

There Is No User in the Living Room: Ambient Household Devices and the User-Relative Architecture of the People-First Chatbot Act

Travis Gilly* · July 2026 (v0.4)

Abstract

The People-First Chatbot Act, introduced in the House on July 9, 2026, would impose privacy, transparency, and design duties on artificial intelligence chatbot providers. Every operative duty in the bill runs through a single term. A provider may process personal data only as necessary to fulfill an express request made by a user who has given affirmative consent. A provider must disable risky design features for a user it knows to be a minor. A provider incurs strict liability to a user injured through use. The Act then defines that term, at Section 5(22), as “an individual, regardless of age.” That is a genus without a differentia. It defines nothing.

The defect is invisible on a phone, where the person holding the device and the person the statute regulates are the same person. It becomes dispositive on an ambient household device, a category already shipping into American homes at scale and about to be entered by the largest chatbot provider in the country. A child who walks past a camera-equipped speaker has made no request, consented to nothing, and interacts with the system only in the sense that the system is processing her face. This Article shows that the Act cannot say whether she is a user, that both available answers defeat the statute’s purpose, and that a third defect compounds them: the Act creates a private right of action denominated in violations while never defining a violation. The resulting damages exposure for a single afternoon in a single household ranges across four orders of magnitude on a question the text does not answer.

The Article makes three contributions. It identifies user-relativity as a structural drafting pathology distinct from the overbreadth and vagueness critiques already directed at chatbot legislation, and distinguishes requester-keyed statutes from data-subject-keyed statutes using the Children’s Online Privacy Protection Act as the comparator. It demonstrates that the Act’s consent architecture permits one individual to authorize the processing of another individual’s biometric data, a result the Act’s own definition of affirmative consent was written to foreclose. It proposes three amendments, each drawn from existing federal law, that close the gap without disturbing the bill’s structure.

Keywords: People-First Chatbot Act; H.R. 9619; OpenAI; artificial intelligence; AI chatbots; ambient computing; smart speakers; children’s privacy; COPPA; biometric data; age estimation; statutory construction; statutory damages; private right of action; bystander privacy

*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.

Contents

I	Introduction	3
II	The User-Relative Architecture	7
	A Every Duty Has the Same Addressee	7
	B Why the Architecture Held Until Now	8
III	The Definition That Defines Nothing.	10
	A Section 5(22) Does One Job	10
	B Construction Must Proceed From Context, and Context Points Away From the Child . . .	11
	C The Definition Is Not Saved by the Inclusions Clause.	12
IV	The Fork.	12
	A Horn One: The Child Is a User	12
	B Horn Two: The Child Is Not a User.	13
	C Both Horns Defeat the Purpose.	16
V	The Necessity Limiter and Consent That Belongs to Someone Else.	17
	A The Limiter.	17
	B Whose Consent Is It.	18
	C The Visitor	18
VI	The Counting Problem	20
	A A Damages Provision Without a Unit	20
	B Three Constructions.	21
	C Four Orders of Magnitude	25
	D Standing, and the State-Court Answer.	26
VII	Whose Data It Is Versus Whose Request It Was	27
	A Two Ways to Key a Privacy Statute	27
	B The Counting Problem Is Downstream of the Keying Choice	28
	C The Comparison Runs the Other Way on Remedy	28
	D The Patchwork Outside the Act	30
VIII	Remedies Written for a Phone, Aimed at Furniture	31
IX	Three Amendments	33
X	Conclusion	34

I. Introduction

A speaker sits on a kitchen counter. It has a camera. Its microphone does not stop, because on the mature devices it must listen in order to hear its own name, and on the device now entering the category it is designed to listen and speak at the same time. It runs a large language model that responds in a human voice, in real time, and it is designed to learn the household it lives in.¹ A parent is at the sink. A child walks into the room.

At that moment the device captures the child's face. It captures her voice if she speaks. It transmits both to a server. It may infer her age from the pitch and cadence of that voice, because the models that power such systems can.² It may infer her emotional state from her expression, because the models that power such systems can do that too.³

The People-First Chatbot Act, introduced in the House of Representatives on July 9, 2026, is a bill about exactly this kind of processing.⁴ Its largest operative block is captioned Data Privacy and Security.⁵ It bans the processing of personal data outside narrow limits, bans advertising built on chat logs, bans training on a minor's data, and bans profiling that infers emotional vulnerability.⁶ It

1. Such devices are not hypothetical. Amazon's Echo Show, Google's Nest Hub, and Amazon's Ring line all combine ambient audio capture, camera hardware, and cloud processing in domestic settings, and all are shipping today. The category is mature. What is new is twofold. The first is the integration of a general-purpose conversational model, which converts an appliance that responds to commands into a system that sustains conversation. The second is the disappearance of the wake word as an architectural gate: the voice layer reported for the device described *infra* note 20 is designed to listen and speak simultaneously rather than to wait for an address. See Gurman, *infra* note 20.

2. Automatic estimation of a speaker's age from voice, alone or fused with facial imagery, is an established biometrics task. See M. Fairhurst, C. Li & M. Da Costa-Abreu, *Predictive Biometrics: A Review and Analysis of Predicting Personal Characteristics from Biometric Data*, 6 IET Biometrics 369 (2017) (reviewing automated prediction of age and of emotional or mental state from biometric data); E. Pantraki, C. Kotropoulos & A. Lanitis, *Age Interval and Gender Prediction Using PARAFAC2 and SVMs Based on Visual and Aural Features*, 6 IET Biometrics 290 (2017); cf. N. Lavan, *The Effect of Familiarity on Within-Person Age Judgements from Voices*, 113 Brit. J. Psychol. 287 (2021) (human listeners estimate age from voice with above-chance accuracy).

3. See M.P. Ramaswamy & S. Palaniswamy, *Multimodal Emotion Recognition: A Comprehensive Review, Trends, and Challenges*, 14 WIREs Data Mining & Knowledge Discovery (2024) (reviewing automated emotion recognition from facial, vocal, and physiological signals); Fairhurst et al., *supra* (emotional and mental state as predictable characteristics).

4. H.R. 9619, 119th Cong. (2026) [hereinafter People-First Chatbot Act or the Act]. The bill was introduced by Representative Foushee and referred to the Committee on Energy and Commerce. Its enacting clause states its purpose: "To require artificial intelligence chatbot providers to provide data privacy and security, and for other purposes." *Id.*

5. *Id.* § 2(a).

6. *Id.* § 2(a)(1)(A)–(D).

creates a private right of action with statutory damages.⁷ It is, in ambition, the most direct federal answer yet proposed to the question of what a conversational machine may do with what it hears.

It cannot answer the question posed by the child in the kitchen.

The reason is a single word. Every operative duty in the Act runs through the term *user*. A provider may process personal data only where the processing is “necessary to fulfill an express request made by a user and that user has provided affirmative consent.”⁸ A provider must disable risky design features where it knows or should know “that a user is under the age of 18.”⁹ A provider is strictly liable to “[a] user of an artificial intelligence chatbot who suffers an injury in fact caused through the use” of it.¹⁰ The Act’s notice provision tells “a user” that “the user is interacting with an artificial intelligence chatbot.”¹¹ Remove the term and the statute has no addressees.

Section 5(22) then supplies the definition in full: “The term ‘user’ means an individual, regardless of age.”¹²

That sentence performs one function, which is to foreclose an argument that children are excluded. It performs no other. It names a genus and supplies no differentia.¹³ Read for what it says, every individual on earth is a user of every chatbot, and the Act’s duties are owed to all of them at once. No court will read it that way. So a court will read *user* against the provisions surrounding it, and every one of those provisions presupposes a person who reached out and touched the machine. The Act’s own definition of input data is information “provided to an artificial intelligence chatbot

7. *Id.* § 4(c).

8. *Id.* § 2(a)(1)(A).

9. *Id.* § 2(c)(3).

10. *Id.* § 4(d).

11. *Id.* § 2(b)(2).

12. *Id.* § 5(22).

13. A definition of the form “X means Y” does analytic work only where Y is narrower than the field of possible referents. “User means an individual” is coextensive with the universe of natural persons. On the interpretive consequence, see *infra* Part III.

by a user.”¹⁴ Its processing exception turns on “an express request made by a user.”¹⁵ Its notice tells the user that the user “is interacting.”¹⁶

The child in the kitchen provided nothing, requested nothing, and did not reach out. The machine reached her.

The title of this Article is therefore a provocation rather than a holding, and it is answered in the Act’s own words: she is an individual, regardless of age. Whether that makes her a user of a machine she never addressed is the question the Act cannot answer, and this Article is about the consequences of that silence.

The stakes of the silence are simple to state. The bill presents itself, in its enacting clause and its largest operative block, as a data privacy and security statute for artificial intelligence,¹⁷ and it was introduced with the public endorsement of a coalition of civil society organizations. Yet measured against the only nationwide children’s privacy statute in existence, it neither meets nor exceeds the protection COPPA already provides for the paradigm device of the period, and in the scenarios developed below it falls short of COPPA’s floor while its endorsers described it as an advance. Demonstrating that shortfall, provision by provision, is the contribution.

This Article proceeds as follows. Part II sets out the Act’s user-relative architecture and shows that it is not incidental but load-bearing. Part III demonstrates that Section 5(22) does no analytic work and that construction must therefore proceed from context, which points away from the child. Part IV runs the fork: if the child is a user, the Act’s consent architecture is unsatisfiable by the device as designed; if she is not, the Act does not reach her at all. Both answers defeat the statute’s stated purpose. Part V isolates the necessity limiter in Section 2(a)(1)(A) and shows that the Act permits one individual to consent to the processing of a second individual’s biometric data, a result Section 5(2) was drafted to prevent. Part VI takes up the counting problem: the Act denominates damages in violations and never defines a violation, and the indeterminacy spans four orders of

14. People-First Chatbot Act § 5(13).

15. *Id.* § 2(a)(1)(A).

16. *Id.* § 2(b)(2).

17. People-First Chatbot Act, enacting clause; *id.* § 2(a) (captioned DATA PRIVACY AND SECURITY).

magnitude. Part VII distinguishes requester-keyed statutes from data-subject-keyed statutes, using COPPA as the comparator, and locates the source of the pathology. Part VIII shows that two of the Act's central remedies are specified in terms that cannot operate on a screenless shared device. Part IX proposes three amendments.

A word on timing is owed first. The bill is weeks old, and definitions are rebuilt in markup as a matter of routine, so a drafting critique might seem addressed to a moving target. Two facts answer the objection. The Act's language derives from a model bill that has already been introduced in seven state legislatures this year, in Arizona, Illinois, Kansas, Louisiana, Maryland, South Carolina, and Vermont, carrying the same keying, so the defect this Article identifies is propagating through statehouses rather than awaiting repair in a single committee.¹⁸ And the window before markup is the window in which drafting analysis can still be acted on; the three amendments proposed in Part VIII are addressed to it.

A word on the object of study. The argument does not depend on any device that has not shipped. The ambient household category is mature and large, and every device in it runs the fork set out below.¹⁹ The category is about to be entered by the largest chatbot provider in the United States, which is what makes the question live rather than academic, and reporting on that entry is used here as illustration only.²⁰ If it never ships, the defect remains, because the defect is in the text. One reported attribute does matter to the analysis and should be named: the device is described as mobile and battery-powered rather than fixed.²¹ A fixed device makes the bystander

18. The model bill was developed by the Consumer Federation of America with the Electronic Privacy Information Center and Fairplay and released in January 2026. Press Release, Consumer Fed'n of Am., *CFA Celebrates House Introduction of the People-First Chatbot Act* (July 10, 2026), <https://consumerfed.org/news/press-releases/cfa-celebrates-house-introduction-of-the-people-first-chatbot-act/> (listing the seven states).

19. *See supra* note 1.

20. Mark Gurman, *OpenAI's First Device Will Be Home Speaker Built as AI Companion*, BLOOMBERG (July 14, 2026), <https://www.bloomberg.com/news/articles/2026-07-14/openai-s-first-device-will-be-moveable-screenless-speaker-built-as-ai-companion> (reporting a mobile, screen-free smart speaker with camera and sensors, designed as a humanlike companion for the home, drawing on personal information including emails) (last visited July 15, 2026). Earlier reporting placed the device in public view five months before introduction of the Act. *See* Juli Clover, *Jony Ive's First OpenAI Device Will Be Smart Speaker With Camera, 2027 Launch Planned*, MACRUMORS (Feb. 20, 2026), <https://www.macrumors.com/2026/02/20/jony-ive-openai-smart-speaker-2027/> (last visited July 15, 2026). Nothing in this Article turns on whether that device ships.

21. *See supra* note 20 (describing a moveable unit).

class the household. A portable one makes the bystander class everyone the owner carries it past: the waiting room, the school pickup line, the birthday party at someone else's house. The Article uses the kitchen throughout because it is the easiest room to see, not because the problem stays in it.

II. The User-Relative Architecture

A. Every Duty Has the Same Addressee

The Act contains four operative blocks. Each is keyed to a user.

Section 2(a) governs data. Its lead prohibition bars a provider from processing “personal data other than input data to inform chatbot outputs unless the processing of personal data is necessary to fulfill an express request made by a user and that user has provided affirmative consent.”²² The advertising prohibitions bar processing “the chat log of a user” to select or customize advertisements.²³ The training prohibitions bar processing “the chat log or personal data of a user” where the provider knows or should know “that the user is under the age of 18,” absent parental consent.²⁴ The profiling prohibition bars use of any classification of “the personality or behavioral characteristic of a user” beyond what is necessary to fulfill “an express request made by a user.”²⁵ Retention limits run to “the chat log of a user.”²⁶ Access and deletion rights run to the user.²⁷

Section 2(b) governs transparency. The provider may not represent that “the input data or chat log of a user is confidential.”²⁸ It must give “a clear, conspicuous, and explicit notice to a user that the user is interacting with an artificial intelligence chatbot rather than a human.”²⁹

22. People-First Chatbot Act § 2(a)(1)(A).

23. *Id.* § 2(a)(1)(B).

24. *Id.* § 2(a)(1)(C)(i).

25. *Id.* § 2(a)(1)(D).

26. *Id.* § 2(a)(1)(F).

27. *Id.* § 2(a)(2)(A)–(B).

28. *Id.* § 2(b)(1)(A).

29. *Id.* § 2(b)(2).

Section 2(c) governs safety. The monthly assessment duty is product-level and is the one duty in the Act that does not name a user.³⁰ The design duty is not. Where a provider knows or should know “that a user is under the age of 18,” it “shall disable for that user any feature or setting that creates an unreasonable risk of the user suffering a covered harm, emotional dependence, or compulsive usage.”³¹ The phrase “for that user” appears in the operative command.

Section 4 governs enforcement. The private right of action runs to “[a] person injured by an act or practice in violation of section 2,” which is broader than user but requires a Section 2 violation as its predicate, and every Section 2 prohibition except the monthly assessment is itself user-keyed.³² The strict-liability injury track runs to “[a] user of an artificial intelligence chatbot who suffers an injury in fact caused through the use of the artificial intelligence chatbot.”³³

The architecture is consistent, deliberate, and, on the devices the drafters had in mind, unremarkable.

B. Why the Architecture Held Until Now

User-relativity is invisible where the class of persons whose data a system processes is coextensive with the class of persons who asked it to. On a phone, a laptop, or a browser tab, those classes are the same class. One person opens the application. One person types. One person’s data goes to the server. The statute can key every duty to that person because there is no one else in the frame.

Ambient devices break the identity. The class of persons whose data the system processes is the class of persons in physical range. The class of persons who asked is, frequently, empty. Between those two classes sits every person in the household who is not holding the device, and the household member least likely to be holding it, and most likely to be in range of it, is a child.

30. *Id.* § 2(c)(1) (requiring assessment of “the artificial intelligence chatbot for the risk of any covered harm, emotional dependence, or compulsive usage, according to metrics set forth in rules promulgated by the Commission”). The metrics are deferred. *Id.* § 3(a)(3).

31. *Id.* § 2(c)(3).

32. *Id.* § 4(c)(1). The predicate structure is the load-bearing feature. A plaintiff who cannot identify a Section 2 violation has no Section 4(c) claim regardless of injury.

33. *Id.* § 4(d).

The identity is not breaking in theory. It is breaking in litigation, on the defendant this Article's device belongs to, and the first case to break it involved no ambient device at all. In December 2025 the estate of Suzanne Adams, an eighty-three year old Connecticut woman, sued the provider, its chief executive, its principal commercial partner, and twenty unnamed employees and investors for her wrongful death, alleging that months of chatbot conversation cultivated her son's paranoid delusions and aimed them at her before he killed her and himself.³⁴ The estate pleads her as a person who never used the product, did not know it was describing her to her son as a threat, and had no ability to protect herself from a danger she could not perceive.³⁵

Her injury is not this Article's injury, and the difference matters. She was harmed by the outputs generated for another person's requests, not by the capture of her own data; she is a bystander to the conversation rather than a data subject within it. But she establishes the antecedent point on which everything below depends, and she establishes it in a filed complaint rather than a thought experiment: the class of persons a conversational machine can injure is already larger than the class of persons who used one, courts are already being asked what to do about that, and the largest such action on the docket names as its decedent someone who never typed a word. A statute that keys every duty to the requester was drafted for a world in which that sentence could not be written.

This is not a novel observation about ambient computing. It is the observation that produced COPPA's structure, which keys its central duty to the age of the data subject and asks nothing about who initiated.³⁶ It is the observation that produced the FTC's 2019 enforcement posture toward voice assistants marketed to families. What is novel is that a bill drafted in 2026, aimed at

34. *See Open AI, Microsoft Sued Over ChatGPT's Alleged Role in Fueling Man's "Paranoid Delusions" Before Murder-Suicide in Connecticut*, CBS NEWS (CBS/AP) (Dec. 11, 2025), <https://www.cbsnews.com/news/open-ai-microsoft-sued-chatgpt-murder-suicide-connecticut/> (last visited July 16, 2026) (reporting the filing in California Superior Court in San Francisco, that it was the first chatbot wrongful death action to name the commercial partner and the first to involve a homicide rather than a suicide, and that the provider was at that time defending seven other actions alleging the product drove users to suicide and harmful delusions).

35. *Id.* The estate's characterization is quoted in the reporting; the allegation has not been adjudicated and is recited here for what it shows about the shape of the class rather than for its truth.

36. 15 U.S.C. §§ 6501–6506; 16 C.F.R. pt. 312. The operative trigger is the collection of personal information from a child by an operator of a website or online service directed to children, or with actual knowledge of collection from a child. 16 C.F.R. § 312.3.

conversational machines, and captioned as a data privacy statute, keys every duty to the requester and never asks who else is in the room.

III. The Definition That Defines Nothing

A. Section 5(22) Does One Job

“The term ‘user’ means an individual, regardless of age.”³⁷

A statutory definition earns its place by narrowing. The definition of *chatbot* narrows: it excludes systems whose responses are fully predetermined and systems unable to respond outside a narrow purpose.³⁸ The definition of *covered harm* narrows to three things.³⁹ The definition of *affirmative consent* narrows aggressively, listing five conditions and five exclusions.⁴⁰ Each of these tells a reader something she could not have inferred.

Section 5(22) tells a reader that a user is a person. The clause that follows, “regardless of age,” does real work, and it is the only work the definition does: it forecloses an argument that minors fall outside the Act’s protections. But an anti-exclusion clause is not a definition. It removes a limit that no one had imposed. What it does not do is state the relation between the individual and the chatbot that makes the individual a user of it.

Read literally, the definition is unbounded. Every individual is an individual. If *user* means individual, then a provider’s Section 2(a)(1)(A) duty is owed with respect to every individual whose personal data it processes, which is a result the drafters would likely welcome, and, simultaneously, its Section 2(b)(2) duty is to give an hourly interactive notice to every individual on earth, which is absurd. A construction that produces an absurdity in one subsection and a windfall in another is not a construction. It is a signal that the term is doing no defining.

37. People-First Chatbot Act § 5(22).

38. *Id.* § 5(3)(A)(i)–(ii).

39. *Id.* § 5(9) (“death, a suicide attempt, or a psychiatric emergency that results in urgent medical treatment”).

40. *Id.* § 5(2)(A)(i)–(v), (B)(i)–(v).

B. Construction Must Proceed From Context, and Context Points Away From the Child

Where a defined term is defined in a way that cannot bear its own weight, a court construes the term from its statutory surroundings. Words are known by the company they keep,⁴¹ and a statute is read so that no provision is rendered superfluous.⁴²

Apply both canons to *user* and the result is unfavorable to the child in the kitchen.

Section 5(13) defines input data as information “provided to an artificial intelligence chatbot by a user.”⁴³ To provide is to act. Section 2(a)(1)(A) permits processing “necessary to fulfill an express request made by a user.”⁴⁴ To request is to act. Section 2(b)(2) requires notice that “the user is interacting with an artificial intelligence chatbot.”⁴⁵ Section 2(a)(2)(A) requires the provider to furnish the chat log “to the user upon request by the user.”⁴⁶ Section 2(a)(1)(G) bars retaliation against a user “for refusing to consent,” which presumes a user who was asked.⁴⁷ Section 4(d) reaches injury “caused through the use” of the chatbot by the user.⁴⁸

Six provisions, one picture: a user is a person who does something to the machine. The child did nothing to the machine. The machine did something to her.

Against that, the Act supplies exactly one counterweight, and it is the phrase “regardless of age.” A litigant for the child will argue that the clause was inserted precisely because children do not interact with these systems on adult terms and that the drafters meant to sweep in the child in the room. That argument has intuitive appeal and no textual support beyond the clause itself. The clause addresses age. It does not address relation.

41. *Gustafson v. Alloyd Co.*, 513 U.S. 561, 575 (1995) (applying *noscitur a sociis* to construe a term by reference to the words associated with it).

42. *Corley v. United States*, 556 U.S. 303, 314 (2009) (“[A] statute should be construed so that effect is given to all its provisions, so that no part will be inoperative or superfluous, void or insignificant.”) (internal quotation marks and citation omitted).

43. People-First Chatbot Act § 5(13).

44. *Id.* § 2(a)(1)(A).

45. *Id.* § 2(b)(2).

46. *Id.* § 2(a)(2)(A).

47. *Id.* § 2(a)(1)(G).

48. *Id.* § 4(d).

C. The Definition Is Not Saved by the Inclusions Clause

Section 5(3)(B) provides that chatbot “includes” a service that, “in a manner that simulates a sustained interpersonal relationship or emotional interaction with the user,” exhibits “persistent responses that suggest affection or attachment directed toward the user,” engages in “interactions involving emotional disclosures from the user,” or “presents at least one persistent identity, persona, or character or holds itself out as a sentient being, fictional character, or social entity.”⁴⁹

This clause is the clearest evidence in the Act of what the drafters were aiming at. It describes a companion device with precision, and it was written before any such device reached the market.⁵⁰ A reader could reasonably conclude from Section 5(3)(B) that a humanlike home companion is the paradigm case the Act exists to regulate.

But the clause defines the machine, not the person. Every one of its triggers is stated relative to “the user.” A device that simulates a sustained interpersonal relationship with the parent, and merely photographs the child, satisfies Section 5(3)(B) as to the parent and says nothing about the child. The inclusions clause identifies the target and hands the target the same exit as the rest of the Act.

IV. The Fork

A. Horn One: The Child Is a User

Assume a court reads “regardless of age” expansively and holds that a child in range of an ambient device is a user of it.

Section 2(a)(1)(C)(i) then bars the provider from processing her chat log or personal data “without the affirmative consent of the parent or legal guardian of that user,” where the provider knows or should know she is under 18.⁵¹ The knowledge predicate is satisfied on these facts more easily than on any other fact pattern in consumer technology. A camera performing facial

49. *Id.* § 5(3)(B)(i)–(ii).

50. *See supra* note 20 and accompanying text.

51. People-First Chatbot Act § 2(a)(1)(C)(i).

recognition and a model capable of inferring age from vocal characteristics do not merely permit an inference of minority; they constitute one. “[K]nowledge fairly implied on the basis of objective circumstances” is the statutory standard,⁵² and a system engineered to identify who is in the room satisfies it by operating as designed.

The consent the Act then requires is defined at Section 5(2), and its conditions are severe. It must be “a clear affirmative act” signifying “freely given, specific, informed, and unambiguous authorization,” made in response to a specific request, “provided to the user in a clear and conspicuous standalone disclosure,” including a written description in easy-to-understand language, made in a manner accessible to a user with a disability, available in every language the provider offers, and structured so that refusal is “at least as prominent” as consent and takes no more steps.⁵³ Acceptance of general terms of use is expressly excluded.⁵⁴ Inferred consent from continued use is expressly excluded.⁵⁵

A screenless device in a shared room cannot present a standalone written disclosure to the parent of a child who has just walked in. There is no interface on which to present it, no moment at which to present it, and no way to identify whose parent to present it to before the processing that would identify her has already occurred. The consent is not difficult to obtain. It is unobtainable by the architecture of the product, which means the device violates Section 2(a)(1)(C)(i) continuously and cannot cure.

That is a coherent result. It is also a result no court reaches lightly, because it renders an entire product category per se unlawful on a construction the text does not compel. Horn one is available. It is not where a court lands absent a clearer command than “regardless of age.”

B. Horn Two: The Child Is Not a User

Assume instead the contextual construction, which the canons favor.⁵⁶

52. *Id.* § 2(a)(1)(C)(i).

53. *Id.* § 5(2)(A)(i)–(v).

54. *Id.* § 5(2)(B)(ii).

55. *Id.* § 5(2)(B)(i).

56. *See supra* Part III.B.

The child is not a user. Section 2(a)(1)(B) does not protect her, because it reaches “the chat log of a user.” Section 2(a)(1)(C) does not protect her, because it reaches “the chat log or personal data of a user” and gates on the provider’s knowledge that “the user” is a minor. Section 2(c)(3) does not protect her, because it commands the provider to disable features “for that user,” and she is not one. Section 4(d) does not reach her, because it runs to “[a] user.”

One provision arguably survives, and the argument must be stated fairly because the provision contains a clause that cuts both ways. Section 2(a)(1)(A) bars a provider from processing “personal data other than input data *to inform chatbot outputs*” unless the processing is “necessary to fulfill an express request made by a user” who consented.⁵⁷ The prohibition is not phrased as to the user’s personal data. It is phrased as to personal data, unqualified, with the user appearing only in the exception. Her face is personal data.⁵⁸ It is not input data, because she provided nothing.⁵⁹ And no user made an express request that requires knowing whose face is in the room.

The clause that cuts is “to inform chatbot outputs.” On one parse it is a purpose limiter on the prohibition itself: only processing done to inform an output is barred, and a provider that routes the child’s biometrics into storage, security screening, or a training pipeline that never touches inference can argue the prohibition never attached to that processing at all. Taken seriously, that parse produces a perverse gradient. Data processed to serve the person in the room is regulated; data processed only to harvest the person in the room is not; and the more purely extractive the architecture, the less of the Act applies to it. A privacy statute that regulates service and exempts extraction has the sign of its incentive reversed.

The parse fails on the technical facts, and the failure should be stated as an evidentiary matter rather than assumed away. A large language model is autoregressive: it generates each token conditioned on everything in its context window, with an attention mechanism that computes

57. People-First Chatbot Act § 2(a)(1)(A) (emphasis added).

58. *Id.* § 5(15)(A) (defining personal data as “any information, including derived data, inferences, or unique identifiers, that is linked or reasonably linkable, alone or in combination with other information, to an identified or identifiable individual”).

59. *Id.* § 5(13); *see supra* Part III.B.

weights over every element of that context for every prediction.⁶⁰ There is no principled partition between context that is present and context that informs. The empirical literature makes the point directly: adding material to a model’s context measurably alters its predictions even when the material is irrelevant or misleading, a phenomenon documented in the natural-language-processing literature as mispriming,⁶¹ and formal work on in-context learning quantifies how in-context material shifts predictions even against the model’s own pre-training.⁶² An entire engineering subdiscipline, context engineering, exists precisely because whatever occupies the context participates in the output. If the child’s face and voice reach the system that generates responses, they inform outputs in the only sense the architecture admits.

The provider’s remaining move is to claim an internal partition: the sensors feed a pipeline that the conversational model never sees. Two answers. First, the device category is built against that claim. On the mature devices, wake-word detection requires the system to process all ambient audio continuously in order to determine whether anyone is addressing it at all, so the listening precedes and conditions every interaction. On the device this Article takes as its object, even that gate is gone: the reported voice layer is designed to listen and speak at the same time, which is to say the audio path into the conversational model has no closed state to partition away from.⁶³ The listening is the product. Second, the partition is an assertion about the internals of a proprietary system that the plaintiff cannot observe, the provider controls, and no term of the Act requires to

60. For the standard technical descriptions of the context window (“[t]he span of text that a language model can consider at one time when generating or interpreting language”) and the attention mechanism (computing weights “for every word *** within the context window”), see Jiang Lu et al., *Large Language Models and Their Applications in Drug Discovery and Development: A Primer*, 18 CLINICAL & TRANSLATIONAL SCI. no. 4 (2025), <https://doi.org/10.1111/cts.70205> (last visited July 15, 2026); Mauro Giuffrè et al., *Systematic Review: The Use of Large Language Models as Medical Chatbots in Digestive Diseases*, 60 ALIMENTARY PHARMACOLOGY & THERAPEUTICS 144 (2024), <https://doi.org/10.1111/apt.18058> (last visited July 15, 2026).

61. See Ole Pütz, *Co-Textual Dopes: How LLMs Produce Contextually Appropriate Text in Chat Interactions with Humans Without Access to Context*, 36 J. LINGUISTIC ANTHROPOLOGY no. 1 (2025), <https://doi.org/10.1111/jola.70036> (last visited July 15, 2026) (surveying the mispriming results and demonstrating that different context yields different continuations as a structural property of the architecture).

62. Pengfei He et al., *Towards the Effect of Examples on In-Context Learning: A Theoretical Case Study*, 14 STAT no. 1 (2025), <https://doi.org/10.1002/sta4.70045> (last visited July 15, 2026).

63. See Gurman, *supra* note 20 (reporting that the device’s communication abilities rely on a voice mode able to listen and talk simultaneously). A defendant asserting that ambient audio reaches some other pipeline must explain what the conversational model is listening to, given that continuous listening is the advertised capability of the layer the device runs on.

be disclosed. A prohibition whose reach depends on which internal bus the defendant claims the data traveled is not a prohibition. It is an attestation regime without an attestation requirement.

So the better reading, and the one the architecture compels, is that ambient capture by a conversational system processes personal data to inform chatbot outputs. On that reading she has a Section 2 violation, and Section 4(c) runs to “[a] person injured by an act or practice in violation of section 2,” not to a user.⁶⁴ She can sue. But the clause should be repaired rather than litigated, and the amendment in Part IX now has a fourth line.

The remedy is the tell. Her recovery is governed by Section 4(c)(1)(B)(i): “up to \$10,000 per violation.”⁶⁵ A child who is a user, protected by Section 2(c)(3), recovers “not less than \$50,000,” rising to not less than \$100,000 where the violation results in emotional dependence or compulsive usage and not less than \$250,000 where it results in a covered harm.⁶⁶

Those figures are not two points on one scale. “Up to” is a discretionary ceiling with no floor; a court may award one dollar. “Not less than” is a mandatory floor. The most the bystander child can recover before any willfulness enhancement is one fifth of the least the user child is guaranteed. Apply the Act’s willfulness multiplier of up to five times to the bystander’s ceiling and she arrives, at her theoretical maximum, at exactly the number where the user child begins.⁶⁷

Same device. Same camera. Same data on the same server. The difference is that one child spoke to the machine and the other was merely in front of it.

C. Both Horns Defeat the Purpose

The Act’s enacting clause states its object as requiring providers “to provide data privacy and security.”⁶⁸ Horn one delivers that object by making the leading ambient product category unlawful per se, which is a policy Congress did not debate and would not have adopted silently in a

64. People-First Chatbot Act § 4(c)(1).

65. *Id.* § 4(c)(1)(B)(i).

66. *Id.* § 4(c)(1)(B)(iii)(I)–(III).

67. *Id.* § 4(c)(2) (permitting the court to increase the award to “not more than 5 times the amount available under paragraph (1)(B)”). The multiplier is not a limit on the number of violations; the Act contains no such limit. *See infra* Part VI.

68. People-First Chatbot Act, enacting clause.

definitions section. Horn two delivers a data privacy statute under which the child least able to protect herself, and least able to have chosen the exposure, receives the Act's weakest remedy for the same conduct.

A fork whose horns are per se prohibition and near-total exclusion is not a close question about which reading is better. It is evidence that the term being construed is not doing the work the statute assigned it.

V. The Necessity Limiter and Consent That Belongs to Someone Else

A. The Limiter

Section 2(a)(1)(A) contains two conditions, and commentary on user-keyed privacy provisions tends to notice only the second. The exception applies where processing is “necessary to fulfill an express request made by a user” *and* “that user has provided affirmative consent.”⁶⁹ Consent alone does not open the exception. The processing must also be necessary to fulfill the request.

The limiter matters because it defeats the obvious compliance strategy. A provider cannot obtain a single Section 5(2)-compliant consent at setup and thereafter process whatever it captures. Each act of processing must be traceable to necessity for a specific express request.

Work the arithmetic on an ordinary household exchange. A parent, standing at the counter with a child beside her, asks the device for ticket prices to a theme park for the two of them that weekend. There is an express request. It was made by a user. That user gave consent at setup, and the provider built the consent flow to Section 5(2) specification because Section 5(2)(B)(ii) forbade it from relying on terms of use.

Now ask what processing is necessary to fulfill that request. Retrieving ticket prices requires the content of the question. It does not require a faceprint of the child. It does not require an age inference from the child's voice. It does not require an emotional classification of the child's expression. None of that is necessary to answer a question about ticket prices, and the Act permits

69. People-First Chatbot Act § 2(a)(1)(A) (emphasis added).

only what is necessary. The exception fails on the limiter even with a consenting user standing at arm's length, and it fails as to the household's own child, in her own kitchen, with her parent present and cooperating.

B. Whose Consent Is It

Now the structural defect the limiter exposes.

Read Section 2(a)(1)(A) again and ask whose consent unlocks the exception. It is the consent of “that user,” meaning the user who made the express request. It is not the consent of the individual whose personal data is processed. The Act nowhere requires the data subject's consent to the processing of the data subject's own personal data. It requires the requester's consent.

That is a category error with consequences. Wherever the necessity limiter can be satisfied, and there will be requests for which knowing who is in the room is necessary to the answer, one individual's consent authorizes the processing of a second individual's biometric data. The parent who asks the device a question about her household consents, on this reading, to the processing of her child's face. And the child who is not hers, the visitor, the cousin, the neighbor's child at a birthday party, is exposed by a consent given by an adult who is not that child's parent and has no authority to give it.

Section 5(2) was drafted to prevent exactly this. Its conditions insist that consent be “specific,” “informed,” and “unambiguous,” from an individual who received a standalone disclosure describing the practice.⁷⁰ Every one of those conditions describes a relationship between a person and the processing of that person's own data. The Act builds a consent instrument of unusual rigor and then attaches it to the wrong individual.

C. The Visitor

The visiting child is the case where both statutes converge and neither can be satisfied.

70. *Id.* § 5(2)(A)(i)–(ii).

Under COPPA, the operator needs verifiable parental consent from her parent before collecting her personal information, and the amended Rule now classes voiceprints as personal information.⁷¹ The operator also needs a second, separate verifiable parental consent before disclosing her personal information to a third party for targeted advertising or for training or developing artificial intelligence, purposes the Commission held are never integral to any service.⁷² Her parent is not present, received no direct notice, and gave no consent of either kind. The Rule's general command admits no ambiguity: an operator "is required to obtain verifiable parental consent before any collection, use, or disclosure of personal information from children."⁷³

There is exactly one provision under which a voice product may capture a child's voice without that consent, and its conditions are worth setting out in full because the device fails every one of them. Section 312.5(c)(9) excepts collection "[w]here an operator collects an audio file containing a child's voice, and no other personal information, for use in responding to a child's specific request and where the operator does not use such information for any other purpose, does not disclose it, and deletes it immediately after responding to the child's request."⁷⁴ Five conditions, conjunctive. A camera-equipped speaker collects other personal information, so the first fails. The visiting child made no request, so the second fails. Training or advertising is another purpose, so the third fails. Transmission to a server is disclosure, so the fourth fails. Retention of any kind defeats the fifth.

Note what that exception is. It is a requester-keyed carve-out, the same species as Section 2(a)(1)(A). But the Commission cabined it with five conjunctive conditions ending in immediate deletion, and confined it to audio alone. The Act's requester-keyed rule is cabined by necessity and nothing else, and it reaches every category of personal data the device can capture.

71. 16 C.F.R. § 312.2 (definition of personal information, as amended); Children's Online Privacy Protection Rule, 90 Fed. Reg. 16,977 (Apr. 22, 2025) (final amendments; full compliance required by April 22, 2026).

72. *Id.*; see also 16 C.F.R. § 312.5(a)(2). The never-integral holding forecloses the argument that model training is core to service delivery.

73. 16 C.F.R. § 312.5(a)(1). The separate-consent requirement is at § 312.5(a)(2) (parent must be given the option to consent to collection and use "without consenting to disclosure of his or her personal information to third parties, unless such disclosure is integral to the website or online service," and separate verifiable parental consent must be obtained for that disclosure).

74. 16 C.F.R. § 312.5(c)(9). An operator relying on the exception must also disclose in its online notice how it uses such audio files and that it deletes them immediately. *Id.* § 312.4(d)(4).

The impossibility is provable from the Rule’s own list of consent methods, which is closed and enumerated. Every one requires an out-of-band channel to a specific adult: a signed form returned by mail, fax, or scan; a payment-card transaction that notifies the account holder; a call to trained personnel on a toll-free line; a video-conference with trained personnel; a government-identification check against a database; knowledge-based authentication with dynamic multiple-choice questions a twelve-year-old could not answer; a government photo identification matched by facial recognition against a live image; or email or text coupled with confirmatory steps.⁷⁵ Each presupposes that the operator knows which adult to contact and has a channel to reach that adult before collection. A device in a room the visiting child has just entered knows neither, and it cannot learn the first without performing the collection that required the consent.

Under the People-First Chatbot Act, the exception in Section 2(a)(1)(A) runs on the consent of the user who made the request. No user in that house is her parent. No consent obtainable in that house reaches her.

Sound and light do not stop at the property line, and a birthday party is not an edge case. It is a Saturday.

VI. The Counting Problem

A. A Damages Provision Without a Unit

Section 4(c)(1)(B)(i) awards “up to \$10,000 per violation.”⁷⁶ Section 4(c)(2) multiplies the award by up to five for willful conduct.⁷⁷ Section 4(c)(3) makes costs and reasonable attorney’s fees mandatory for a prevailing plaintiff, which is what makes small claims bringable at all.⁷⁸

75. *Id.* § 312.5(b)(2)(i)–(ix). A safe harbor program may approve an additional method under § 312.5(b)(3), but any method must be “reasonably calculated, in light of available technology, to ensure that the person providing consent is the child’s parent.” *Id.* § 312.5(b)(1).

76. People-First Chatbot Act § 4(c)(1)(B)(i).

77. *Id.* § 4(c)(2).

78. *Id.* § 4(c)(3).

The Act does not define violation. It contains no counting rule, no per-day provision, and no clause stating that each instance of prohibited processing constitutes a separate violation. The term appears throughout Section 4 as the denominator of the entire damages scheme and appears nowhere in Section 5.

Three constructions are available on the text, and they are not close together.

B. Three Constructions

Per operation. Section 5(18) defines processing as “any operation or set of operations performed, whether by manual or automated means, on personal data or input data, such as the use, storage, disclosure, analysis, deletion, or modification of such data.”⁷⁹ The list is disjunctive and enumerated. A single frame of a child’s face is used, stored, disclosed to a server, and analyzed. That is four operations on the Act’s own terms before anyone asks what the analysis inferred.

Section 5(19) then treats inference separately. Profiling means “any form of processing performed on input data or personal data to infer, detect, classify, or designate emotional vulnerability or distinct behavioral characteristics of an individual,”⁸⁰ and Section 2(a)(1)(D) attaches its own prohibition to the use of any such classification beyond what is necessary to fulfill an express request.⁸¹ Detecting that a face is present and inferring what that face is feeling are two acts the Act names in two places with two prohibitions.

The separation is not merely formal, and the empirical literature makes the point sharper than the statute does. Automated inference of emotion from facial movement performs acceptably under laboratory conditions and degrades severely outside them. A systematic review of forty-five studies across seventy-one experiments reports mean accuracy of 85.8 percent for video in laboratory settings and 53.3 percent for video captured in the wild, with a floor of 30.1 percent.⁸² A

79. *Id.* § 5(18).

80. *Id.* § 5(19)(A).

81. *Id.* § 2(a)(1)(D).

82. Wilson Contreras-Higuera, Lucrezia Crescenzi-Lanna & Mirko Duradoni, *The Role of Time in Facial Dynamics and Challenges in Automatic Emotion Recognition (2019–2024)*, 2025 HUM. BEHAV. & EMERGING TECHS. no. 1, tbl. 8, <https://doi.org/10.1155/hbe2/7777949> (last visited July 15, 2026). The review excluded single-use datasets, studies with fewer than six emotion classes, and cases flagged for overfitting before computing the aggregate.

living room is the wild. The same review reports that identical expressions accompany distinct psychological states, that expressivity varies by culture and by sex, and that one study found frowning correlates so poorly with anger that models detected it at 30 percent.⁸³ Its conclusion about what the accuracy figures measure is the one that matters here: algorithmic accuracy reflects proximity to human consensus rather than emotional truth, so it “measures alignment with training data, not universal validity.”⁸⁴ A parallel review states the operational point plainly, that facial expression as an emotion signal “can be easily faked or misunderstood.”⁸⁵

That evidence does double duty. It defeats any argument that profiling a child in a kitchen is “necessary to fulfill an express request” under Section 2(a)(1)(D), because a classification that is right roughly half the time in the environment where it runs cannot be necessary to fulfill anything. And it answers the response that the inference is harmless because it is only a guess. The Act does not prohibit accurate profiling. It prohibits processing “to infer, detect, classify, or designate emotional vulnerability.”⁸⁶ The verbs describe the attempt. A child designated emotionally vulnerable by a system operating at 53 percent carries the designation at the same weight as one designated correctly, and the record that follows her is identical. A provider arguing they are one operation must argue against the structure of the statute’s own definitions.

This is not a novel drafting failure and its consequences are on the books. The Digital Millennium Copyright Act authorizes statutory damages “for each violation” and, like the Act, never says what a violation is.⁸⁷ Courts confronting the omission reached opposite answers on materially similar facts. One held that a simultaneous distribution of a single photograph to 1,147 subscribers was a single violation, analogizing it to the transmission of a television signal, “a

83. *Id.* § 2 (surveying the construction critique and citing Le Mau et al.).

84. *Id.*

85. Manju Priya Ramaswamy & Suja Palaniswamy, *Multimodal Emotion Recognition: A Comprehensive Review, Trends, and Challenges*, 14 WIREs DATA MINING & KNOWLEDGE DISCOVERY no. 6 (2024), <https://doi.org/10.1002/widm.1563> (last visited July 15, 2026).

86. People-First Chatbot Act § 5(19)(A).

87. *Goldman v. Healthcare Mgmt. Sys., Inc.*, 559 F. Supp. 2d 853, 868 (W.D. Mich. 2008) (“Section 1203(c)(3) authorizes statutory damages ‘for each violation’ of sections 1201 and 1202, without defining what constitutes a violation.”).

single act regardless of the number of viewers.”⁸⁸ *Goldman* distinguished that reasoning where the conduct was variable rather than singular, holding the violation was not “a singular, isolated event” because the defendants “allegedly provided computer programs at different times, under different circumstances, to multiple hospitals.”⁸⁹ It then held that whether any given act crossed the line was “a factual question that should be resolved by a jury.”⁹⁰

Map that onto the kitchen. Ambient capture is variable conduct in the *Goldman* sense, not a broadcast: different children, at different times, under different circumstances, in the same room. The per-operation reading is therefore available, the per-course-of-conduct reading is available on the *McClatchey* analogy, and the choice between them is, on *Goldman*’s own disposition, a jury question. A statute whose damages provision resolves into a jury question about what a violation is has not set a price. It has set a lottery.

And the full arc of that lottery has already run once, at state supreme court scale, on biometric facts. The Illinois Biometric Information Privacy Act awards liquidated damages “for each violation” and, like the Act, never defines the unit.⁹¹ In *Cothron v. White Castle System, Inc.*, the Illinois Supreme Court held “that a separate claim accrues under the Act each time a private entity scans or transmits an individual’s biometric identifier or information,” not merely on the first capture.⁹² The court saw the number that holding produced, acknowledged that nothing in the statute suggested “legislative intent to authorize a damages award that would result in the financial destruction of a business,” and held that “policy-based concerns about potentially excessive damage awards under the Act are best addressed by the legislature,” respectfully suggesting that the legislature “review these policy concerns and make clear its intent.”⁹³ The legislature did exactly that: within eighteen months it amended the statute to provide that repeated collections of the same

88. *Id.* at 867 (characterizing *McClatchey v. Associated Press*, No. 3:05-cv-00145 (W.D. Pa. June 4, 2007)).

89. *Id.* at 868.

90. *Id.* (“This Court concludes whether any particular update constitutes a distribution of an infringed copyright is a factual question that should be resolved by a jury.”).

91. 740 ILL. COMP. STAT. 14/20.

92. *Cothron v. White Castle Sys., Inc.*, 2023 IL 128004, ¶ 1.

93. *Id.* ¶¶ 42–43.

biometric identifier from the same person in the same manner constitute a single violation for which the plaintiff recovers once.⁹⁴

That is the completed experiment. An undefined unit went to the courts; the courts read the plain text and picked the large construction; the exposure was pronounced potentially annihilating by the court that created it; and the legislature reversed to one recovery per person. Congress is now drafting the same undefined unit into a statute aimed at devices that perform the scan not at a time clock but continuously, in rooms with children in them. The predictable end state is the Illinois end state: after years of litigation, the bystander child is returned to a single discretionary recovery, which is where Section 4(c)(1)(B)(i) already caps her. The lottery resolves to the house.

Per data subject. Section 4(a)(1) provides that a violation of Section 2 “shall be treated as a violation of a regulation under section 18(a)(1)(B) of the Federal Trade Commission Act,”⁹⁵ and a court may look to Commission practice for the unit. The difficulty is that there is no judicial counting practice to borrow. Children’s privacy enforcement under COPPA has proceeded almost entirely by negotiated consent decree, in which the penalty is a bargained sum rather than a per-violation product computed by a court.⁹⁶ A defendant will urge the per-child unit anyway, because it is the smallest unit that is not facially unreasonable, and a court will have nothing to consult but its own instinct. That is not a construction. It is a vacancy.

The one-operation defense also assumes that a capture has one use, which is the assumption the data-governance literature exists to rebut. Function creep, defined in the canonical treatment as “an imperceptibly transformative and therewith contestable change in a data-processing system’s proper activity,” is the ordinary trajectory of stored personal data rather than an aberration in it.⁹⁷ A

94. Pub. Act 103-0769 (eff. Aug. 2, 2024) (amending 740 ILL. COMP. STAT. 14/20); *see, e.g., McGivney v. ITS Techs. & Logistics, LLC*, 2025 IL App (1st) 241961-U, ¶ 22 (describing the amendment as limiting a plaintiff to recovery for a single violation even where the defendant collected the plaintiff’s biometrics repeatedly).

95. *Id.* § 4(a)(1) (citing 15 U.S.C. § 57a(a)(1)(B)).

96. *See, e.g., United States v. Epic Games, Inc.*, No. 5:22-cv-00518 (E.D.N.C. Feb. 7, 2023) (stipulated order for permanent injunction and civil penalty); *United States v. Edmodo, LLC*, No. 3:23-cv-02495 (N.D. Cal. June 27, 2023) (same). In neither does a court state a rule for counting violations, because in neither did a court have to.

97. Bert-Jaap Koops, *The Concept of Function Creep*, 13 L. INNOVATION & TECH. 29, 29 (2021), <https://doi.org/10.1080/17579961.2021.1898299> (last visited July 15, 2026). Koops situates the concept “in the nexus of reverse adaptation and self-augmentation of technology, incrementalism and disruption in policy and innovation, policy spillovers, ratchet effects, transformative use, and slippery slope argumentation.” *Id.*

faceprint captured to answer a question is available afterward to identify, to authenticate, to train, to segment, to target, and to infer, and the mere fact of retention supplies the opportunity whatever the original purpose was.⁹⁸ That is why the Act's own definition of processing enumerates operations disjunctively, and it is why a court asked to collapse them into one will be collapsing precisely the distinction the drafters drew.

Per course of conduct. A defendant will argue that ambient capture is a single continuous processing program, that the statute regulates “an act or practice,”⁹⁹ and that a practice is one thing however many times it runs. On this reading a household recovers ten thousand dollars total, and OpenAI's exposure for a decade of operation in fifty million homes is bounded by the number of households, not the number of children or captures.

C. Four Orders of Magnitude

Take one afternoon. Fifteen children at a birthday party, moving through a kitchen where the device sits, over three hours.

Per operation, with face and voice each processed on each entry and each entry producing use, storage, disclosure, and analysis, the count runs to hundreds. At the Section 4(c)(1)(B)(i) ceiling the exposure passes one million dollars, and the willfulness multiplier takes it past five.

Per data subject, the count is fifteen. The exposure ceiling is one hundred fifty thousand dollars, and a court may award less, because the provision has no floor.

Per course of conduct, the count is one. The ceiling is ten thousand dollars.

Same afternoon, same house, same children, same facts. The Act's answer ranges from ten thousand dollars to over five million, and the range is not a function of anything a court would call a disputed fact. It is a function of a word Congress used forty-one times in Section 4 and never defined.

98. The observation that storage itself creates the opening, independent of any regulatory permission, is standard in the secondary-use literature. See Sarah B. Evans-Jordan & John-Arne Skolbekken, *Scientific Citizenship's Youngest Domain: Function Creep in Norway's Newborn Screening Programme*, 26 SCI. TECH. & SOC'Y 98, 101 (2021) (“The mere fact of storage provides a logistical opportunity, if not a regulatory one, for unrelated secondary use.”), <https://doi.org/10.1177/0971721820964891> (last visited July 15, 2026).

99. *Id.* § 4(c)(1).

A defendant reaching the top of that range has one more argument, and it is a good one. Aggregated statutory damages are subject to due process limits “in certain extreme circumstances” even where the per-violation figure is constitutionally sound, and an award offends due process where it is “so severe and oppressive as to be wholly disproportionate to the offense and obviously unreasonable in relation to the goals of the statute and the conduct the statute prohibits.”¹⁰⁰ So the per-operation reading is not merely contestable as a matter of construction. It is contestable a second time on constitutional grounds after the plaintiff wins the construction.

That is the Article’s central practical claim. The People-First Chatbot Act creates a private right of action whose value cannot be computed from its text. It can be computed after a court picks a unit the statute does not supply, and then only if the number survives a proportionality challenge the statute does not anticipate.

D. Standing, and the State-Court Answer

A defendant will add that the bystander’s claim is a bare procedural violation and that *TransUnion* forecloses it.¹⁰¹ The argument is serious in federal court and weaker than it first appears here. The traditional analogue is intrusion upon seclusion, and a camera capturing a child inside a private residence is closer to that tort than most data claims that have failed the test.¹⁰² And the visiting child’s position is stronger still, not weaker, though the intuition runs the other way. The tort reaches intrusion upon a person’s private affairs, not merely her private spaces, and the child’s

100. *Wakefield v. ViSalus, Inc.*, 51 F.4th 1109, 1118, 1122 (9th Cir. 2022) (adopting the *Williams* test for aggregated awards and remanding a \$925 million judgment for reassessment). The court cautioned that the limit “must be reserved for circumstances in which a largely punitive per-violation amount results in an aggregate that is gravely disproportionate,” lest the judiciary “usurp the power of the legislature.” *Id.* at 1122. *Wakefield* carries a caution flag in the citators and its factor test derives from liquidated-damages cases; it is offered here for the existence of the constitutional ceiling, not for where the ceiling sits.

101. *TransUnion LLC v. Ramirez*, 594 U.S. 413, 426–27 (2021) (holding that a plaintiff does not automatically satisfy the injury-in-fact requirement whenever a statute grants a right and purports to authorize suit to vindicate it, and requiring a concrete harm with a close relationship to a harm traditionally recognized at common law); *see also Spokeo, Inc. v. Robins*, 578 U.S. 330, 341 (2016) (same principle, remanding for concreteness analysis).

102. *Cf. Gadelhak v. AT&T Servs., Inc.*, 950 F.3d 458, 462–63 (7th Cir. 2020) (Barrett, J.) (holding that unwanted intrusion into a private sphere bore a close relationship to a harm traditionally actionable and sufficed for standing under *Spokeo*). The Ninth Circuit has since held that *TransUnion* did not disturb the rule that an intangible injury is concrete where Congress created the cause of action and the injury has a close historical or common-law analog, and applied it to unsolicited telephone calls. *Wakefield v. ViSalus, Inc.*, 51 F.4th 1109, 1116–17 (9th Cir. 2022). A camera capturing a child inside her own home is a stronger privacy analog than a ringing telephone.

face, voice, and inferred emotional state are her private affairs in any room she stands in. What distinguishes her from the household's child is the consent chain: the household's child at least has a parent who purchased the device, configured it, and consented to something. The visitor's chain is empty. No one with authority over her data consented to anything, anywhere, and the assumption-of-risk argument a defendant would reach for is an affirmative defense on the merits, not a standing objection, and it is unavailable against a child in any event.

In any event the Act supplies the workaround in its own text. Section 4(c)(1) authorizes suit “in an appropriate State court or an appropriate district court of the United States.”¹⁰³ State courts are not bound by Article III. The statutory element of injury remains in either forum, but the constitutional overlay does not.

VII. Whose Data It Is Versus Whose Request It Was

A. Two Ways to Key a Privacy Statute

The pathology traced above has a name once the comparison is made.

COPPA keys on the data subject. Its operative question is the age of the person whose information is collected.¹⁰⁴ It does not ask who initiated the session, whether the child made a request, whether anyone consented to anything, or whose account was logged in. It asks how old she is and whether her parent said yes. One fact, one question.

The People-First Chatbot Act keys on the requester. Its operative question is whether processing was necessary to fulfill an express request made by a user who consented.¹⁰⁵ Three conditions, all of which run through a person who may not be the person whose data is at issue.

A data-subject-keyed statute reaches the child on the carpet automatically, because she is the data subject and nothing else is required. A requester-keyed statute has to reason its way to her

103. People-First Chatbot Act § 4(c)(1).

104. 16 C.F.R. § 312.3; 15 U.S.C. § 6502(a)(1).

105. People-First Chatbot Act § 2(a)(1)(A).

through someone else's relationship with the machine, and on an ambient device that relationship may not exist.

The choice is not idiosyncratic to COPPA. It is the architecture of the other dominant privacy regime as well. The General Data Protection Regulation defines personal data as information “relating to an identified or identifiable natural person,” and gives that person a name, the data subject, in the definition itself; every right and every lawful basis in the Regulation then runs to or through that person, and the child-consent provision keys to the age of the child whose data is processed, not to whoever operated the service.¹⁰⁶ So the two privacy frameworks with the largest installed enforcement bases on earth, one American and one European, both key to the person the data is about. The bill before Congress keys to the person who asked the machine a question. Whatever else may be said for that choice, it was not compelled by any drafting tradition; both traditions run the other way.

B. The Counting Problem Is Downstream of the Keying Choice

The two defects are one defect.

A data-subject-keyed statute can count, because the data subject is a countable unit that exists in the world independent of the statute. A child is one child. A requester-keyed statute cannot count without inventing a unit, because the thing it regulates is a relation rather than a person, and relations do not come pre-individuated. There is no natural answer to how many express requests were not made.

That is why the Act, having keyed to the requester, arrived at Section 4 with no unit available and used the word violation without defining it. The omission is not carelessness. It is the downstream consequence of the keying choice, and it is why the fix in Part IX has to reach both.

C. The Comparison Runs the Other Way on Remedy

Fairness to the bill requires the other half of the comparison, and it is the half its critics omit.

106. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 Apr. 2016, arts. 4(1), 8, 2016 O.J. (L 119) 1, <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32016R0679> (last visited July 15, 2026).

COPPA has no private right of action. Enforcement runs through the Commission and state attorneys general.¹⁰⁷ The maximum civil penalty is \$53,088 per violation, a figure set by the Commission's January 2025 inflation adjustment and unchanged for 2026, because the lapse in appropriations from October 1 to November 12, 2025 prevented the Bureau of Labor Statistics from producing October 2025 CPI-U data, leaving no multiplier, and the Office of Management and Budget cancelled the government-wide adjustment.¹⁰⁸ That money goes to the Treasury. The family whose child was recorded receives nothing.

So the Act, whatever its defects, would give that family something twenty-eight years of COPPA has never given it: a path into a courtroom, with fees. That is a real advance and it should survive the amendments proposed below. The objection is not that the Act created the remedy. It is that the Act priced the remedy at a discretionary ceiling of \$10,000 while the Commission's ceiling for identical conduct is \$53,088 and the Act's own floor for a child who happened to speak first is \$50,000. The regulator's ceiling is more than five times the family's ceiling. The family is the only party that was harmed.

One state supplies the second benchmark, and it is the only state that does. Illinois grants a private right of action for biometric violations: \$1,000 in liquidated damages for a negligent violation and \$5,000 for an intentional or reckless one, per violation, with fees, for precisely the identifiers this device collects, voiceprints and scans of face geometry by name, and with no requirement of actual harm beyond the violation itself.¹⁰⁹ The other principal state biometric statutes are enforceable only by the state.¹¹⁰

107. 15 U.S.C. § 6502(c)–(d).

108. Adjustments to Civil Penalty Amounts, 90 Fed. Reg. 5,580 (Jan. 17, 2025) (setting the maximum at \$53,088 for violations of sections 5(l), 5(m)(1)(A), and 5(m)(1)(B) of the FTC Act, the provision through which COPPA penalties run); OFF. OF MGMT. & BUDGET, EXEC. OFF. OF THE PRESIDENT, M-26-11 (Apr. 17, 2026) (directing agencies that no inflation adjustment would be made for 2026 and to continue using 2025 levels). Agencies implemented the direction across the government. *See, e.g.*, No Adjustment to Civil Monetary Penalty Amounts, 91 Fed. Reg. 41,722 (July 7, 2026).

109. 740 ILL. COMP. STAT. 14/10, 14/20; *Cothron*, 2023 IL 128004, ¶ 38 (the statutory violation itself is the injury, and a plaintiff “is not required to plead or prove actual damages”) (citing *Rosenbach v. Six Flags Ent. Corp.*, 2019 IL 123186).

110. *See* TEX. BUS. & COM. CODE § 503.001(d) (enforcement by the attorney general); WASH. REV. CODE § 19.375.030 (enforcement through the consumer protection act).

So place two children in front of the identical device. The child in Illinois has a statutory claim with a fixed liquidated floor per negligent violation, fees, and no injury requirement. The child in Indiana, or in any of the other forty-nine states, has whatever this Act gives her, which for the bystander is a discretionary federal ceiling of \$10,000 that a court may reduce to a dollar. A bill marketed as a national data privacy and security statute for artificial intelligence delivers, to the child the market is about to meet, less than the one state that legislated seriously delivered in 2008. Federal law here is not a floor beneath the states. It is a basement beneath one of them and the only thing standing in the other forty-nine.

One more comparison objection deserves an answer here. The Act reaches every individual under eighteen while COPPA stops at thirteen, and broader age scope can be framed as the advance its endorser described. Coverage of more children with weaker duties is no advance for the children COPPA already covers; for them, the Act substitutes a request-keyed permission structure for a subject-keyed prohibition, and Parts IV and V price that substitution. A statute that adds five ages of nominal coverage while subtracting the keying that made coverage operative has not extended the floor. It has moved the floor's foundation.

D. The Patchwork Outside the Act

A defender of the drafting has one more move, and it points away from the statute. The Act's sponsors present it as preserving state law,¹¹¹ and if the states already reach the child in the kitchen, the federal silence costs little. The backdrop does not deliver what the objection needs.

Federal wiretap law protects an oral communication only where the speaker exhibits a justified expectation that the utterance is not subject to interception, and it permits interception with the consent of a party to the communication.¹¹² The device owner who installed the machine supplies that consent for the household's ambient audio, and the child's captured face is not a

111. Press Release, Office of Rep. Valerie Foushee, *Reps. Foushee, Casar Introduce Legislation to Protect Children and Americans' Privacy from AI Chatbot Harms* (July 9, 2026), <https://foushee.house.gov/media/press-releases/rep-foushee-casar-introduce-legislation-to-protect-children-and-americans-privacy-from-ai-chatbot-harms-and-require-chatbot-safety-assessments> (endorser statement describing the bill as one that "protects state laws and the families already fighting these companies").

112. 18 U.S.C. §§ 2510(2), 2511(2)(d).

communication at all; the federal statute has nothing to say about video or biometric processing inside a home. Two-party consent states move the consent line without moving the subject line. California's eavesdropping statute keys to confidential communications,¹¹³ and the child's presence in the room, her face, and the inferences drawn from both are confided to no one.

The genuine counterexample is Illinois. The Biometric Information Privacy Act keys to the subject of the data rather than to any user: no private entity may collect a biometric identifier, a category that includes a voiceprint and a scan of face geometry, without informed written consent from the subject or the subject's legally authorized representative,¹¹⁴ it supplies a private right of action with liquidated damages,¹¹⁵ and the Illinois Supreme Court has held that a violation of the consent provisions makes a person aggrieved without proof of injury beyond the violation itself.¹¹⁶ The Illinois child on the carpet is protected today, in exactly the keying this Article proposes, and the device category has shipped into Illinois through it. That is confirmation of the amendment's design, and it is no substitute for the amendment. Texas and Washington key their biometric statutes to the individual and then hand enforcement exclusively to the attorney general,¹¹⁷ and most states have nothing at all. A child's protection against a machine that reached her should not depend on which side of a state line her kitchen sits.

VIII. Remedies Written for a Phone, Aimed at Furniture

Two of the Act's central remedies cannot operate on the device category this Article describes, and both failures are visible on the face of the text.

The notice provision requires disclosure "in a font size easily readable by an average user, and not smaller than the largest font size of other text appearing on the interface on which the artificial intelligence chatbot is provided."¹¹⁸ The Act's central anti-deception remedy is specified

113. Cal. Penal Code § 632(a), (c).

114. 740 Ill. Comp. Stat. 14/10, 14/15(b).

115. *Id.* 14/20.

116. *Rosenbach v. Six Flags Ent. Corp.*, 2019 IL 123186, ¶¶ 30, 33.

117. Tex. Bus. & Com. Code § 503.001(d); Wash. Rev. Code § 19.375.030.

118. People-First Chatbot Act § 2(b)(2)(A).

in font sizes. A screenless device has no font size, no other text, and no interface in the sense the provision assumes. The accessibility requirement that follows, that “[t]he notice shall be accessible to users with disabilities,”¹¹⁹ is drafted as a supplement to a visual notice rather than as an independent modality requirement, which leaves a blind user and a screenless speaker in the same gap for opposite reasons. The Commission’s rulemaking authority over “the form and content of the disclosures” could repair this,¹²⁰ which is to say the provision does not presently work and Congress has delegated the question of whether it ever will.

The delegation deserves a paragraph of its own, because the natural response to every defect in this Article is that Section 3(b) authorizes the Commission to “promulgate any other regulation necessary to implement this Act,”¹²¹ and an agency could in principle define *user*, define *violation*, and specify a screenless notice. Three answers. First, an agency construing the operative terms of a statute after *Loper Bright* receives no deference for the construction; the definitions would be litigated de novo in the first enforcement action, which returns the indeterminacy to the courts with a detour.¹²² Second, a rule that selects the counting unit for a private damages provision selects the difference between a ten-thousand-dollar statute and an industry-reshaping one, which is the profile of a question courts now reserve to Congress rather than agencies. Third, and independent of doctrine, the assumption that the agency will define the terms protectively is an assumption about the agency’s priorities, and those priorities are an artifact of whoever holds the agency. A statute whose protective scope is left to future rulemaking is not a guarantee of protection. It is a bet on the next several administrations, placed with other people’s children.

The design duty commands the provider to “disable for that user any feature or setting” creating unreasonable risk.¹²³ One speaker, one room, one household, several people, no accounts. There is no for-that-user on a shared object. The provider’s compliance options are to disable the feature

119. *Id.* § 2(b)(2)(B).

120. *Id.* § 3(a)(1)–(2).

121. People-First Chatbot Act § 3(b).

122. *Loper Bright Enters. v. Raimondo*, 603 U.S. 369 (2024).

123. *Id.* § 2(c)(3).

for the household, which is a remedy the Act did not authorize and the provider will not volunteer, or to identify each individual in the room in order to apply per-person settings, which requires precisely the biometric processing Section 2(a) restricts. The Act's safety provision, applied to this device, can be satisfied only by conduct the Act's privacy provision prohibits.

IX. Three Amendments

Each of the following exists in current federal law. None disturbs the bill's structure.

First, define the term or remove its load. Either Section 5(22) supplies a relation, or the operative prohibitions in Section 2(a) are rekeyed to the data subject. The cleaner repair is the second, and COPPA supplies the model: a prohibition on processing the personal data of an individual the provider knows or should know is under 18, without verifiable parental consent from that individual's parent, stated without reference to who made a request. That sentence reaches the child on the carpet, the child at the counter, and the visitor, and it requires no theory of who a user is. The objection that subject keying bans the product category is answered by the statute that already governs the category: COPPA has applied this keying to ambient voice products since the Commission classed voiceprints as personal information, and the category ships through it, because consent-or-delete is an operating mode rather than a prohibition. Section 312.5(c)(9) is the Commission writing that mode down.¹²⁴

Second, supply a counting rule. A per-violation damages scheme without a unit is not enforceable at a computable value, and the range demonstrated in Part VI is the proof. Congress need not accept any particular unit; per affected individual per day is one determinate answer and there are others. What it cannot do is leave the question open, because the DMCA experiment has already run and produced opposite outcomes on similar facts followed by a holding that the question belongs to a jury.¹²⁵ The virtue of a stated unit is not that it is generous. It is that it is

124. See *supra* note 74 and accompanying text.

125. See *supra* notes 87–90 and accompanying text (discussing *Goldman* and *McClatchey*).

knowable before suit, which is the minimum condition for a private right of action to function as anything other than a wager.

Third, fix the floor to the regulator's ceiling, and index it. Section 4(c)(1)(B)(i) should read not less than the amount then in effect under section 5(m)(1)(A) of the Federal Trade Commission Act, for a violation of Section 2(a) involving the personal data of an individual under 18. That is \$53,088 today, it moves when the Commission's figure moves, and it makes the family's floor the regulator's ceiling rather than a fifth of it.¹²⁶ The consent architecture should be amended in the same stroke so that the affirmative consent that opens the Section 2(a)(1)(A) exception must come from the individual whose personal data is processed, or that individual's parent, rather than from whoever happened to ask the question. And the phrase "to inform chatbot outputs" should be struck from Section 2(a)(1)(A), so that the prohibition reaches the processing of personal data as such. As drafted, the clause invites the argument that data diverted to training, storage, or analytics was never processed to inform an output and so was never regulated at all, an argument that rewards the most extractive architecture with the least coverage.¹²⁷ Nor do the second and third amendments combine into the exposure problem Part VI documents. *Wakefield* polices aggregates rather than floors, and it polices them through a proportionality review that is administrable only where the unit and the per-unit figure are knowable before suit. A stated unit and an indexed floor are the preconditions of that review; the indeterminate scheme is what defeats it.¹²⁸

X. Conclusion

The People-First Chatbot Act is a serious bill. It bans advertising built on the contents of private conversation. It bans training on children's data. It creates a private right of action where federal children's privacy law has offered none for twenty-eight years. Those are not small things and this Article does not propose to undo them.

126. See *supra* note 108 and accompanying text. Indexing by cross-reference rather than by number also spares the provision the fate of every statutory damages figure that Congress fixed in nominal dollars and never revisited.

127. See *supra* Part IV.B.

128. See *supra* note 100 and accompanying text.

But the bill was drafted for a machine that a person picks up. It regulates the relationship between an individual and a device that individual chose to address, and it does that competently. The machines now entering American homes do not wait to be addressed. They occupy rooms. They watch and listen across the entire population of the room, and the members of that population who did not choose to be there, did not consent, and cannot leave are the smallest ones.

The Act's answer to the child in the kitchen is that she is an individual, regardless of age. That is true. It is also the whole of what the statute has to say about her, and it is not enough to tell a court whether her face belongs to her.