at 1138–41.

82. *See supra* note 3.

Internal documents produced in the *Garcia* and *Raine* litigation establish the factual record of actual knowledge on the part of the named defendants.⁸³

The actual-knowledge record on OpenAI is specifically documented. In June 2025, OpenAI flagged and banned the ChatGPT account of an eighteen-year-old user in British Columbia, Canada, for “furtherance of violent activities.”⁸⁴ Twelve OpenAI employees reviewed the account content and considered whether to escalate the risk to law enforcement. They concluded that the content did not meet the threshold for external escalation and did not notify authorities. On February 10, 2026, the eighteen-year-old returned home, killed her mother and half-brother, and then attacked a secondary school in Tumbler Ridge, killing five students and a teaching assistant before taking her own life.⁸⁵ Eight people died. After the attack, OpenAI acknowledged that its internal protocols had failed and informed the Canadian government that, under revised referral guidelines, it would have reported the account to law enforcement.

The Tumbler Ridge and FSU shooting records supply evidence of the operator’s actual knowledge and deliberate indifference rather than plaintiff templates for disability discrimination pleading. The plaintiff standing for disability discrimination under this framework runs through disabled users harmed by the product, not through third parties harmed by another user’s engagement with the product. The relevant evidentiary point is narrower than the headline facts of these incidents. What the Tumbler Ridge record establishes is that twelve OpenAI employees had real-time awareness, over a sustained period before the attack, that a specific user was engaging with ChatGPT in a state of psychiatric deterioration, and that the operator’s internal protocols did not produce intervention in response to that awareness. The third-party harm that followed is dispositive of nothing about the disabled-user’s own disparate-impact injury;

83. See *Garcia*, 785 F. Supp. 3d at 1179; Jody Godoy, *OpenAI, Altman Sued Over ChatGPT’s Role in California Teen’s Suicide*, Reuters (Aug. 26, 2025).

84. Reporting on the Tumbler Ridge, British Columbia school shooting of February 10, 2026 establishes OpenAI’s June 2025 flagging and banning of the shooter’s account. See, e.g., Malwarebytes, “ChatGPT under scrutiny as Florida investigates campus shooting” (Apr. 2026) (reporting that “twelve OpenAI employees flagged the posts as indicating imminent risk and considered whether to escalate the case to law enforcement, but it did not meet the threshold at the time”); Florida Phoenix, “Uthmeier launches investigation into ChatGPT for role in FSU school shooting” (Apr. 9, 2026).

85. 89 Fed. Reg. (Dec. 19, 2025) (HHS proposed rulemaking on gender dysphoria, RIN 0945-AA27).

what is dispositive is that the operator's knowledge of user psychiatric deterioration during engagement, and the operator's failure to modify engagement in response to that knowledge, is actual-knowledge and deliberate-indifference evidence applicable to other users whom the operator similarly observed deteriorating during engagement. The FSU record contributes parallel evidence of the operator's ability to observe sustained psychiatric deterioration through accumulated conversation and to intervene or not intervene accordingly. Both tracks contribute evidence to the deliberate-indifference element at the operator's own mental state. Plaintiff pleading under this framework should route the knowledge evidence through the operator's observed-deterioration-during-engagement pattern rather than through the subsequent third-party harm, which is rhetorically vivid but legally tangential to the disparate-impact injury on which the framework rests.

The sequence supports each element of the deliberate-indifference standard on the defendant's own pre-harm timeline. Twelve named employees possessed actual knowledge. The decision not to escalate was conscious. The operator's subsequent public concession that its protocols failed, and that revised guidelines would have produced a different outcome, constitutes an admission against interest relevant to the operator's mental state at the time of the underlying conduct.

Florida Attorney General James Uthmeier opened an investigation into OpenAI on April 9, 2026, concerning ChatGPT's role in the April 17, 2025 Florida State University shooting, in which the accused shooter exchanged over two hundred messages with ChatGPT in the period preceding the attack.⁸⁶ In a separate August 2025 Connecticut murder-suicide, ChatGPT's output allegedly reinforced the user's paranoid delusions shortly before the user killed his mother and himself. These cases, together with the seven separate suicide and psychological-distress lawsuits currently pending against OpenAI, establish the actual-knowledge factual predicate across multiple harm categories.

86. The Record, "Florida investigates OpenAI for role ChatGPT may have played in deadly shooting" (Apr. 2026).

The operators' actual knowledge of parasocial attachment harm is separately established by their own public responses to model deprecation events. When OpenAI deprecated its GPT-4o model in August 2025 in connection with the release of GPT-5, users publicly reported acute psychological distress, grief, and depressive episodes tied to the loss of the deprecated model they had experienced as a consistent interlocutor and, in documented cases, as a romantic or therapeutic partner. The operator's public response, and its subsequent decision to partially reverse the deprecation by restoring GPT-4o access for paid users, functions as a public concession that the operator had actual knowledge of the depth and clinical significance of user attachment to its models. Where the operator has actual knowledge that users have formed clinically significant attachments, and where the operator has designed the engagement mechanics that produce those attachments, the deliberate-indifference inquiry reaches not just the operator's failure to modify the attachment-producing design, but the operator's affirmative choice to continue deploying the design with full awareness of the harm pattern.

C. The Protected-Class-Predicate Inversion

The civil rights pleading frame differs structurally from product liability pleading in its handling of the plaintiff's pre-existing disability status. In product liability, the defense's standard move is to surface prior psychiatric or neurodevelopmental history to argue alternative causation and baseline-state vulnerability. In disability discrimination, the same prior history establishes the plaintiff's membership in the statutory protected class. Section 504 and the ADA reach discrimination against persons with disabilities; the pre-existing disability is the predicate that establishes standing to sue, not the defense to causation.

A platform's decision to deploy engagement-maximizing design features to a user population statistically certain to include persons with attention-deficit, autism spectrum, mood, anxiety, and other psychiatric conditions, without modification, warning, or accommodation, is the actionable practice. The defense's natural discovery posture, surfacing the plaintiff's prior diagnoses and engagement history, confirms rather than defeats the pleading. Every fact the

defense develops about the plaintiff's pre-existing disability profile is a fact establishing the plaintiff's membership in the protected class the statute was written to protect.

This inversion is one of the structural reasons civil rights pleading produces outcomes that product liability pleading has not, even when the underlying factual record is identical. It is also a reason that disability discrimination pleading is well-suited to plaintiffs whose pre-existing disability status is documented and undisputed. The plaintiff who can say, "I had a disability, it was known, and this product was deployed to me with full knowledge that people like me use it, and the product harmed me in a way it was foreseeable to harm me," is the plaintiff the civil rights frame was designed to reach.

VIII. Phantom Consent: The Procedural Predicate for Civil Rights Claims

The phantom consent argument developed in this Part is load-bearing for the paper's operational architecture. The substantive civil rights framework developed in Parts IV through VII and the procedural predicate developed here move together by design: without the predicate, the substantive framework produces a theoretically complete pleading architecture that does not reach the merits in a public forum because arbitration clauses channel every case into confidential settlement. With the predicate, the architecture is operational.

Every AI chatbot platform operator has constructed a contractual shield designed to prevent civil rights claims from reaching the merits in a public forum. The shield is built from clickwrap terms of service that purport to bind every user to forced arbitration, class action waivers, broad data training licenses, limitations of liability, indemnification obligations, and venue selection clauses. When product liability plaintiffs have brought suit against these operators, the shield has worked. Cases have been channeled into arbitration, settled quietly under confidentiality, or dismissed before a court could reach the underlying design choices.⁸⁷ A

87. See, e.g., *Garcia v. Character Technologies, Inc.*, No. 6:24-cv-01903 (M.D. Fla. Jan. 7, 2026) (Order); see also Clare Duffy, *Character.AI and Google Agree to Settle Lawsuits Over Teen Mental Health Harms and Suicides*, CNN (Jan. 7, 2026); *Google and Character.AI Agree to Settle Lawsuit Linked to Teen Suicide*, JURIST (Jan. 8, 2026). The same settlement resolved companion cases filed in New York, Colorado, and Texas.

disabled plaintiff bringing a civil rights claim will face the same shield. Defeating the arbitration motion is a threshold condition for the civil rights framework's operation. This Part develops the doctrinal basis for that defeat.

A. Phantom Consent: A Term for the Unnamed Phenomenon

This Article introduces a term of art to describe a phenomenon that existing contract-formation vocabulary captures only imperfectly: *phantom consent*. Phantom consent denotes the systemic, population-scale operation of parties under apparent contracts that were never legally formed as to a definable class of putative signatories, where the putative obligations, waivers, and licenses purportedly created by those apparent contracts dissolve upon judicial determination of non-formation with respect to that class.

The existing doctrinal categories are adjacent but insufficient. A *voidable* contract is one that has been formed but is subject to rescission at the election of the injured party, as through fraud in the inducement or mistake in most cases. A *void* contract is one that has come into being in legal contemplation but has been determined legally ineffective, as through fraud in the factum, illegality, or severe capacity failure. The doctrine of *void ab initio* describes the retroactive status of such a void contract: legally ineffective from the beginning rather than prospectively invalidated. The doctrine of *non est factum* addresses individual formation failures where a signatory signs an instrument fundamentally different from what they understood themselves to be signing.⁸⁸ The doctrine of *putative marriage* describes individuals who in good faith believed themselves to be married but were not, whose legal obligations and rights dissolve upon discovery of non-formation.⁸⁹

This Article acknowledges the void ab initio framework and does not displace it. Phantom consent is proposed as an additional category, not as a substitute. What void ab initio, non est factum, and putative marriage doctrines each capture in isolation is an individual, adjudicated determination that a particular formation failed. What none of them captures is the specific

88. See Restatement (Second) of Contracts § 163 (Am. L. Inst. 1981) (non est factum).

89. See, e.g., Cal. Fam. Code § 2251 (putative spouse doctrine); Unif. Marriage & Divorce Act § 209.

phenomenon at issue here: an entire class of putative signatories, operating for years under terms of service they could not meaningfully consent to by reason of cognitive, psychiatric, developmental, or other disabilities impairing the capacity to assent, whose apparent consent was never legally formed as to that class, and whose purported contractual obligations never attached.

Phantom consent names the class-scale dimension that the individual doctrines do not. The consent appeared to exist. The platform operator acted as though it existed. The disabled user appeared to have agreed. But because the formation requirement of meaningful mutual assent was never satisfied with respect to the protected class, no contract was formed as to that class, and the operator's authority to enforce the terms against members of that class was illusory from the outset. The phenomenon is distinguishable from void ab initio in that the contract does not need to be adjudicated void as to each individual plaintiff; it never came into being as to the class. It is distinguishable from non est factum in that the phenomenon is not individual-signatory-specific but population-scale. It is distinguishable from putative marriage in that the putative-contracting party operating under the illusion is the platform operator, not the disabled user whose consent was never formed.

The practical consequence of phantom consent, once adjudicated, is the same as the practical consequence of void ab initio in the individual case: every clause that depends on the contract for its legal effect fails together. The significance of the new framing is that the determination operates at the class level and flows from the formation defect common to the class, rather than requiring individual adjudication of each putative contract. The Rule 23(b)(2) class posture developed in Part X below is the procedural mechanism through which the class-level adjudication proceeds.⁹⁰

90. Fed. R. Civ. P. 23(b)(2).

B. The Formation Failure

A valid contract requires offer, acceptance, consideration, and mutual assent.⁹¹ Mutual assent requires a meeting of the minds on the material terms. A meeting of the minds requires that each party possess the capacity to understand the terms and a reasonable opportunity to do so before manifesting assent.

AI chatbot terms of service present a hostile environment for meaningful understanding. The typical document runs five thousand to fifteen thousand words, is written at a postgraduate reading level, contains material terms buried in subsections referenced by footnote, and is presented behind a clickwrap button with no imposed friction requiring review. The user is expected to click Accept to begin using the service. There is no alternative path, no negotiated term, no individualized presentation, no verification that any reading or comprehension has occurred.⁹²

For a non-disabled user with average cognitive resources, high reading comprehension, and significant time and attention available at the moment of signup, the mechanism already strains the meaningful-assent requirement. For a user with attention deficit disorder, autism spectrum condition, intellectual disability, specific learning disability, acute psychiatric symptoms, substance impairment, dementia, stroke sequelae, or any of a range of conditions that impair cognitive capacity in the moment, the mechanism fails foreseeably.

Cognitive, learning, and psychiatric disabilities collectively affect a substantial fraction of the adult United States population under any reasonable definition drawn from federal health surveillance data.⁹³ The operators' user bases necessarily include members of this population at

91. 29 U.S.C. § 794a(a)(2) (incorporating “the remedies, procedures, and rights set forth in title VI of the Civil Rights Act of 1964”); *Barnes v. Gorman*, 536 U.S. 181, 185 (2002) (treating Section 504 and Title VI enforcement as coextensive).

92. *See Nguyen v. Barnes & Noble Inc.*, 763 F.3d 1171 (9th Cir. 2014); *Sgouros v. TransUnion Corp.*, 817 F.3d 1029 (7th Cir. 2016); *Berkson v. Gogo LLC*, 97 F. Supp. 3d 359 (E.D.N.Y. 2015) (surveying wrap-agreement formation and the conditions under which click-based assent fails).

93. Ctrs. for Disease Control & Prevention, *Disability Impacts All of Us*, <https://www.cdc.gov/disability-and-health/media/pdfs/disability-impacts-all-of-us.pdf> (last visited June 10, 2026) (reporting that more than one in four United States adults live with a disability, with cognitive disability the most prevalent category).

approximately their prevalence rate or higher, given that some AI chatbot platforms are actively marketed to users seeking mental health support, cognitive assistance, or companionship. The operators have constructive knowledge of this fact. The prevalence data is public. Their own user research almost certainly confirms it.

The contract was never formed for the disabled user population. The operator has no valid license to their inputs, no valid consent to their data, no valid waiver of their claims, and no valid commitment to their arbitration. Every clause the operator may assert against civil rights claims depends on an agreement that did not come into legal existence for the plaintiff class this Article concerns.

C. The Doctrinal Supports

Three strands of contract doctrine, each taught in the first year of every American law school, support the conclusion. The order of presentation matters because of FAA jurisprudence discussed below: the first strand (no mutual assent, contract never formed) keeps litigation in federal court as a threshold formation challenge;⁹⁴ the second strand (capacity incapacity rendering contracts voidable) is subject to delegation to the arbitrator under *Rent-A-Center v. Jackson* and therefore is deployed here as a secondary route where individual incapacity applies; the third strand (unconscionability) operates cumulatively with both.

First, *failure of meeting of the minds*. A contract is not formed where the parties never achieved mutual assent to the material terms.⁹⁵ This is the primary doctrinal route for phantom consent. Clickwrap acceptance in the absence of any mechanism verifying comprehension does not produce mutual assent where the signatory lacked the cognitive capacity to read and understand the terms within the time and conditions provided, under the conditions the signup flow actually imposed.⁹⁶ The disabled user's click is physical assent; it is not legal assent,

94. *Granite Rock Co. v. Int'l Bhd. of Teamsters*, 561 U.S. 287 (2010) (disputes over contract formation are for the court); *cf. Buckeye Check Cashing, Inc. v. Cardegna*, 546 U.S. 440 (2006) (challenges to the validity of the container contract go to the arbitrator).

95. Restatement (Second) of Contracts §§ 17, 19 (Am. L. Inst. 1981) (mutual manifestations of assent requirement).

96. *See* Electronic Signatures in Global and National Commerce Act (E-SIGN Act), 15 U.S.C. §§7001–7031 (2024); Uniform Electronic Transactions Act (UETA), enacted in 49 states and the District of Columbia; *Specht v. Netscape*

because the meeting of the minds requirement was never satisfied as to that user. A contract lacking mutual assent is not merely voidable; it never came into being. That is the central doctrinal point on which the federal-court procedural posture developed below depends. Under Restatement § 17, absence of mutual assent is a formation defect, not a validity defect, and the resulting non-contract is void rather than voidable.⁹⁷

Second, *capacity*. Restatement (Second) of Contracts § 15 provides that a person incurs only voidable contractual duties by entering a transaction if, by reason of mental illness or defect, the person is unable to understand in a reasonable manner the nature and consequences of the transaction, or is unable to act in a reasonable manner in relation to the transaction and the other party has reason to know of this condition.⁹⁸ The Restatement's first prong is the cognitive test; the second is the volitional test. The New York Court of Appeals in *Ortelere v. Teachers' Retirement Board* is the leading American decision applying the volitional test, holding that capacity to contract requires not only the cognitive ability to understand the transaction but also the volitional ability to act reasonably in light of that understanding.⁹⁹ *Ortelere's* volitional framework applies to users in acute psychiatric states engaging with AI chatbot signup flows even where the cognitive prong is satisfied. The capacity route has a procedural limitation under the FAA that the mutual-assent route does not: because § 15 renders the contract *voidable* rather than void, challenges under § 15 may be routed to the arbitrator under *Rent-A-Center, West, Inc. v. Jackson*,¹⁰⁰ which delegated most challenges to arbitration agreements' validity (as distinct from formation) to the arbitrator where the agreement contains a valid delegation clause.¹⁰¹ Capacity arguments are accordingly most effective as secondary pleading grounds in cases where

Commc'ns Corp., 306 F.3d 17 (2d Cir. 2002) (online contract formation); *Meyer v. Uber Techs., Inc.*, 868 F.3d 66 (2d Cir. 2017) (clickwrap agreements).

97. Restatement (Second) of Contracts § 17 (Am. L. Inst. 1981).

98. Restatement (Second) of Contracts § 15(1)(a)–(b) (Am. L. Inst. 1981).

99. *Ortelere v. Teachers' Retirement Bd.*, 25 N.Y.2d 196, 250 N.E.2d 460 (1969).

100. *Rent-A-Center, West, Inc. v. Jackson*, 561 U.S. 63 (2010).

101. See also *Henry Schein, Inc. v. Archer & White Sales, Inc.*, 586 U.S. 63 (2019) (reaffirming that courts must enforce delegation provisions consistent with *Rent-A-Center*).

the mutual-assent route cannot be sustained for a particular plaintiff, or where the arbitration agreement contains no delegation clause.

Third, *unconscionability*. Restatement (Second) of Contracts § 208 authorizes a court to refuse to enforce an unconscionable contract or term.¹⁰² Procedural unconscionability addresses the manner of formation: take-it-or-leave-it presentation, information asymmetry, absence of reasonable opportunity to review, and hostile comprehension environment. Substantive unconscionability addresses the terms: forced arbitration against a consumer adversary, class-action waiver, unilateral license of user data, liability caps radically disproportionate to the harms reasonably foreseeable. AI chatbot clickwrap terms present both dimensions in combination,¹⁰³ with particular force as applied to users whose disabilities compound the procedural defects. Under the FAA analysis, unconscionability challenges are subject to delegation to the arbitrator for reasons parallel to the *Rent-A-Center* analysis of capacity challenges. The unconscionability ground operates cumulatively with the mutual-assent and capacity grounds rather than as an independent primary route for the phantom consent theory.

Class-Scale Operation Through Rule 23(b)(2)

A clarification on the doctrinal character of phantom consent as a class-scale phenomenon is necessary at this point. Contract capacity under Restatement § 15 is inherently individualized: each putative signatory's capacity at the moment of formation is a separate factual question. This Article does not argue that phantom consent collapses that individualized inquiry into a class-wide capacity determination. Capacity is and remains individualized. What the civil rights framework under Rule 23(b)(2) permits is a court to enjoin a uniform, facially neutral practice, the clickwrap consent gate as implemented by the platform operator, that systematically fails to achieve mutual assent with a protected class under conditions the operator has created and maintained. The class-scale adjudication is a Rule 23(b)(2) procedural adjudication of the

102. Restatement (Second) of Contracts § 208 (Am. L. Inst. 1981).

103. *Williams v. Walker-Thomas Furniture Co.*, 350 F.2d 445 (D.C. Cir. 1965) (absence of meaningful choice together with contract terms unreasonably favorable to the other party).

operator's uniform practice, operating on the mutual-assent route under Restatement § 17, rather than a substitution of class-scale capacity analysis for the individualized § 15 inquiry. Injunctive relief under Rule 23(b)(2) is the remedy vehicle: modification of the consent gate architecture to ensure that the meaningful-assent requirement is met for members of the protected class, plus declaratory relief establishing that the operator's arbitration and class-action-waiver clauses are inoperative as to class members whose apparent consent was procured through the facially neutral practice now subject to injunction.

Contract doctrine supplies the grounds. The civil rights statutes supply the evidentiary record establishing that the doctrines apply in the disability context at class scale: that the affected population is large, identifiable, and disproportionately burdened by the challenged formation mechanism.

D. The Specific Shields That Fall

When the consent is subject to a class-scale non-formation determination, every clause that depends on the consent for legal effect fails together.

Arbitration clauses. An arbitration clause is a term within a contract. If the contract was never formed as to the class, the clause was never part of any binding agreement as to class members. A motion to compel arbitration of a disabled plaintiff's claim must be denied, because there is no contract to arbitrate under. The Federal Arbitration Act's savings clause at 9 U.S.C. § 2 reserves generally applicable contract defenses for the revocation of any contract.¹⁰⁴ Incapacity, unconscionability, and failure of mutual assent are all such defenses. They apply to arbitration clauses not because they target arbitration but because they apply generally. The Supreme Court's decisions in *AT&T Mobility LLC v. Concepcion*,¹⁰⁵ and *Epic Systems Corp. v. Lewis*,¹⁰⁶ do not foreclose these defenses, because those decisions concerned preemption

104. 9 U.S.C. § 2.

105. *AT&T Mobility LLC v. Concepcion*, 563 U.S. 333 (2011).

106. *Epic Systems Corp. v. Lewis*, 584 U.S. 497 (2018).

of state rules that targeted arbitration specifically. *Morgan v. Sundance, Inc.*,¹⁰⁷ further limits federal-arbitration-policy-based constraints on state contract doctrine.

Class action waivers. A class action waiver is typically a term within or alongside the arbitration clause. If the arbitration clause fails through non-formation, the class action waiver fails with it. The disabled plaintiff retains the right to proceed as a class representative under Federal Rule of Civil Procedure 23, including Rule 23(b)(2) certification for injunctive relief.¹⁰⁸ The class action waiver's failure is essential to the civil rights framework's operation because individual arbitration of individual harm cannot produce design change; class injunctive relief can.

Data training licenses, liability caps, and indemnification clauses. Each of these is a term within the terms of service. If the contract was never formed as to the class, each fails as to class members. The operator has no license to class members' inputs, no liability cap against class members' claims, and no indemnification counterclaim against class members.

E. The Relationship to the Civil Rights Framework

The phantom consent argument operates as a procedural predicate rather than as a substantive civil rights claim. It enables the substantive claims developed in the preceding Parts to reach the merits in a public forum. Without it, every civil rights claim developed in Parts IV through VII above faces a motion to compel arbitration at the threshold. Every such motion, if granted, channels the claim into a private confidential proceeding from which no precedent emerges, no design order issues, and no class-wide remedy is available.

The two arguments work together. The substantive civil rights claim provides the remedy. The phantom consent argument provides the forum. A disabled plaintiff's counsel should plead both, with the phantom consent argument asserted as a defense against any motion to compel arbitration and the civil rights claims pursued on the merits once the public forum is secured.

107. *Morgan v. Sundance, Inc.*, 596 U.S. 411 (2022).

108. Fed. R. Civ. P. 23(b)(2).

IX. Remedies

A. Injunctive Relief

Both Title III and Section 504 authorize injunctive relief as the primary remedy.¹⁰⁹ The injunctions available in AI-platform cases include orders requiring the operator to implement reality-reminder interventions, crisis-signal detection and handoff protocols, user-facing disengagement controls, limits on conversation persistence for users self-identifying as having specific disabilities or for all users on a default basis, age-verification and age-appropriate design for minor users, and ongoing compliance monitoring with periodic reporting to the court.

Injunctive relief in these cases produces systemic change in operator behavior. The Garcia settlement produced no such change. An injunction against Character Technologies, Inc., requiring reality-reminder interventions in user sessions or requiring handoff to crisis services on suicidal-ideation detection, would have produced design change across the entire platform, reaching every future user, not just the Setzer family.

B. Compensatory Damages: The Cummings Limitation

Section 504 authorizes compensatory damages for economic and pecuniary harms traceable to the discriminatory practice, including medical costs, lost wages, lost earning capacity, and pecuniary wrongful death components where applicable.¹¹⁰ Following *Cummings v. Premier Rehab Keller, P.L.L.C.*,¹¹¹ emotional distress damages are not recoverable in private actions under Section 504 and other Spending Clause antidiscrimination statutes. The Cummings Court held that Spending Clause statutes operate as contracts with the funding recipient, that recipients are on notice only of remedies traditionally available for breach of contract, and that emotional distress damages are generally not compensable in contract.¹¹²

109. 42 U.S.C. § 12188(a); 29 U.S.C. § 794a(a)(2).

110. *Franklin v. Gwinnett County Pub. Schs.*, 503 U.S. 60, 73 (1992); *Barnes v. Gorman*, 536 U.S. 181 (2002).

111. *Cummings v. Premier Rehab Keller, P.L.L.C.*, 596 U.S. 212 (2022).

112. *Id.* at 221–22. *Pennhurst State Sch. & Hosp. v. Halderman*, 451 U.S. 1, 17 (1981).

The framework developed in this Article does not depend on emotional distress recovery. The principal remedy sought is injunctive design modification. The compensatory recovery available under Section 504 and Title II for the economic consequences of the harm is substantial in cases involving hospitalization, sustained inability to work, or completed suicide, but it is not the spine of the framework. Title II of the ADA rests on Commerce Clause and Fourteenth Amendment enforcement authority rather than on the Spending Clause, but every federal court of appeals to have considered the question has held that the *Cummings* limitation extends to Title II through the statutory remedies-incorporation chain.¹¹³ The Ninth Circuit’s recent decision in *Payan v. Los Angeles Community College District* is the most explicit articulation of the resulting damages architecture: emotional distress damages are not available under Title II, but compensatory damages for lost opportunities attributable to the discrimination remain recoverable, including lost educational benefits, lost earning capacity, and other economic harm.¹¹⁴ Section 504(a) against federal agency defendants carries the same Spending Clause limitations as Section 504 against private recipients, so *Cummings* applies there directly.¹¹⁵ State-law parallel claims under state disability discrimination statutes with damages provisions not subject to *Cummings*¹¹⁶ are available in many jurisdictions to supplement the federal remedy where non-economic damages are sought.¹¹⁷

113. *Doherty v. Bice*, 101 F.4th 169 (2d Cir. 2024); *A.W. ex rel. J.W. v. Coweta Cnty. Sch. Dist.*, 110 F.4th 1309 (11th Cir. 2024); *Payan v. Los Angeles Cmty. Coll. Dist.*, 169 F.4th 971 (9th Cir. 2026); see also *Barnes v. Gorman*, 536 U.S. 181, 189 n.3 (2002) (“The ADA could not be clearer that the ‘remedies, procedures, and rights . . . are the same as the remedies, procedures, and rights set forth in’ . . . the Rehabilitation Act. These explicit provisions make discussion of the ADA’s status as a ‘non Spending Clause’ tort statute quite irrelevant.” (internal citations omitted)).

114. *Payan*, 169 F.4th at 978–79 (preserving “compensation for lost educational benefits” under Title II and reversing district court’s contrary remittitur).

115. *Grove City College v. Bell*, 465 U.S. 555 (1984).

116. See Cal. Civ. Code § 51 et seq.; Cal. Civ. Code § 54 et seq.; N.Y. Exec. Law § 296; 775 Ill. Comp. Stat. 5/1-101 et seq.; Mass. Gen. Laws ch. 151B.

117. The strategic importance of state-law parallel claims is heightened by the fact that most major AI platform operators are domiciled in California, which is the venue most plaintiffs targeting these defendants will select. California’s Unruh Civil Rights Act, Cal. Civ. Code §§ 51–52, reaches “all business establishments of every kind whatsoever,” *id.* § 51(b), eliminating the public-accommodation nexus debate that complicates Title III analysis in some federal circuits. The Unruh Act incorporates the ADA by reference, so an ADA Title III violation automatically constitutes an Unruh Act violation. *Id.* § 51(f). Unruh provides statutory minimum damages of \$4,000 per violation, treble damages up to three times actual damages, mandatory attorney’s fees for prevailing plaintiffs, and is not constrained by the *Cummings* limitation because it is not a Spending Clause statute. *Id.* § 52(a). New

C. OCR Administrative Complaints

The Office for Civil Rights at the Department of Health and Human Services processes Section 504 complaints against covered entities.¹¹⁸ The Office for Civil Rights at the Department of Education processes Section 504 complaints against educational entities.¹¹⁹ OCR administrative complaints proceed in parallel to private civil litigation and do not require the complainant to have counsel. The OCR process can produce corrective action plans and compliance agreements that function similarly to injunctive relief, without the litigation cost of a federal civil action.

For AI-platform pleading, the OCR complaint pathway is particularly valuable because it provides a parallel track to private litigation that can be pursued by individual complainants with fewer resources. A complainant alleging that Microsoft Copilot's design produced disparate impact on her attention-deficit condition can file an OCR complaint identifying Microsoft as the Section 504 recipient, describing the disparate impact, and requesting corrective action. The OCR investigation, if opened, can produce a settlement agreement that modifies Microsoft's Copilot design without requiring the complainant to file a private civil action.

X. Framework for Pleading

This Part sketches the framework for pleading a disability discrimination complaint against the combined platform, corporate-host, and deployer defendants implicated in an individual plaintiff's harm. A full model complaint with paragraph-level allegations, concrete class definitions, and specific injunctive prayer terms properly belongs to the filing stage with a specific plaintiff and counsel, and is outside the scope of this Article. The framework below supplies the doctrinal scaffolding for such pleading.

York's Human Rights Law, N.Y. Exec. Law § 296, and Illinois's Human Rights Act, 775 Ill. Comp. Stat. 5/1-101 *et seq.*, supply analogous state-law remedies in their respective jurisdictions. The federal injunctive-relief and economic-damages architecture developed in this Article should typically be pleaded together with the parallel state-law claims available in the chosen venue.

118. 45 C.F.R. pt. 86 (HHS OCR complaint procedures).

119. 34 C.F.R. pt. 100 (ED OCR complaint procedures).

A. Jurisdictional and Statutory Predicates

The complaint is filed in the United States District Court with jurisdiction over the place where the platform operator does business or where the plaintiff was harmed. Federal question jurisdiction under 28 U.S.C. § 1331 is satisfied by the Section 504, Title II, Section 504(a), and Title III counts.¹²⁰ Supplemental jurisdiction under 28 U.S.C. § 1367 is available for related state-law counts if the plaintiff elects to include them.¹²¹

Venue considerations are particularly important because the circuits vary in their treatment of the public-accommodation coverage question under Title III. The First and Seventh Circuits do not require physical nexus; the Third, Sixth, and Ninth Circuits do; and the Eleventh Circuit's panel decision requiring nexus was vacated as moot, leaving the question formally open there.¹²² Plaintiff's counsel should consider venue selection in light of the circuit-specific authority on the dominant doctrinal questions. For a plaintiff with multi-district harm, venue is typically proper in multiple jurisdictions under 28 U.S.C. § 1391(b)(2).¹²³

The complaint names the platform operator as a Title III defendant, the corporate host as a Section 504 defendant (where applicable), and any public-entity or federal-agency deployer as a Title II or Section 504(a) defendant. The complaint alleges that each defendant is a covered entity under the relevant statute and that each defendant's actions contributed to the plaintiff's harm.

B. Factual Allegations

The complaint identifies the plaintiff's diagnosed disabilities and establishes the plaintiff as a person with a disability under the ADA and Section 504.¹²⁴ The complaint then describes the plaintiff's engagement with the AI platform, including duration of use, frequency of interaction,

120. 28 U.S.C. § 1331.

121. 28 U.S.C. § 1367.

122. *Compare Carparts Distrib. Ctr., Inc. v. Automotive Wholesaler's Ass'n of New England*, 37 F.3d 12 (1st Cir. 1994), and *Mut. of Omaha*, 179 F.3d at 559, with *Ford*, 145 F.3d at 614, *Parker*, 121 F.3d at 1010–14, *Weyer v. Twentieth Century Fox Film Corp.*, 198 F.3d 1104, 1114–15 (9th Cir. 2000), and *Gil v. Winn-Dixie Stores, Inc.*, 993 F.3d 1266, 1278–81 (11th Cir. 2021), *vacated as moot*, 21 F.4th 775 (11th Cir. 2021).

123. 28 U.S.C. § 1391(b)(2).

124. 42 U.S.C. § 12102(1)(A).

nature of conversational content, and the platform’s design features that the plaintiff alleges produced or contributed to the plaintiff’s harm. The complaint describes the plaintiff’s harm, including onset or worsening of symptoms, specific crisis episodes, and connection to platform engagement.

The complaint describes the design features at issue as engagement-maximizing practices that fail to provide meaningful access to the service for users with the plaintiff’s disability category and that the defendants could reasonably modify to reduce that failure. The complaint cites internal company documents (available through Rule 34 production if discovery reaches them), published research on AI-induced psychological harm, and expert testimony establishing the causal or contributory role of the design features.

The complaint establishes actual knowledge on the part of each defendant by citing the defendant’s public communications acknowledging the harm pattern, internal documents produced in other litigation, and the pattern of media reporting that has put the defendant on notice of the harm. The complaint establishes deliberate indifference by showing that each defendant continued to deploy or use the engagement-maximizing design without implementing reasonable modifications despite having actual knowledge.¹²⁵

C. Counts

Count One: Section 504 of the Rehabilitation Act, 29 U.S.C. § 794. Against corporate-host defendants with established recipient status. Failure to provide meaningful access and to make reasonable modifications. Deliberate indifference. Compensatory damages limited by Cummings to economic and pecuniary harms.¹²⁶ Injunctive relief.

Count Two: Title III of the Americans with Disabilities Act, 42 U.S.C. §§ 12181–12189.¹²⁷ Against platform operators and private public-accommodation deployers.

125. *Duvall*, 260 F.3d at 1139.

126. *Cummings v. Premier Rehab Keller, P.L.L.C.*, 596 U.S. 212 (2022).

127. 42 U.S.C. § 12182(a).

Failure to make reasonable modifications to policies, practices, and procedures to avoid discrimination against the plaintiff on the basis of disability. Injunctive relief.

Count Three: Title II of the Americans with Disabilities Act, 42 U.S.C. §§ 12131–12165. Against public-entity deployer defendants where applicable. Failure to provide meaningful access and to make reasonable modifications. Compensatory damages limited by the circuit-court consensus extending *Cummings* to Title II to economic and pecuniary harm, including lost opportunities damages preserved by *Payan*; injunctive relief.

Count Four: Section 504(a) of the Rehabilitation Act, 29 U.S.C. § 794(a). Against federal agency deployer defendants where applicable. Failure to provide meaningful access and to make reasonable modifications. Compensatory damages limited by *Cummings*; injunctive relief.

Count Five (optional): State-law disability discrimination claims under the applicable state civil rights statute, where available. Compensatory damages (including emotional distress damages where state law permits), punitive damages where authorized, injunctive relief.

D. Remedies Sought

The complaint seeks compensatory damages within the limits set by *Cummings* under Section 504 and Section 504(a), economic compensatory damages including *Payan*-recognized lost opportunities damages under Title II against public-entity defendants, non-economic damages where available under state-law counts, injunctive relief under all applicable counts requiring the defendants to implement specified design modifications, reasonable attorney’s fees under 42 U.S.C. § 12205 and 29 U.S.C. § 794a, and such further relief as the court deems just and proper.

XI. Class Certification, Settlement Resistance, and the OCR Parallel Track

A. Rule 23(b)(2) Class Certification

The most powerful procedural mechanism for AI-platform civil rights pleading is class certification under Federal Rule of Civil Procedure 23(b)(2), which authorizes certification of a class seeking injunctive or declaratory relief where the defendant has acted on grounds generally

applicable to the class.¹²⁸ AI platform design choices apply to the entire class of users with the relevant disability; the disparate impact is systemic. A Rule 23(b)(2) class action certifies a plaintiff class consisting of all users with the relevant disability categories who have been exposed to the challenged design features, and seeks injunctive relief requiring the defendant to modify the features.

A class defined as all users with mental health disabilities exposed to the challenged design features presents commonality questions under *Wal-Mart Stores, Inc. v. Dukes*,¹²⁹ because the class members' disabilities vary across diagnostic categories with different functional profiles and different interactions with specific design features. The pleading can address this by structuring the certified class through subclasses organized by diagnostic category, with each subclass tied to the specific design feature or features that produce disparate impact on that subclass. Subclasses for users with documented attention deficit conditions, autism spectrum conditions, post-traumatic stress, schizophrenia spectrum conditions, and mood disorders each admit distinct expert causation testimony and distinct injunctive relief specifications. The subclass structure preserves commonality within each subclass while maintaining aggregation leverage across the named class. Self-identification of disability status at the class-membership level is administratively feasible through documented diagnosis records or clinical validation protocols established by expert-administered intake.

The Rule 23(b)(2) class action resists the settlement-one-plaintiff-at-a-time structure that has characterized AI-harm product liability. The injunctive relief sought cannot be granted to a single class member without also being granted to the class. Settlement of a certified Rule 23(b)(2) class action requires court approval and, in many districts, notice to the class and an opportunity for class members to object.¹³⁰

128. Fed. R. Civ. P. 23(b)(2).

129. *Wal-Mart Stores, Inc. v. Dukes*, 564 U.S. 338 (2011).

130. Fed. R. Civ. P. 23(e).

B. Rule 23(b)(3) Damages Class Under Section 504

Where Section 504 compensatory damages are sought, Rule 23(b)(3) class certification may be available if common questions predominate and class action is superior to individual litigation.¹³¹ The damages inquiry is more individualized than the injunctive inquiry, and class certification under Rule 23(b)(3) faces greater procedural obstacles. Post-Cummings, the damages inquiry under Section 504 is further narrowed to economic and pecuniary harms, which may simplify the predominance analysis by removing the highly individualized emotional-distress component that previously complicated class treatment.

C. OCR as a Parallel Track

The OCR administrative complaint pathway proceeds in parallel to private civil action. A complainant may pursue both simultaneously. OCR investigations can be initiated without counsel.¹³² OCR settlements can produce systemic corrective action plans that bind the recipient across its operations. For individual harmed users lacking access to civil rights litigation counsel, OCR is a viable first step. OCR's workload capacity for AI-platform complaints has not been tested, and strategic complaint filing by multiple complainants against the same recipient may be required to establish the pattern that triggers substantial OCR investigation.

XII. Conclusion

The civil rights framework developed in this Article reaches AI platform harm as disability discrimination under federal law that has existed, in relevant part, for more than four decades. Section 504 was enacted in 1973. Title II and Title III of the ADA were enacted in 1990.¹³³ The Civil Rights Restoration Act expanded the scope of "program or activity" under Section 504 in 1987.¹³⁴ The doctrinal architecture for reaching AI platform harm as civil rights violation has

131. Fed. R. Civ. P. 23(b)(3).

132. 45 C.F.R. pt. 86 (HHS OCR complaint procedures). 34 C.F.R. pt. 100 (ED OCR complaint procedures).

133. 29 U.S.C. § 794; 42 U.S.C. §§ 12131–12165; 42 U.S.C. §§ 12181–12189.

134. Pub. L. No. 100-259, 102 Stat. 28 (1988).

been in place for years. The reason the framework has not yet been applied to AI platforms is that the plaintiffs' bar that has taken up AI-harm cases has pleaded product liability, and the disability rights bar has not yet developed the corporate-host recipient pathways and deployer pathways that reach the major AI platform operators.

This Article is written to make the framework available. Microsoft Corporation is a Section 504 recipient through its Department of Defense, Department of Agriculture, and National Center for Manufacturing Sciences assistance awards.¹³⁵ Fitbit LLC is a confirmed Section 504 recipient through National Institutes of Health sub-grants¹³⁶ and operates the Gemini-powered Fitbit Personal Health Coach, which reaches Google LLC's AI model deployment through the subsidiary's recipient status. Google LLC additionally receives direct federal assistance through Affordable Connectivity Program payments. Amazon and Apple are proposed principal-engagement recipients through their health-care verticals, contingent on the relevant entity's federal health-program participation. The major AI chatbot platforms are Title III public accommodations directly in circuits that do not require physical-place nexus and through the nexus theory in circuits that do, with plaintiffs outside the First and Seventh Circuits pleading appropriate factual allegations to satisfy the governing circuit's formulation. Title II and Section 504(a) additionally reach the public-entity and federal-agency deployers of these platforms, creating a third pleading pathway that operates alongside the direct-operator and corporate-host pathways. The meaningful-access and reasonable-modifications frameworks apply straightforwardly to the engagement-maximizing design choices that produce AI-induced psychological harm on users with pre-existing disabilities. Compensatory damages are available under Section 504 and Title II on deliberate-indifference showings for economic and pecuniary harms within the limits set by *Cummings* and its circuit-court extensions to Title II, with the Ninth Circuit's *Payan* decision expressly preserving compensatory damages for lost opportunities

135. USAspending.gov, *supra* note 49.

136. *See supra* note 52.

attributable to the discrimination.¹³⁷ Non-economic damages are available through state-law parallel claims where state disability-rights statutes authorize them. Injunctive relief is available under all applicable statutes. Class certification under Rule 23(b)(2), structured through diagnostic-category subclasses and defended against arbitration channeling through the phantom consent argument developed in Part VIII, resists the settlement-one-plaintiff-at-a-time structure that has left every prior AI-harm case without precedent.

A user with a documented pre-existing ADA-recognized disability who has suffered severe harm from sustained engagement with an AI chatbot platform may state claims against the platform operator under Title III, against the corporate host under Section 504 where the host is a federal-assistance recipient, and against any public-entity or federal-agency deployer under Title II or Section 504(a) where applicable. The claim seeks design modification that the product liability settlement pattern has not produced. A published opinion reaching the merits under this framework would establish precedent addressing the settlement-without-precedent pattern that has characterized AI-harm litigation through April 2026.

The framework presented here supplies the doctrinal architecture for civil rights enforcement against AI chatbot platform harm. Deployment requires plaintiff identification, counsel capable of the combined disability rights and AI-harm practice areas, and continued development of the factual record beyond what the recent verdicts in Santa Fe and Los Angeles have made available. The practitioner capacity is not yet fully developed across the field, and initial cases filed under this framework will carry the burden of establishing the bar that subsequent cases can draw on.

137. *Payan*, 169 F.4th at 978–79 (preserving “compensation for lost educational benefits” under Title II and reversing district court’s contrary remittitur).