

**IN THE UNITED STATES DISTRICT COURT
FOR THE NORTHERN DISTRICT OF ILLINOIS
EASTERN DIVISION**

PHILIP ROGERS, CAROL MARIN,	)	
ALISON FLOWERS, ROBIN AMER,	)	
LINDSEY DORCUS, YOHANCE	)	
LACOUR, and VICTORIA NASSIF,	)	
each individually and on behalf of	)	Case No.
all others similarly situated,	)	
	)	DEMAND FOR JURY TRIAL
Plaintiffs,	)	
	)	
v.	)	
	)	
NVIDIA CORPORATION, a	)	
Delaware corporation,	)	
	)	
Defendant.	)	

COMPLAINT

Plaintiffs, Philip Rogers, Carol Marin, Alison Flowers, Robin Amer, Lindsey Dorcus, Yohance Lacour, and Victoria Nassif, individually and on behalf of all others similarly situated, by their attorneys Loevy & Loevy, and for their complaint against Defendant NVIDIA Corporation (“NVIDIA”), allege as follows:

NATURE OF THE CASE

1. NVIDIA built a suite of commercial voice AI products using the voices of real people, including some of the most accomplished broadcast journalists, investigative podcasters, and audiobook narrators working in the United States today. To build the foundational voice synthesis models that now power Magpie TTS Multilingual, Magpie TTS Zeroshot, Magpie TTS Flow, FUGATTO, PersonaPlex, Nemotron 3 VoiceChat, and the Canary and Parakeet

speech models — and to power the commercial voice services NVIDIA distributes through NVIDIA AI Enterprise, the NVIDIA API Catalog, NVIDIA NIM microservices, and open-weight model releases on Hugging Face — NVIDIA ingested hundreds of thousands of hours of human speech and extracted the unique biometric signatures, the voiceprints, of the speakers from those recordings. Plaintiffs' voices are among them.

2. Plaintiffs are seven Illinois residents whose recorded voices are among the most distinguished in their fields. They include Carol Marin, a five-decade broadcast journalist, a three-time Peabody winner, and a recipient of the 2025 Order of Lincoln, Illinois's highest civilian honor; Yohance Lacour, whose investigative podcast *You Didn't See Nothin'* won the 2024 Pulitzer Prize for Audio Reporting; Alison Flowers, a 2021 Pulitzer finalist for the investigative podcast *Somebody*; Robin Amer, a three-time duPont–Columbia honoree and creator of USA Today's *The City*; Philip Rogers, an Emmy- and Murrow Award-winning broadcaster who covered four decades of Chicago news; Lindsey Dorcus, a SOVAS and Independent Audiobook Award-winning narrator of more than two hundred audiobooks for the major American publishers; and Victoria Nassif, a Lebanese-Palestinian American audiobook narrator and actor whose Arabic-accented narrations of works by Arab and Palestinian American authors have been commercially released by Penguin Random House, Hachette, and Simon & Schuster. None of them was told that their voice was being used to train NVIDIA's commercial voice AI. None of them was asked. None of them consented.

3. A voiceprint is a digital fingerprint of the human voice. It is a mathematical representation of the acoustic features — pitch, timbre, resonance — that arise from a person's distinctive physiology, combined with the speech patterns developed over a lifetime: accent, cadence, articulation. Like a fingerprint, a voiceprint identifies the individual and cannot be changed. A Social Security number can be reissued. A credit card can be canceled. A person whose voiceprint has been taken cannot recover it by altering their voice — the biological and behavioral patterns that produced the voiceprint are the same ones used to speak every day.

4. The Illinois General Assembly enacted the Biometric Information Privacy Act, 740 ILCS 14/1 *et seq.* ("BIPA"), to address this very danger. BIPA recognizes that biometric identifiers, expressly including voiceprints, are "biologically unique to the individual" and that, once compromised, "the individual has no recourse." 740 ILCS 14/5(c). Before any private entity may collect a voiceprint, BIPA requires written notice of the specific purpose and duration of collection, along with a written release. 740 ILCS 14/15(b). NVIDIA failed to comply with any of those requirements with respect to the Plaintiffs.

5. NVIDIA's noncompliance was not a misreading of BIPA. Its commercial terms, model cards, and public statements show that it understood that voice synthesis training implicates the legal rights of the source speakers. NVIDIA's Product Specific Terms for AI Products define a "Custom TTS Application" as "a custom text-to-speech application that enables the creation of synthetic voice output that resembles an input voice," and require NVIDIA's

enterprise customers to "have sufficient rights and licenses for content used with a Custom TTS Application to generate new content." NVIDIA's October 2025 Magpie TTS launch advised users that voice cloning should be used only with "target speakers who have consented to such use." NVIDIA's Trustworthy AI terms, incorporated by reference into the licenses under which NVIDIA distributes its open-weight voice models, expressly prohibit use of those models for "unlawful biometric data collection." NVIDIA placed the consent obligation on its downstream customers and end users, and built nothing comparable for itself regarding the voices it used to train the foundational models in the first place.

6. Nor was NVIDIA's noncompliance a failure of safety capability. When NVIDIA recognized that its FUGATTO audio model — a 2.5-billion-parameter generative model trained on more than 50,000 hours of audio and capable of generating and transforming human voices — posed risks of voice impersonation and deepfake misuse, it withheld the model from public release. When NVIDIA released Magpie TTS Zeroshot and Magpie TTS Flow in October 2025, it announced a collaboration with Pindrop, a voice-fraud detection company, to develop deepfake-detection capabilities for NVIDIA's models. When NVIDIA released PersonaPlex in January 2026, it architecturally restricted the model to a fixed set of eighteen pre-packaged voice embeddings rather than supporting arbitrary voice cloning from user-supplied audio. NVIDIA built each of these safeguards to prevent the downstream misuse of its products. It built

none of them for the people whose voices it took to train those products in the first place.

7. NVIDIA's noncompliance was, instead, a deliberate institutional decision. NVIDIA had every commercial reason to know that voice training extracts biometric identifiers. NVIDIA's published research and product documentation confirm this. NVIDIA's Magpie TTS Flow model card describes the model as analyzing "a speaker's voice" to "replicate voice qualities such as pitch, timbre and speech rate" and as "[m]aintain[ing] the original characteristics that capture unique voice audio signature." NVIDIA's PersonaPlex model card describes the system as conditioned on "a voice prompt" consisting of "audio tokens that establish the target vocal characteristics and speaking style." Compliance with BIPA would have required NVIDIA to identify the source speakers, provide written notice of the specific purpose and duration of collection, and obtain a written release from each speaker before ingesting that speaker's recording into the training pipeline. With foundational models trained on a corpus measured in hundreds of thousands of hours of speech and millions of distinct speakers, that compliance burden would have constrained the speed and scale of NVIDIA's voice AI development. NVIDIA chose speed and scale over compliance. That was not an oversight. It was a business decision.¹

¹ NVIDIA's published research describes the architecture by which voice characteristics are extracted from speech recordings and encoded into the parameters of its voice models. NVIDIA's NeMo and Riva pipelines tokenize audio waveforms into discrete representations that Magpie's documentation describes as capturing "speaker characteristics," and Magpie's preference-alignment training uses "speaker verification (SV) models within a reward system to create a preference dataset" — that is, computational comparison of speaker acoustic

8. The voiceprints NVIDIA extracted from Plaintiffs are not stored in a database that can be deleted on request. They are encoded in the parameters of NVIDIA's commercial voice models and reproduced in the audio those models generate. NVIDIA has also published the parameters of several of those models — including PersonaPlex, Canary, and Parakeet — as open-weight releases on Hugging Face under the NVIDIA AI Open Model License and the NVIDIA Nemotron Open Model License, available for download by any party worldwide. At this point, the biometric data and the product are the same thing. The product has already been distributed.

9. The technology NVIDIA built using Plaintiffs' voices now competes with Plaintiffs in the markets where they earn their living. NVIDIA markets Magpie TTS Flow as "ideal for studio dubbing and podcast narration" — directly displacing the audiobook narration Plaintiffs Dorcus and Nassif provide and the long-form investigative audio journalism Plaintiffs Lacour, Flowers, and Amer have spent their careers developing. Magpie TTS Zeroshot can replicate a target speaker's voice from a five-second audio sample. Each of these products is offered to enterprise customers at a fraction of the cost of human narration. Each was built using the vocal characteristics of the human performers it now

signatures during training. See NVIDIA, *Magpie TTS Flow Model Card*, <https://web.archive.org/web/20260414085841/https://build.nvidia.com/nvidia/magpie-tts-flow/modelcard> (last visited May 12, 2026); NVIDIA, *PersonaPlex Model Card*, <https://huggingface.co/nvidia/personaplex-7b-v1> (last visited May 12, 2026); NVIDIA Developer Blog, *Enhancing Multilingual Human-Like Speech and Voice Cloning with NVIDIA Riva TTS* (Oct. 9, 2025), <https://developer.nvidia.com/blog/enhancing-multilingual-human-like-speech-and-voice-cloning-with-nvidia-riva-tts/> (last visited May 12, 2026). Plaintiffs allege on information and belief that the same architecture and processing pipeline was applied to the voice training data ingested for the foundational models that power NVIDIA's commercial voice products.

displaces, including, on information and belief, the vocal characteristics of every Plaintiff in this case.

10. Plaintiffs' injuries are concrete and particularized. NVIDIA extracted Plaintiffs' voiceprints without notice or consent, depriving them of the right BIPA guarantees to make an informed decision about the collection and use of their biometric data. NVIDIA retains those voiceprints in its commercial models and continues to profit from them. NVIDIA has further disseminated those voiceprints, encoded in model parameters, through its open-weight model releases on Hugging Face. The voiceprints cannot be recovered or replaced. The technology built on those voiceprints now displaces Plaintiffs in the markets where they earn their living.

11. Plaintiffs bring this action under BIPA, 740 ILCS 14/15(a)–(e), alleging that NVIDIA unlawfully collected, retained, commercialized, and disseminated their voiceprints, failed to protect them from disclosure, and did so without notice, informed written consent, a written release, or any publicly available retention and destruction policy applicable to non-users. Plaintiffs also assert that NVIDIA's commercial use of their voices and identities to build and sell AI products that mimic them violates the Illinois Right of Publicity Act (“IRPA”). Plaintiffs further assert claims under the Illinois Consumer Fraud and Deceptive Business Practices Act (“ICFA”), the Illinois Uniform Deceptive Trade Practices Act (“IUDTPA”), and the common law of unjust enrichment.

12. Plaintiffs seek (i) statutory damages under 740 ILCS 14/20 for violations of BIPA's notice, consent, retention, profiting, dissemination, and

protection requirements; (ii) actual damages and disgorgement of profits NVIDIA has earned from the commercial exploitation of Plaintiffs' biometric data; (iii) injunctive relief requiring NVIDIA to (a) cease collecting biometric identifiers from voice recordings produced or recorded in Illinois without BIPA-compliant consent, (b) identify the sources of the voice training data used to build its foundational voice synthesis models, (c) destroy all voiceprints and biometric information unlawfully obtained from Plaintiffs and the class, and (d) destroy or retrain, without the unlawfully obtained data, the foundational voice models and downstream commercial products that contain that data; and (iv) reasonable attorneys' fees, costs, and expenses.

PARTIES

13. Plaintiff Philip Rogers ("Rogers") is a citizen of Illinois and resides in this District. Rogers is a broadcast journalist whose four-decade career was conducted in and from Chicago, primarily at WBBM Newsradio (CBS) and WMAQ-TV (NBC 5 Chicago). Rogers's body of professional voice work, the public availability of his recordings, and the basis for Plaintiffs' allegation that NVIDIA extracted his voiceprint are described at ¶¶ 89-94.

14. Plaintiff Carol Marin ("Marin") is a citizen of Illinois and resides in this District. Marin is a broadcast journalist whose career has been conducted primarily in Chicago. Marin's body of professional voice work, the public availability of her recordings, and the basis for Plaintiffs' allegation that NVIDIA extracted her voiceprint are described at ¶¶ 95-99.

15. Plaintiff Alison Flowers ("Flowers") is a citizen of Illinois and resides in this District. Flowers is an investigative journalist and audio producer who produces her audio reporting from Chicago through her production company Spiralbound. Flowers's body of professional voice work, the public availability of her recordings, and the basis for Plaintiffs' allegation that NVIDIA extracted her voiceprint are described at ¶¶ 100-104.

16. Plaintiff Robin Amer ("Amer") is a citizen of Illinois and resides in this District. Amer is a journalist, podcast creator, audio producer, and on-air host whose work has been produced substantially in and from Chicago, including as creator and host of USA Today's *The City* and as Managing Editor of *Love + Radio*. Amer's body of professional voice work, the public availability of her recordings, and the basis for Plaintiffs' allegation that NVIDIA extracted her voiceprint are described at ¶¶ 105-109.

17. Plaintiff Lindsey Dorcus ("Dorcus") is a citizen of Illinois and resides in this District. Dorcus is a professional audiobook narrator who has recorded more than 200 audiobooks for major American publishers from her home recording studio in Chicago. Dorcus's body of professional voice work, the public availability of her recordings, and the basis for Plaintiffs' allegation that NVIDIA extracted her voiceprint are described at ¶¶ 110-114.

18. Plaintiff Yohance Lacour ("Lacour") is a citizen of Illinois and resides in this District. Lacour is a journalist and audio storyteller whose investigative podcast work, including the Pulitzer Prize-winning *You Didn't See Nothin'*, is produced in Chicago. Lacour's body of professional voice work, the

public availability of his recordings, and the basis for Plaintiffs' allegation that NVIDIA extracted his voiceprint are described at ¶¶ 115-119.

19. Plaintiff Victoria Nassif ("Nassif") is a citizen of Illinois and resides in this District. Nassif is a first-generation Lebanese-Palestinian American actor, audiobook narrator, voiceover artist, and intimacy director whose professional voice work is produced primarily in Illinois. Nassif's body of professional voice work, the public availability of her recordings, and the basis for Plaintiffs' allegation that NVIDIA extracted her voiceprint are described at ¶¶ 120-124.

20. Defendant NVIDIA Corporation ("NVIDIA") is a Delaware corporation with its principal executive offices at 2788 San Tomas Expressway, Santa Clara, California 95051.

JURISDICTION AND VENUE

21. This Court has subject-matter jurisdiction over this action under the Class Action Fairness Act, 28 U.S.C. § 1332(d). The amount in controversy exceeds \$5,000,000 in the aggregate, exclusive of interest and costs. The proposed Class includes more than 100 members. Minimal diversity is satisfied: Plaintiffs are citizens of Illinois, and Defendant NVIDIA Corporation is a Delaware corporation with its principal place of business in California. None of the exceptions to CAFA jurisdiction set forth in 28 U.S.C. § 1332(d)(3)–(5) applies.

22. The aggregate amount in controversy substantially exceeds \$5,000,000. BIPA provides that an aggrieved person may recover the greater of

liquidated damages or actual damages — \$1,000 for a negligent violation or \$5,000 for an intentional or reckless violation — and may recover those statutory damages on a per-person, per-subsection basis where multiple distinct provisions of § 15 are violated, consistent with the statute as amended and as construed in *Clay v. Union Pacific Railroad Co.*, No. 25-2185 (7th Cir. Apr. 1, 2026). Plaintiffs seek recovery under five distinct BIPA subsections — § 15(a), (b), (c), (d), and (e) — each of which creates a distinct duty and supports a distinct per-person statutory or actual damages recovery, as well as recovery under IRPA, ICFA, IUDTPA, and common-law unjust enrichment.

23. This Court has specific personal jurisdiction over NVIDIA. NVIDIA has purposefully directed its conduct at Illinois and this District in ways that bear directly on Plaintiffs' claims. NVIDIA markets, sells, and delivers to Illinois customers the voice AI products at issue in this Complaint — including Magpie TTS Multilingual, Magpie TTS Zeroshot, Magpie TTS Flow, FUGATTO, PersonaPlex, Nemotron 3 VoiceChat, Canary, and Parakeet — through NVIDIA AI Enterprise (a paid enterprise software subscription), the NVIDIA API Catalog at build.nvidia.com (a cloud-hosted inference service offered on a pay-for-use basis), NVIDIA NIM containerized microservices (for production deployment by enterprise customers), the NVIDIA NeMo open-source toolkit on GitHub, and open-weight model releases on Hugging Face. NVIDIA enters into recurring enterprise license agreements with Illinois-based businesses for these products, accepts subscription and API usage payments from Illinois customers, and provides ongoing service performance — including delivery of

voice synthesis outputs in response to user requests — to users located in Illinois. On information and belief, NVIDIA has substantial enterprise customer relationships in Illinois, including with Fortune 500 companies headquartered in this District and other Illinois-based enterprises, generating substantial Illinois revenue.

24. Plaintiffs' claims arise directly from NVIDIA's Illinois-directed conduct. NVIDIA extracted the voiceprints of Plaintiffs and the Class, persons whose voice recordings were produced or recorded in Illinois, without the notice, consent, or written release BIPA requires. The voiceprints are now encoded in the parameters of NVIDIA's commercial voice synthesis models, which NVIDIA sells, delivers, and disseminates to Illinois customers and Illinois users through the channels identified above. Every commercial transaction NVIDIA completes in Illinois involving those products monetizes the biometric data NVIDIA obtained from Plaintiffs and the Class without their consent. Each open-weight model release that places NVIDIA's voice synthesis models on Hugging Face for download by any party — including Illinois developers, Illinois enterprises, and Illinois users — constitutes a continuing dissemination of biometric data originating in Illinois. NVIDIA's Illinois contacts are not incidental to this Complaint; they are the downstream commercialization and the continuing dissemination of the upstream biometric extraction that gives rise to it.

25. Exercising specific personal jurisdiction over NVIDIA in this District is consistent with the Illinois long-arm statute, 735 ILCS 5/2-209, and

with due process requirements. Illinois has a strong interest in providing a forum for redress of unlawful biometric data collection from voice recordings produced or recorded in Illinois, an interest the Illinois General Assembly expressly identified in enacting BIPA, 740 ILCS 14/5. Plaintiffs, as Illinois residents whose biometric privacy was invaded in Illinois and whose recordings were produced or recorded in Illinois, have a corresponding interest in litigating their claims in their home forum. The burden on NVIDIA of litigating in Illinois is not undue, given NVIDIA's substantial Illinois commercial activity, its resources as a publicly traded technology company with fiscal-year-2026 revenue of approximately \$215.94 billion and market capitalization exceeding \$5.4 trillion, and its experience as an active litigant in federal courts across the United States. NVIDIA could reasonably anticipate being haled into Illinois courts for claims like Plaintiffs' that arise from its extraction and commercial exploitation of biometric identifiers from voice recordings produced or recorded in Illinois.

26. Venue is proper in this District under 28 U.S.C. § 1391(b)(2) because a substantial part of the events giving rise to Plaintiffs' claims occurred here: Plaintiffs are residents of this District, their recorded vocal performances were produced and distributed from this District, the biometric privacy violations Plaintiffs allege occurred in this District, and the commercial exploitation of Plaintiffs' voiceprints continues through NVIDIA's sale and delivery of the voice AI products at issue to customers in this District. Venue is independently proper under 28 U.S.C. § 1391(b)(1) because Defendant NVIDIA

is subject to personal jurisdiction in this District for the reasons set forth above and therefore "resides" in this District for venue purposes under 28 U.S.C. § 1391(c)(2).

FACTUAL BACKGROUND

NVIDIA and Its Voice Business

27. NVIDIA is the world's largest publicly traded technology company by market capitalization, and is the dominant supplier of the graphics processing units, AI accelerators, and AI computing platforms that power large-scale machine learning, including the training and deployment of generative AI models. NVIDIA's commercial AI business spans GPU and hardware sales, AI Enterprise software subscriptions, cloud-hosted inference services, containerized microservice deployments, and open-weight model releases.

28. NVIDIA's voice AI research is conducted principally through its Applied Deep Learning Research group ("ADLR"), an internal research organization that has specialized in generative models for audio, speech, language, and vision since at least 2018. ADLR and affiliated NVIDIA research groups have published peer-reviewed research and released a sustained line of voice synthesis, voice cloning, and voice-conversion models, beginning with WaveGlow (2018) and continuing through Mellotron (2019), Flowtron (2020), FastPitch (2020), RAD-TTS (2021–2022), Mixer-TTS (2022), Incremental FastPitch (2023–2024), and the commercial models at issue in this Complaint. ADLR's sustained focus on voice synthesis research, extending back more than seven years before NVIDIA's October 2025 release of Magpie TTS Zeroshot and

Magpie TTS Flow, demonstrates that NVIDIA's entry into commercial voice cloning was not incidental but the product of a deliberate, long-term institutional investment.

29. NVIDIA has identified its voice AI products and the underlying foundational voice synthesis models as core components of its commercial AI strategy in disclosures to the United States Securities and Exchange Commission and in public statements by its Chief Executive Officer, Jensen Huang. Huang has personally identified NVIDIA's commercial voice AI capabilities, including the Magpie TTS family and PersonaPlex, in NVIDIA's quarterly earnings calls and in keynote presentations at the GPU Technology Conference, NVIDIA's global AI conference.

30. NVIDIA distributes its commercial voice AI products and the underlying foundational voice synthesis models through an integrated commercial channel that NVIDIA operates from the United States, including through NVIDIA AI Enterprise (a paid enterprise software subscription), the NVIDIA API Catalog at build.nvidia.com (a cloud-hosted inference service offered on a pay-for-use basis), NVIDIA NIM containerized microservices (for production deployment by enterprise customers), the NVIDIA NeMo open-source toolkit on GitHub, and open-weight model releases on Hugging Face under the NVIDIA AI Open Model License and the NVIDIA Nemotron Open Model License. The voice AI research, model training, and commercial deployment alleged in this Complaint were carried out by NVIDIA personnel in

the United States, on infrastructure operated by NVIDIA and its affiliated entities, in the ordinary course of NVIDIA's business.

31. NVIDIA's commercial voice AI products at issue in this Complaint include the following:

(a) The Magpie TTS family, released in October 2025, comprising Magpie TTS Multilingual (which synthesizes speech in multiple languages with emotion control), Magpie TTS Zeroshot (which enables voice cloning from a five-second audio sample, achieving "speaker similarity of over 70%"), and Magpie TTS Flow (a higher-fidelity voice cloning and synthesis model NVIDIA describes as "ideal for studio dubbing and podcast narration" and which NVIDIA trained on approximately 70,000 hours of paired speech data).

(b) FUGATTO, announced in November 2024, a 2.5-billion-parameter foundational audio generative model trained on more than 50,000 hours of audio and capable of generating and transforming human voices, including by modifying accent and emotion. NVIDIA withheld FUGATTO from public release in part because of concerns about misuse, including voice impersonation and deepfake generation.

(c) PersonaPlex, released in January 2026, a real-time full-duplex speech-to-speech model with voice and persona control, distributed as an open-weight model on Hugging Face under the NVIDIA AI Open Model License.

(d) Nemotron 3 VoiceChat, a 12-billion-parameter end-to-end full-duplex voice model integrating a conformer audio encoder, a Nemotron Nano V2 9B language model backbone, and a TTS decoder.

(e) Canary-1b-v2 and Parakeet-tdt-0.6b-v3, speech models trained on the Granary speech corpus and distributed as open-weight models on Hugging Face.

32. On information and belief, the training of NVIDIA's foundational voice synthesis models occurred on NVIDIA-operated computing infrastructure in the United States, including on GPU clusters owned and operated by NVIDIA. The production voice models that NVIDIA deploys through its commercial products are hosted on NVIDIA-controlled cloud infrastructure and served to United States users, including users in Illinois, from data center facilities operated by NVIDIA and its cloud and content-delivery partners. NVIDIA's voice AI revenue is reported in the Data Center segment of its consolidated financials.

How NVIDIA's Voice AI Extracts and Encodes Voiceprints

33. Modern AI voice synthesis works by training neural networks on large quantities of recorded human speech. During training, the network learns to identify and reproduce the acoustic features that make individual voices distinctive — pitch, timbre, resonance, accent, cadence, articulation, and the dynamics of emotional expression. The network encodes those features as mathematical representations and stores them in the model's parameters. The same parameters are then used to generate new speech that exhibits the acoustic features of the training voices.

34. NVIDIA's published research describes this process in technical detail. NVIDIA's WaveGlow paper, published in 2018, introduced a flow-based

generative network for raw audio waveform synthesis, trained directly on speech recordings.²

35. NVIDIA's RAD-TTS line of research, published in 2021–2022, introduced the disentanglement of voice characteristics from linguistic content — the technical mechanism by which the same model can generate the same words in different speakers' voices, and the same speaker's voice saying different words.³

36. NVIDIA's October 2025 Riva TTS blog post, describing the Magpie TTS training pipeline, states that Magpie's training process uses "speaker verification (SV) models within a reward system to create a preference dataset" — that is, computational comparison of speaker acoustic signatures during training.⁴

² Ryan Prenger, Rafael Valle & Bryan Catanzaro, WaveGlow: A Flow-based Generative Network for Speech Synthesis, arXiv:1811.00002 (Oct. 31, 2018), <https://arxiv.org/abs/1811.00002> (last visited May 12, 2026); see also NVIDIA Applied Deep Learning Research, WaveGlow: A Flow-based Generative Network for Speech Synthesis, <https://nv-adlr.github.io/WaveGlow> (last visited May 12, 2026); NVIDIA, waveglow (GitHub repository), <https://github.com/NVIDIA/waveglow> (last visited May 12, 2026).

³ Kevin J. Shih, Rafael Valle, Rohan Badlani, Adrian Łańcucki, Wei Ping & Bryan Catanzaro, RAD-TTS: Parallel Flow-Based TTS with Robust Alignment Learning and Diverse Synthesis, ICML Workshop on Invertible Neural Networks, Normalizing Flows, and Explicit Likelihood Models (2021), <https://openreview.net/forum?id=ONQwnnwAORi> (last visited May 12, 2026); see also NVIDIA Applied Deep Learning Research, RAD-TTS: Parallel Flow-Based TTS with Robust Alignment Learning and Diverse Synthesis, <https://nv-adlr.github.io/RADTTS> (last visited May 12, 2026); Rohan Badlani, Rafael Valle, Kevin J. Shih, João Felipe Santos, Siddharth Gururani & Bryan Catanzaro, Multilingual Multiaccented Multispeaker TTS with RADTTS, arXiv:2301.10335 (Jan. 24, 2023), <https://arxiv.org/abs/2301.10335> (last visited May 12, 2026); NVIDIA Applied Deep Learning Research, RADMMM: Multilingual Multiaccented Multispeaker TTS with RADTTS, <https://research.nvidia.com/labs/adlr/projects/radmmm/> (last visited May 12, 2026).

⁴ NVIDIA, Enhancing Multilingual Human-Like Speech and Voice Cloning with NVIDIA Riva TTS, NVIDIA Technical Blog (Oct. 9, 2025), <https://developer.nvidia.com/blog/enhancing-multilingual-human-like-speech-and-voice-cloning-with-nvidia-riva-tts/> (last visited May 12, 2026).

37. NVIDIA's product documentation confirms that the speaker-identifying output of this pipeline is, in form and substance, a voiceprint. NVIDIA's Magpie TTS Flow model card describes the model as analyzing "a speaker's voice" to "replicate voice qualities such as pitch, timbre and speech rate to achieve a speaker similarity of over 70%" from a five-second audio prompt, and as "[m]aintain[ing] the original characteristics that capture unique voice audio signature."⁵

38. NVIDIA's PersonaPlex model card describes the system as conditioned on "a voice prompt" consisting of "audio tokens that establish the target vocal characteristics and speaking style." A representation of voice data sufficient to clone the speaker's voice and to maintain the speaker's unique audio signature is, in any ordinary understanding of the term, a voiceprint.⁶

39. BIPA's definitions confirm the legal classification. The statute defines "biometric identifier" to include "voiceprint." 740 ILCS 14/10. It separately defines "biometric information" to include "any information, regardless of how it is captured, converted, stored, or shared, based on an

⁵ NVIDIA, Magpie TTS Flow Model Card, <https://web.archive.org/web/20260414085841/https://build.nvidia.com/nvidia/magpie-tts-flow/modelcard> (last visited May 12, 2026); see also NVIDIA, Speech Synthesis: Multilingual Zero-shot Voice Characterization — A2-Flow Model Overview, https://developer.download.nvidia.com/assets/ace/model_card/Riva_TTS_A2-Flow_for_Nv_IGI_SDK.pdf (last visited May 12, 2026) (technical model card for the A2-Flow architecture deployed in production as Magpie TTS Flow, describing the model's analysis of a speaker's voice, replication of pitch, timbre, and speech rate, speaker similarity exceeding 70% from a five-second audio prompt, and retention of the speaker's unique audio signature).

⁶ NVIDIA, PersonaPlex-7B-v1 Model Card, Hugging Face, <https://huggingface.co/nvidia/personaplex-7b-v1> (last visited May 12, 2026); see also NVIDIA, personaplex-7b-v1 README.md, <https://huggingface.co/nvidia/personaplex-7b-v1/blob/main/README.md> (last visited May 12, 2026).

individual's biometric identifier used to identify an individual." *Id.* The audio tokens, speaker embeddings, voice prompts, and speaker-acoustic-signature representations produced by NVIDIA's voice synthesis pipeline fall within both definitions. The label the AI research community uses for these representations does not determine their statutory classification; the function of the data does.

40. NVIDIA's Product Specific Terms for AI Products confirm this functional understanding. Those terms define a "Custom TTS Application" as "a custom text-to-speech application that enables the creation of synthetic voice output that resembles an input voice." NVIDIA's contractual definition acknowledges that NVIDIA's voice cloning capability is built on the extraction and encoding of speaker-identifying acoustic characteristics — exactly the capability the Magpie TTS Flow and PersonaPlex documentation describes, and exactly the capability that BIPA's voiceprint and biometric information definitions reach.

41. On information and belief, the same category of biometric processing that occurs in NVIDIA's consumer-facing voice-cloning features — Magpie TTS Zeroshot, Magpie TTS Flow, and PersonaPlex — also occurs during base-model training. The consumer-facing features are applications of capabilities built during training. The training pipeline is where extraction of voiceprints and biometric information occurs first, at the largest scale, and on the broadest set of voice recordings — including, on information and belief, the recordings of Plaintiffs and the Class.

*NVIDIA Built Safeguards Against Misuse,
Not Consent for the Voices It Took*

42. NVIDIA understands that voice synthesis technology built on human voice data creates the categories of harm that BIPA was enacted to prevent. NVIDIA has built and publicly deployed five distinct categories of safeguard measures over the period 2024–2026, each directed at the downstream misuse of NVIDIA's commercial voice products. NVIDIA built none of those measures for the upstream voices it took to construct those products in the first place.

43. In November 2024, NVIDIA announced FUGATTO, a 2.5-billion-parameter foundational audio generative model trained on more than 50,000 hours of audio and capable of generating and transforming human voices, including by modifying accent and emotion. NVIDIA did not release FUGATTO publicly. Industry reporting, citing NVIDIA, confirms that NVIDIA withheld public release in part because of concerns about misuse, including voice impersonation and deepfake generation.

44. In October 2025, when NVIDIA released Magpie TTS Zeroshot and Magpie TTS Flow — voice cloning models that can replicate a target speaker's voice from a five-second audio sample — NVIDIA simultaneously announced a collaboration with Pindrop, a voice-fraud detection company, to develop deepfake detection capabilities against NVIDIA's voice models. NVIDIA further advised users that voice cloning should be used only with "target speakers who have consented to such use."

45. In January 2026, when NVIDIA released the PersonaPlex conversational voice model, NVIDIA architecturally restricted the model to a fixed set of eighteen pre-packaged voice embeddings rather than supporting arbitrary voice cloning from user-supplied audio.

46. NVIDIA's Product Specific Terms for AI Products define a "Custom TTS Application" as "a custom text-to-speech application that enables the creation of synthetic voice output that resembles an input voice," and require NVIDIA's enterprise customers to "have sufficient rights and licenses for content used with a Custom TTS Application to generate new content."

47. NVIDIA's Trustworthy AI terms, incorporated by reference into the licenses under which NVIDIA distributes its open-weight voice models, expressly prohibit use of NVIDIA's models for "unlawful biometric data collection."

48. Each of the five measures described in ¶¶ 42 through 47 is a constraint at the downstream end of NVIDIA's commercial pipeline, directed at restricting how end users use NVIDIA's products, or at imposing the upstream-rights obligation on NVIDIA's enterprise customers. None of the measures is directed at the upstream issue: that NVIDIA's foundational voice synthesis models were trained on voiceprints extracted without the knowledge or consent of the speakers whose voice recordings were ingested during training. NVIDIA built each of these safeguards against the downstream misuse of its products. It built none of them for the people whose voices it took to train those products in the first place.

NVIDIA Trained on Data It Knew It Had No Right to Use

49. NVIDIA has never published a comprehensive disclosure of the sources, scale, or provenance of the voice data used to train its commercial voice synthesis models.

50. It has not published a model card, data sheet, training data manifest, licensing inventory, or transparency report identifying all of the voice recordings used to train Magpie TTS Flow (approximately 70,000 hours of paired speech data), FUGATTO (more than 50,000 hours of audio), or Nemotron 3 VoiceChat (unquantified in publicly available materials), where it obtained those recordings, or whether any of the speakers consented.

51. Although NVIDIA has identified certain narrower datasets — including LJSpeech, LibriTTS, Hi-Fi TTS, Mozilla Common Voice, and the Fisher English Corpus — in specific model cards and research papers, no equivalent public-domain or licensed-sourcing claim has been made for the much larger corpora on which NVIDIA's commercial voice models were trained.

52. NVIDIA's silence about the bulk of its voice training data is made more conspicuous by the contrast with its disclosures for narrower and smaller voice datasets. In April 2021, NVIDIA published a research paper describing its Hi-Fi TTS dataset as drawn from LibriVox audiobooks and Project Gutenberg texts, public-domain sources that NVIDIA specifically identified.

53. NVIDIA's Magpie TTS Multilingual model card identifies LibriTTS, an open, publicly available speech dataset, as one of the datasets used to train the model.⁷

54. But no equivalent public-domain or licensed-sourcing claim has ever been made for the approximately 70,000 hours of paired speech data on which the released Magpie TTS Flow model was trained, or for the more than 50,000 hours of audio on which FUGATTO was trained. If NVIDIA's commercial voice models were trained entirely on properly licensed or consented data, NVIDIA would have an obvious commercial incentive to say so. It has not.

55. The contemporaneous public record establishes that NVIDIA obtained substantial portions of its training data, across multiple foundational-model families, from publicly accessible internet sources without the consent of the source speakers, authors, or content creators, and that NVIDIA did so at industrial scale.

56. In August 2024, 404 Media reported on internal NVIDIA communications describing a company-wide initiative to aggregate "a huge curated video dataset" for generative modeling that processed approximately "a human lifetime" of video content per day. The communications describe a sustained, systematized effort to ingest internet-accessible video content at industrial scale, conducted across NVIDIA's research and engineering organizations.

⁷ <https://build.nvidia.com/nvidia/magpie-tts-multilingual/modelcard> (last visited on May 12, 2026)

57. NVIDIA is the defendant in two pending putative class actions concerning the training-data provenance of its other foundational AI models. In *Nazemian v. NVIDIA Corp.*, No. 4:24-cv-01454-JST (N.D. Cal. filed Mar. 8, 2024), three professional authors allege that NVIDIA trained its NeMo Megatron large language models on the Books3 dataset, a collection of approximately 196,640 pirated e-books drawn from the Bibliotik shadow library, without notice, consent, or compensation. The court permitted leave to file an amended complaint adding additional NeMo Megatron model families in January 2026. In *Ted Entertainment, Inc. v. NVIDIA Corp.*, No. 5:25-cv-10287 (N.D. Cal. filed Nov. 26, 2025), the operators of YouTube channels including H3H3 Productions and MrShortGame Golf, and Golfholics, Inc., allege that NVIDIA bypassed YouTube's technical protections to scrape millions of YouTube videos to train its Cosmos foundational video model, in violation of YouTube's terms of service and the anti-circumvention provisions of the Digital Millennium Copyright Act.⁸

58. In August 2025, NVIDIA publicly released the Granary dataset, an approximately one-million-hour corpus of speech audio assembled by processing pre-existing open-source corpora — including YODAS and YouTube-Commons (both YouTube-derived), VoxPopuli (European Parliament proceedings), and LibriLight (LibriVox audiobooks) — through NVIDIA's NeMo Speech Data Processor pipeline. NVIDIA used Granary to train its Canary-1b-

⁸ Plaintiffs cite the public allegations in *Nazemian* and *Ted Entertainment* as circumstantial evidence of NVIDIA's training-data acquisition practices across its foundational model families, not as adjudicated fact.

v2 and Parakeet-tdt-0.6b-v3 speech models, which NVIDIA distributes as open-weight models on Hugging Face.

59. The Granary release demonstrates, in NVIDIA's public documentation, that NVIDIA maintains the infrastructure for the large-scale ingestion and processing of human voice audio from publicly accessible platforms; that NVIDIA willingly ingests identifiable voice recordings from those platforms without obtaining the consent of the speakers whose voices appear in the recordings; and that NVIDIA's voice-audio processing pipeline operates at a scale, approximately one million hours, well in excess of what is necessary to produce the commercial voice cloning models at issue in this case.

60. NVIDIA has not published, and on information and belief does not maintain, any mechanism that would allow speakers like Plaintiffs, whose voices were ingested into NVIDIA's foundational voice synthesis models, to discover that fact, withdraw consent, or request deletion. NVIDIA's general Privacy Policy provides certain rights to users with an existing relationship with NVIDIA, but it does not address voice data used in AI model training and does not apply to non-users like Plaintiffs.

61. NVIDIA trained its foundational voice models on a training corpus it has refused to identify. NVIDIA's contemporaneous training-data acquisition practices, documented in internal communications and pending litigation across multiple foundational-model families, reflect a sustained institutional pattern of ingesting identifiable source material at industrial scale without the consent of the people whose work the source material represents. NVIDIA

operates the infrastructure required to do so for voice audio, has used it for the open-source-derived Granary corpus, and has provided no mechanism by which the speakers whose voices were ingested can verify, withdraw consent, or request deletion. Plaintiffs allege, on the basis of these public facts, that NVIDIA trained its foundational voice synthesis models on voice recordings that included Plaintiffs' recordings, obtained from publicly accessible platforms, and extracted voiceprints from those recordings without the knowledge or consent of the speakers.

The Biometric Data Was Generated in Illinois

62. The conduct giving rise to Plaintiffs' claims is localized in Illinois in three independent and mutually reinforcing respects: (i) the biometric data at issue was generated in Illinois; (ii) NVIDIA's acquisition of that biometric data targeted material that was identifiably Illinois-origin; and (iii) NVIDIA's ongoing retention, dissemination, and commercial exploitation of the biometric data is directed at, or felt in, Illinois. Each ground supports the application of BIPA, IRPA, and the related Illinois statutes invoked in this Complaint to NVIDIA's conduct as to Plaintiffs and the Class.

63. Plaintiffs' voices, the biological characteristics from which voiceprints and biometric information are derived, were produced by Plaintiffs while they were physically present in Illinois.

64. The voice itself, as a unique biometric signature, came into existence in Illinois.

65. The audio recordings embodying Plaintiffs' voices, from which NVIDIA is alleged to have extracted voiceprints and biometric information, were created in Illinois — Plaintiffs were physically present in Illinois at the time of the recordings, and the recordings captured the biometric characteristics of speakers who were in Illinois at the time of recording.

66. The recordings were published from Illinois to publicly accessible platforms, including platforms on which Plaintiffs' identity, geographic origin, professional biography, and content catalog were and are publicly visible to any person accessing the recordings. The metadata, descriptions, and contextual information accompanying the recordings publicly identified Plaintiffs as Illinois-based speakers and identified the recordings as Illinois-origin content at the time the recordings were made and continuously thereafter.

67. The biometric source material — the voice itself, the audio encoding of the voice, and the publicly distributed recordings embodying both — is therefore of Illinois origin. The biometric data NVIDIA is alleged to have extracted from those recordings is Illinois-origin, regardless of where any subsequent processing occurred.

68. NVIDIA accessed Plaintiffs' voice recordings from publicly accessible third-party platforms on which the recordings were hosted with metadata identifying each Plaintiff by name and identifying the recordings' origin. To a person with knowledge of Illinois broadcast, journalism, audiobook, podcast, and voice-performance markets, the metadata accompanying Plaintiffs' recordings identified each Plaintiff as a speaker whose work

originated from Illinois. NVIDIA's act of accessing and processing those recordings was therefore not the passive receipt of an anonymized dataset assembled by a third party. It was the affirmative acquisition of identifiable, attributed voice content — voice content that, at the time of acquisition, was publicly associated with Illinois-based speakers and Illinois-origin production. The Illinois affiliation of each Plaintiff and of each Plaintiff's recordings was knowable from publicly visible information at the time NVIDIA accessed and processed the recordings.

69. NVIDIA's commercial exploitation of voiceprints and biometric information generated in Illinois is itself Illinois-directed. NVIDIA sells subscriptions and AI-generated voice outputs to Illinois subscribers and Illinois-based commercial users on a continuing basis through NVIDIA AI Enterprise, the NVIDIA API Catalog, NVIDIA NIM microservices, NeMo, and open-weight model releases on Hugging Face and GitHub. The commercial value of those subscriptions and outputs derives, on information and belief, from voice models developed through unconsented extraction of biometric data, including biometric data generated in Illinois. Each Illinois transaction in those products is therefore an exercise of commercial value that NVIDIA derived from the same biometric extraction Plaintiffs allege is unlawful.

70. NVIDIA's continuing dissemination of the foundational voice models — including through the open-weight releases of PersonaPlex, Canary, and Parakeet on Hugging Face, available for download by any person in Illinois or worldwide — is likewise Illinois-directed. The injury suffered by Plaintiffs

and the Class is felt in Illinois: the privacy interest BIPA protects — the right to control the collection, retention, and commercial use of one's biometric identity — is exercised by Plaintiffs in Illinois. Plaintiffs' loss of control over their biometric data, their loss of the licensing and consent rights BIPA preserves, and the dilution of the commercial value of their voices in markets in which they participate are all injuries Plaintiffs sustain in connection with their personal and professional activities, including activities in Illinois.

71. BIPA's notice and consent obligations under § 15(b) are duties that run to the subject whose biometric data is collected. The statute requires that the collecting entity inform "the subject" in writing and receive "a written release executed by the subject" before collection. 740 ILCS 14/15(b). The duty is owed to the subject, not to the location of the collecting entity's computational infrastructure.

72. The same is true of BIPA's retention obligations under § 15(a), profiting prohibitions under § 15(c), dissemination prohibitions under § 15(d), and reasonable-care obligations under § 15(e), each of which protects the persons from whom biometric data is taken. The subjects of the biometric data at issue in this Complaint, Plaintiffs and the Class, are persons whose voices and recordings were generated in Illinois. NVIDIA's continuing possession of voiceprints and biometric information generated in Illinois likewise runs afoul of obligations BIPA imposes for the benefit of the subjects of that data, including NVIDIA's failure to publish a retention schedule and destruction policy compliant with 740 ILCS 14/15(a) covering biometric data obtained from

non-user training-data sources, and NVIDIA's failure to provide any mechanism by which Plaintiffs or any other non-user training-data subject may seek removal or destruction of their biometric data.

NVIDIA's Monetization of the Voice Models

73. NVIDIA's commercial exploitation of the foundational voice synthesis models is not limited to a single product line. NVIDIA monetizes voice models through a vertically integrated commercial chain spanning enterprise software subscriptions, cloud-hosted inference services, containerized microservice deployments, hardware and platform sales, open-weight model distribution, and the broader strategic value of voice AI as part of NVIDIA's Data Center segment. The voice characteristics learned during training, the same characteristics extracted from Plaintiffs' and Class members' voice recordings, enable each of these revenue streams.

74. NVIDIA monetizes the voice models through, among other channels:

(a) *Enterprise software subscriptions.* NVIDIA AI Enterprise is a paid enterprise software subscription that includes Riva, the speech AI platform that packages Magpie TTS Multilingual, Magpie TTS Zeroshot, and Magpie TTS Flow. NVIDIA charges enterprise customers recurring subscription fees for AI Enterprise access. The voice synthesis capabilities bundled into AI Enterprise are derived from the foundational voice models.

(b) *Cloud-hosted inference fees.* The NVIDIA API Catalog at build.nvidia.com provides cloud-hosted inference access to NVIDIA's voice

models — including Magpie TTS, PersonaPlex, and Nemotron 3 VoiceChat — on a pay-for-use basis. Each inference request processed through the API Catalog generates revenue that NVIDIA captures because the voice models exist.

(c) *Containerized microservice deployments.* NVIDIA NIM provides containerized voice synthesis microservices for production deployment by enterprise customers. NVIDIA charges enterprise customers for NIM deployments, including the Riva TTS NIM that packages the Magpie TTS models. NIM customer deployments are recurring commercial transactions, the value of which depends on the underlying voice synthesis capabilities.

(d) *Hardware sales supported by demonstrated voice AI capability.* NVIDIA's commercial value as a hardware company depends materially on its ability to demonstrate, through state-of-the-art AI software, the value of NVIDIA-architecture compute for AI workloads. NVIDIA's voice AI products serve this purpose: by demonstrating that NVIDIA-architecture compute produces commercially capable voice synthesis, NVIDIA supports the demand for the GPUs, AI accelerators, and DGX systems that NVIDIA sells to enterprise and hyperscale customers. NVIDIA's GPU and AI hardware sales generated approximately \$197.3 billion of NVIDIA's \$215.94 billion fiscal year 2026 revenue in the Data Center segment, of which the voice AI products are an integral demonstration component.

(e) *Open-weight model distribution and ecosystem value.* NVIDIA distributes several of its voice synthesis models — including PersonaPlex, Canary, and Parakeet — as open-weight models on Hugging Face under the

NVIDIA AI Open Model License or the NVIDIA Nemotron Open Model License. Open-weight distribution is itself a monetization channel: it builds developer adoption of NVIDIA's voice AI stack, drives downstream demand for NVIDIA-architecture compute for inference and fine-tuning, supports NVIDIA's positioning as the dominant commercial AI platform, and reinforces the commercial value of NVIDIA's paid AI Enterprise tier through demonstrated capability of the underlying foundational models.

(f) Strategic deployment value across NVIDIA's commercial footprint.

NVIDIA deploys voice synthesis capability across its broader product portfolio, including in customer-facing demonstrations, partnership integrations, developer-relations programming, and industry conferences such as the GPU Technology Conference. The competitive and strategic value of best-in-class voice AI deployment across NVIDIA's product portfolio is itself a benefit NVIDIA captures from the foundational voice models.

75. Each monetization channel depends on the foundational voice models. Each foundational voice model depends on the training data used to build it. And the training data, on information and belief, includes the voice recordings of Plaintiffs and Class members from which NVIDIA extracted voiceprints without notice or consent. The commercial exploitation NVIDIA captures across the full chain is, in operational substance, the monetization of biometric data NVIDIA obtained from Plaintiffs without authorization.

*NVIDIA's Dissemination of the Voice Models Across Its Corporate,
Cloud, and Open-Weight Distribution Infrastructure*

76. The voice synthesis models in which Plaintiffs' and Class members' voiceprints are encoded are not stored in a single location or accessed by a single legal entity. They are deployed, served, processed, transferred, and — in the case of NVIDIA's open-weight model releases — published worldwide, across multiple categories of infrastructure NVIDIA operates in the ordinary course of its voice AI business.

77. NVIDIA conducts its global voice AI business through a corporate structure that includes multiple subsidiaries. NVIDIA's Form 10-K identifies, as significant subsidiaries, NVIDIA International, Inc. (Delaware), NVIDIA Singapore Pte. Ltd. (Singapore), and Mellanox Technologies, Ltd. (Israel), among others. The development, training, refinement, and commercial deployment of NVIDIA's voice models requires, on information and belief, the transfer of those models, and the voiceprints encoded within them, among NVIDIA's affiliated entities in the ordinary course of business. Plaintiffs and Class members did not consent to any such transfer.

78. NVIDIA's voice synthesis pipeline operates within a sub-processor framework involving cloud-infrastructure providers, content-delivery-network operators, hardware partners, and other vendors. The pipeline's training, evaluation, deployment, and operation involve, on information and belief, the transmission of model parameters, and of the voiceprints encoded within those parameters, among NVIDIA's vendors and service providers in the ordinary course of business. The delivery of voice synthesis services to enterprise

customers, including Illinois customers, necessarily involves transmitting model parameters and inference outputs over that infrastructure.

79. NVIDIA distributes several of its voice and speech synthesis models as open-weight model releases on Hugging Face under the NVIDIA AI Open Model License or the NVIDIA Nemotron Open Model License. Through Hugging Face, NVIDIA makes the parameters of those models — including PersonaPlex (released January 2026), Canary-1b-v2, Parakeet-tdt-0.6b-v3, and other speech models, available for download by any party in the world. The biometric characteristics extracted from training-data speakers are encoded in those parameters and persist in every downloaded copy. Once a model's weights have been published on Hugging Face under a permissive license, the encoded biometric data has been disseminated to a downstream ecosystem that NVIDIA cannot retract. NVIDIA cannot recall the downloaded copies. NVIDIA cannot prevent third parties from redistributing, modifying, or fine-tuning them on additional data while preserving the underlying biometric encodings. The open-weight release is a unilateral, irreversible act of dissemination by NVIDIA.

80. Taken together, the pattern of NVIDIA's dissemination practices places the voiceprints and biometric information of Plaintiffs and Class members in the hands of NVIDIA's affiliated entities, vendors, service providers, and any party worldwide who downloads NVIDIA's open-weight voice synthesis models. Plaintiffs and Class members did not consent to any of these disseminations. No enumerated exception to BIPA's dissemination prohibition

under 740 ILCS 14/15(d) applies to the disseminations described in this subsection.

*The Voice Products Compete With the
People Whose Voices Built Them*

81. The voice AI products built on Plaintiffs' biometric data are now sold and deployed into the markets where Plaintiffs and the Class earn their livelihoods.

82. Professional audiobook narration historically has been performed by human narrators at industry rates of approximately \$250 to \$400 per finished hour, meaning a typical ten-hour novel costs \$3,000 to \$4,000 to narrate. NVIDIA's voice cloning models, available to enterprise customers through NVIDIA AI Enterprise and to developers through the NVIDIA API Catalog, can generate synthetic speech in a target voice at a fraction of that cost. NVIDIA's October 2025 marketing describes Magpie TTS Flow as "ideal for studio dubbing and podcast narration," and Magpie TTS Zeroshot can replicate a target speaker's voice from a five-second audio sample. The major American audiobook publishers, including Penguin Random House, Hachette, Simon & Schuster, and Macmillan, for whom Plaintiffs Dorcus and Nassif have narrated books, have begun to integrate AI-generated narration into their production pipelines.

83. NVIDIA's commercial voice products generate podcast-format audio outputs at scale. Magpie TTS Flow, as marketed by NVIDIA itself, is "ideal for...podcast narration." Magpie TTS Multilingual generates expressive

multilingual speech with emotion control. PersonaPlex enables real-time conversational voice generation.

84. NVIDIA has publicly demonstrated the use of its voice synthesis models to replace human narration in its own creative production, including converting a male producer's recorded speech into a female narrator's voice using RAD-TTS for the "I AM AI" video series previously narrated by a human voice actor.

85. NVIDIA's voice synthesis products directly substitute for the long-form investigative audio journalism, podcast hosting, narrative non-fiction production, and on-air reporting that Plaintiffs Marin, Rogers, Lacour, Flowers, and Amer have spent their careers developing — produced, until now, by human reporters and narrators at substantial professional cost.

86. NVIDIA's voice synthesis products are marketed for studio dubbing, e-learning narration, advertising voiceover, video game voice acting, and customer service voice work. Each of these markets historically has been served by professional voice actors, including Plaintiffs Dorcus and Nassif. NVIDIA's voice synthesis products generate substitutable voice content in those markets at per-character and per-audio-second pricing that is orders of magnitude lower than human voice talent rates.

87. NVIDIA's Magpie TTS Multilingual generates speech in multiple languages with emotion control, and the released paired training corpus covers, by NVIDIA's own description, "23 of 25 languages." NVIDIA's voice synthesis capability, on information and belief, across multiple languages and

accents includes generating Arabic-accented English narration, Levantine Arabic narration, and other culturally and linguistically specific voice content. This capability competes specifically with Plaintiff Nassif's distinctive professional market position as a Lebanese-Palestinian American audiobook narrator providing authentic Arabic-accented narration of works by Arab and Palestinian American authors.

88. Each of NVIDIA's voice products was built using the vocal characteristics of the human performers it now displaces. Plaintiffs allege, on information and belief, that the voiceprints of every Plaintiff in this case are among the voiceprints encoded in the foundational voice synthesis models on which these products depend. The market substitution is therefore not merely temporal but causal. The machines generate the audio they generate because the performers' voiceprints were extracted from their recordings and embedded in the model parameters that produce the output.

Named Plaintiffs' Individual Experiences

Philip Rogers

89. Rogers is a broadcast journalist whose four-decade career was conducted in and from Chicago, primarily at WBBM Newsradio (CBS) and WMAQ-TV (NBC 5 Chicago). His on-air work spans radio reporting at WBBM Newsradio, television reporting and anchoring at WMAQ-TV, and live broadcast coverage from conflict zones, disaster scenes, the Olympic Games, mass shootings, corruption trials, and major national and international events. Rogers has been recognized with a National Emmy Award, the Edward R.

Murrow Award, five Associated Press Best Reporter honors, and multiple Peter Lisagor Awards from the Chicago Headline Club. The Lisagor Awards, conferred by Chicago's professional journalism organization, recognize excellence in journalism conducted within the Chicago metropolitan region.

90. Rogers's broadcast catalog comprises thousands of hours of single-speaker, studio-quality audio, continuously available through the NBC 5 Chicago digital archive at nbcchicago.com, where his on-air reports, investigative segments, and broadcast news stories are archived and searchable, and on YouTube, including a career-retrospective interview conducted by the Illinois News Broadcasters Association in which Rogers reflects on four decades of on-air reporting. The platforms display metadata identifying Rogers as the speaker and Chicago as the location of production. Rogers's body of work has been continuously available through these channels for years preceding NVIDIA's training of the foundational voice synthesis models alleged in this Complaint.

91. On information and belief, Rogers's voice recordings were among the audio that NVIDIA ingested to train its foundational voice synthesis models, and voiceprints derived from his recordings are encoded in those models' parameters. Rogers's broadcast catalog matches the profile of training audio NVIDIA's documentation identifies as optimal: long-form, single-speaker, studio-quality, professionally produced, identifiable by name and source. The volume and continuity of his on-air work — four decades of single-speaker

professional broadcasting — represents a category of source material particularly valuable for voice synthesis training.

92. NVIDIA's released Granary speech corpus, assembled through NVIDIA's NeMo Speech Data Processor pipeline, was drawn in substantial part from YouTube-derived corpora (YODAS and YouTube-Commons), and Rogers's broadcast work, including the INBA career-retrospective interview, is publicly accessible on YouTube. NVIDIA has refused to disclose the sources of the training corpus powering its commercial voice products, placing the specific records of training data ingestion under NVIDIA's exclusive control.

93. Rogers never created an account with NVIDIA, downloaded an NVIDIA voice model, invoked an NVIDIA voice AI through any API, or accepted any NVIDIA terms of service related to its voice products. He never received notice that NVIDIA had collected his voiceprint, never received any disclosure of the purpose or duration of that collection, and never executed a written release. NVIDIA's collection of Rogers's voiceprint, and NVIDIA's continuing possession and commercial exploitation of that voiceprint, occurred and continues without his knowledge or consent.

94. Rogers's injury is concrete and particular to him. NVIDIA extracted his voiceprint from recordings he produced over four decades of professional work, encoded that voiceprint in commercial models, and continues to profit from those models. The voiceprint cannot be recovered or replaced; in its operational form, it is the same biological and behavioral signature Rogers used to speak every day during his career. The technology NVIDIA built using

Rogers's voiceprint — including Magpie TTS Multilingual, Magpie TTS Zeroshot, and Magpie TTS Flow, the last of which NVIDIA itself markets as "ideal for studio dubbing and podcast narration" — now operates in the same broadcast journalism market in which Rogers built his career, against the same colleagues with whom he worked, in a competitive position Rogers neither chose nor authorized.

Carol Marin

95. Carol Marin is a five-decade investigative broadcast journalist whose career has been conducted substantially in and from Chicago. Her on-air work has aired on NBC (WMAQ-TV), CBS News (*60 Minutes*, *60 Minutes II*, the *CBS Evening News*), WTTW (*Chicago Tonight*), CNN, and the Discovery Channel, and includes television anchoring, investigative reporting, debate moderation, documentary narration, and long-form interviewing. Marin has been recognized with three George Foster Peabody Awards (including a Personal Peabody), two Alfred I. duPont–Columbia University Awards, two National Emmy Awards, fifteen Regional Emmy Awards, the George Polk Award, the Gracie Award, and the Sigma Delta Chi Ethics in Journalism Award. She has been inducted into the Chicago Journalism Hall of Fame. In 2025, the Governor of Illinois designated Marin a Lincoln Laureate and awarded her the Order of Lincoln, the State of Illinois's highest civilian honor.

96. Marin's broadcast catalog comprises thousands of hours of single-speaker, studio-quality audio. Substantial portions of that catalog are continuously and publicly available through the Media Burn Independent

Video Archive (mediaburn.org), a Chicago-based nonprofit archive that preserves Marin's broadcast investigative reporting and documentary work; through the WTTW digital archive at news.wttw.com; through the NBC 5 Chicago digital archive at nbcchicago.com; and on YouTube, where Marin's Peabody acceptance speeches, debate moderation, and archived broadcast segments are publicly accessible. Each platform displays metadata identifying Marin as the speaker, the producing program, and Chicago as the production location. Marin's body of work has been continuously available through these channels for years preceding NVIDIA's training of the foundational voice synthesis models alleged in this Complaint.

97. On information and belief, Marin's voice recordings were among the audio that NVIDIA ingested to train its foundational voice synthesis models, and voiceprints derived from her recordings are encoded in those models' parameters and reproduced in the audio they generate. Marin's body of work matches, on every material dimension, the profile of training audio that NVIDIA's technical documentation identifies as optimal: long-form, single-speaker, studio-quality, professionally produced, identifiable by name and source. Her recordings are publicly accessible on the same categories of platforms — broadcast archives, YouTube, streaming video — from which NVIDIA, on information and belief, sourced training data, including through its Granary pipeline drawn substantially from YouTube-derived corpora (YODAS and YouTube-Commons). And NVIDIA has refused to disclose the sources of

the training corpus that powers its commercial voice products, placing the specific records of training-data ingestion under NVIDIA's exclusive control.

98. Marin never created an account with NVIDIA, downloaded an NVIDIA voice model, invoked an NVIDIA voice AI through any API, or accepted any NVIDIA terms of service related to its voice products. She never received notice that NVIDIA had collected her voiceprint, never received any disclosure of the purpose or duration of that collection, and never executed a written release. NVIDIA's collection of Marin's voiceprint, and NVIDIA's continuing possession and commercial exploitation of that voiceprint, occurred and continues without her knowledge or consent.

99. Marin's injury is concrete and particular to her. NVIDIA extracted her voiceprint from recordings she produced over five decades of professional work, encoded it into commercial models, and continues to profit from those models. The voiceprint cannot be recovered or replaced; in its operational form, it is the same biological and behavioral signature Marin uses to speak every day. The technology NVIDIA built using Marin's voiceprint — including the Magpie TTS family that NVIDIA itself markets as "ideal for studio dubbing and podcast narration" — now operates in the same broadcast and journalism markets in which Marin built her career, against the same colleagues with whom she worked, in a competitive position Marin neither chose nor authorized.

Alison Flowers

100. Alison Flowers is an investigative journalist and audio producer based in Chicago. Before founding her Chicago-based production company Spiralbound, Flowers served as Head of Production at the Invisible Institute, a Chicago nonprofit investigative journalism organization on the South Side, where she built and led journalism production teams whose work received four nominations from the Pulitzer Prize Board over a four-year period. Flowers produced and reported the seven-part investigative podcast *Somebody*, which premiered on March 31, 2020, investigates the 2016 murder of Courtney Copeland in Chicago, and was a 2021 Pulitzer Prize finalist for Audio Reporting. *Somebody* received the National Magazine Award (the "Ellie") for Podcasting from the American Society of Magazine Editors, the Scripps Howard Award for Excellence in Radio/Podcast Coverage, the International Documentary Association award for Best Audio Documentary, a National Headliner Award, a Gracie Award, and the 2020 Third Coast International Audio Festival Award for Best Serialized Story. *Rolling Stone* named *Somebody* to "The 25 Best True-Crime Podcasts of All Time" and ranked it first on its "Best Podcasts of 2020" list; *The New York Times* ranked it first on its list of true crime podcasts at the intersection of race; *The Atlantic* ranked it second on its "50 Best Podcasts of 2020"; and the podcast reached the number two position among true crime podcasts on the Apple Podcasts chart.

101. *Somebody* is publicly distributed across Apple Podcasts, Spotify, iHeartRadio, YouTube, Stitcher, and other major podcast platforms, with

metadata identifying Flowers as a producer and on-air journalist and Chicago as the location of production. Flowers's additional audio reporting on *Reveal* (from the Center for Investigative Reporting), *The Heist* (from the Center for Public Integrity), *Vox*, *Dateline NBC*, and *Democracy Now!* is likewise continuously available on those programs' respective platforms. Flowers's body of audio work has been continuously available across these channels for years preceding NVIDIA's training of the foundational voice synthesis models alleged in this Complaint.

102. On information and belief, Flowers's voice recordings were among the audio NVIDIA ingested to train its foundational voice synthesis models, and voiceprints derived from her recordings are encoded in the parameters of those models. Flowers's audio reporting matches the profile of training audio NVIDIA's documentation identifies as optimal — long-form, single-speaker, studio-quality, professionally produced. Her Pulitzer-finalist work was a readily identifiable high-value source of professional audio journalism, recognized commercially as among the best podcasts of its release year by *Rolling Stone*, *The New York Times*, and *The Atlantic*. Her recordings are distributed on the same podcast platforms — Apple Podcasts, Spotify, YouTube, and others — from which NVIDIA, on information and belief, sourced training data, including through its Granary pipeline drawn substantially from YouTube-derived corpora. NVIDIA has refused to disclose the sources of the training corpus that powers its commercial voice products.

103. Flowers never created an account with NVIDIA, downloaded an NVIDIA voice model, invoked an NVIDIA voice AI through any API, or accepted any NVIDIA terms of service related to its voice products. She never received notice that NVIDIA had collected her voiceprint, never received any disclosure of the purpose or duration of that collection, and never executed a written release.

104. Flowers's injury is concrete and particular to her. Flowers continues to produce investigative audio journalism through Spiralbound, and the long-form audio markets in which she earns her livelihood — investigative podcast journalism, narrative audio documentary — are precisely the markets that NVIDIA's Magpie TTS Flow, marketed by NVIDIA as "ideal for studio dubbing and podcast narration," is designed to serve. The technology NVIDIA built using Flowers's voiceprint now competes with Flowers in her own active markets, with the capabilities of her own voice contributing to the products that compete against her.

Robin Amer

105. Robin Amer is a journalist, podcast creator, audio producer, and on-air host who has produced the substantial majority of her audio work from Chicago. Amer is the creator, host, narrator, and showrunner of *The City*, the investigative podcast produced by USA Today over two seasons. Season 1 of *The City*, which focused on Chicago, peaked at No. 6 on the Apple Podcasts charts and was named Best Podcast of the Year by *The New Yorker*, *The New York Times*, *Quartz*, and Apple Podcasts. Amer's earlier production work for

Gravy, the podcast produced by the Southern Foodways Alliance, contributed to that program's 2015 James Beard Award for Best Podcast. Amer subsequently spent three years as Senior Producer for Audio Features at *The Washington Post*, where she edited the *Post Reports* daily news podcast and produced the standalone narrative series *Field Trip*. During her tenure at *The Washington Post*, Amer won or was a finalist for the Alfred I. duPont–Columbia University Award for three consecutive years. Amer currently serves as Managing Editor of *Love + Radio*, where she managed production of *Blood Memory*, a ten-part narrative series that won the 2025 Tribeca Festival Audio Storytelling prize for Best Independent Non-Fiction and was shortlisted for the Whickers Prize at the Sheffield Documentary Festival.

106. *The City* is publicly distributed across Apple Podcasts, Spotify, iHeartRadio, and other major podcast platforms, with metadata identifying Amer as the host, narrator, and showrunner. *Post Reports*, *Field Trip*, *Love + Radio*, *Blood Memory*, and Amer's contributions to *Reveal*, *The Heist*, *Vox*, and *Gravy* are likewise continuously available on the respective programs' platforms. Amer's audio work has been continuously available across these channels for years preceding NVIDIA's training of the foundational voice synthesis models alleged in this Complaint. The substantial majority of Amer's voice work was recorded in Chicago while Amer was an Illinois resident, and Amer has been an Illinois resident at all times relevant to this Complaint.

107. On information and belief, Amer's voice recordings were among the audio NVIDIA ingested to train its foundational voice synthesis models, and

voiceprints derived from her recordings are encoded in the parameters of those models. Amer's hosted, narrated, and showrun audio work — totaling hundreds of hours of single-speaker, studio-quality audio — matches the profile of training audio NVIDIA's documentation identifies as optimal. *The City's* commercial peak position on the Apple Podcasts charts and Amer's three-time duPont–Columbia recognition established her work as a readily identifiable high-value source of professional audio journalism. Her recordings are distributed on the same podcast platforms — Apple Podcasts, Spotify, YouTube, and others — from which NVIDIA, on information and belief, sourced training data, including through its Granary pipeline drawn substantially from YouTube-derived corpora. NVIDIA has refused to disclose the sources of the training corpus that powers its commercial voice products.

108. Amer never created an account with NVIDIA, downloaded an NVIDIA voice model, invoked an NVIDIA voice AI through any API, or accepted any NVIDIA terms of service related to its voice products. She never received notice that NVIDIA had collected her voiceprint, never received any disclosure of the purpose or duration of that collection, and never executed a written release.

109. Amer's injury is concrete and particular to her. Amer is presently a working audio producer, host, and managing editor in long-form investigative podcasting — directly the market that NVIDIA's Magpie TTS Flow, marketed by NVIDIA as "ideal for studio dubbing and podcast narration," is designed to serve at substantially lower cost than the human production Amer provides.

The technology NVIDIA built using Amer's voiceprint now competes with Amer in her own active market.

Lindsey Dorcus

110. Lindsey Dorcus is a professional audiobook narrator who has recorded more than 200 audiobooks for major American publishers, including Penguin Random House, Simon & Schuster, Macmillan, Hachette, Disney Hyperion, Audible Studios, Blackstone Publishing, Tantor Media, Harper Audio, Podium, and Scribd. She is a 2020 Society of Voice Arts and Sciences Voice Arts Award winner (as part of the full-cast ensemble for *Wild Monsters Dance About: Stories from an Unruly Mind*) and a 2021 Independent Audiobook Award winner for LGBTQ+ audiobook narration. Her narration spans young adult fiction, science fiction and fantasy, romance, mystery and thriller, and works featuring LGBTQ+ themes and characters. *AudioFile Magazine*, the principal independent trade publication reviewing audiobook narration in the United States, has reviewed her work favorably, describing her performances as "silky," "joyful," and capable of "drawing listeners in with the haunting cadence of her voice." Dorcus possesses a professionally recognized range of accents and dialects used across her narration work, including General American, British (Received Pronunciation, Estuary, and Cockney), Scottish, Irish (Dublin and Northern Irish), French, American Southern, and Greek for main characters and narration, and New England, New York, German, Indian, and Russian for supporting characters. Dorcus operates a professional home recording studio in Chicago, where she records the audiobooks at issue in this

Complaint to broadcast-quality, professionally edited specifications that meet the technical requirements of the major audiobook publishers.

111. Dorcus's audiobook catalog is continuously distributed across Audible, Apple Books, Google Play Books, Spotify, Libro.fm, Chirp, Scribd/Everand, and other major audiobook platforms. Each platform displays metadata identifying Dorcus as the narrator, the publisher, and the title. The complete catalog is searchable on Audible.com. Dorcus has also appeared as a voice actor in multiple fiction podcasts performing dramatic voice work for serialized audio storytelling distributed through major podcast platforms. Dorcus's narrated catalog has been continuously available through these channels for years preceding NVIDIA's training of the foundational voice synthesis models alleged in this Complaint.

112. On information and belief, Dorcus's voice recordings were among the audio NVIDIA ingested to train its foundational voice synthesis models, and voiceprints derived from her recordings are encoded in the parameters of those models. Dorcus's catalog represents precisely the long-form, single-speaker, studio-quality audio that NVIDIA's own documentation identifies as optimal training material — more than two hundred audiobooks, each comprising hours of consistent single-speaker narration, professionally produced, publicly searchable by narrator, distributed on major commercial audio platforms. Few professional voice catalogs match NVIDIA's training-data profile more directly than Dorcus's does. The technical sophistication of her voice work — including her professionally recognized range across more than a dozen accents and

dialects — represents an additional category of training-data value to a multilingual voice synthesis system designed to generate speech across dozens of languages and accents. NVIDIA's published Granary speech corpus — assembled through NVIDIA's NeMo Speech Data Processor pipeline at approximately one million hours of audio, drawn in substantial part from the LibriLight corpus of LibriVox audiobooks and used to train NVIDIA's Canary-1b-v2 and Parakeet-tdt-0.6b-v3 models — demonstrates that NVIDIA actively ingests audiobook-format speech audio into its voice training pipelines. NVIDIA has refused to disclose whether commercial-publisher audiobook content was likewise ingested into the Magpie TTS family or other commercial voice synthesis models, placing the specific records of training-data ingestion under NVIDIA's exclusive control.

113. Dorcus never created an account with NVIDIA, downloaded an NVIDIA voice model, invoked an NVIDIA voice AI through any API, or accepted any NVIDIA terms of service related to its voice products. She never received notice that NVIDIA had collected her voiceprint, never received any disclosure of the purpose or duration of that collection, and never executed a written release.

114. Dorcus's injury is concrete, particular to her, and direct. NVIDIA's Magpie TTS Flow — which NVIDIA itself markets as "ideal for studio dubbing and podcast narration" and which can replicate a target speaker's voice from a five-second audio sample at "speaker similarity of over 70%" — is licensed to enterprise customers, including audiobook publishers, at per-character and

per-audio-second fees that are orders of magnitude lower than human narrator rates. The market displacement is not abstract. The technology NVIDIA built using Dorcus's voiceprint, on information and belief, is now sold to the same publishers that have historically paid Dorcus to narrate audiobooks, as a substitute for the human narration she provides. Dorcus's voiceprint, taken without her consent, is now part of the commercial product that competes against her in her own market.

Yohance Lacour

115. Yohance Lacour is a journalist, audio storyteller, writer, and playwright from the South Side of Chicago whose work centers on the lives and stories of Black Chicago. He is affiliated with the Invisible Institute, a nonprofit investigative journalism organization based on Chicago's South Side. Lacour is the creator, host, writer, and lead reporter of *You Didn't See Nothin'*, a seven-part investigative podcast produced by the Invisible Institute and USG Audio that revisits the 1997 hate-crime attack on Lenard Clark on the South Side of Chicago, tracks down key players in the case a quarter-century later, and examines how the case shaped Lacour's own life. The series was awarded the 2024 Pulitzer Prize for Audio Reporting and a 2024 Peabody Award. It was named to Apple Podcasts' "Podcasts We Love" and recognized as one of Apple Podcasts' "100 Best Podcasts of All Time," received nominations at the Signal Podcasting Awards in the Limited Series & Specials – Best Host and Limited Series & Specials – Documentary categories, and received four nominations at the Black Podcasting Awards in the Best Sound Design, Best History Podcast,

Best True Crime Podcast, and Best Limited Series Podcast categories. Lacour serves as the on-air voice, narrator, and lead reporter throughout the series.

116. *You Didn't See Nothin'* is publicly distributed across Apple Podcasts, Spotify, iHeartRadio, Amazon Music, YouTube, Overcast, and other major podcast platforms, with metadata identifying Lacour as the host and Chicago as the location of production. Lacour's interview appearances on NPR's *Fresh Air* with Tonya Mosley, NPR's *All Things Considered* with Adrian Florido, the Canadian Broadcasting Corporation's *Crime Story* with Kathleen Goldhar, and the Pulitzer Prize Board's *Pulitzer on the Road* podcast (produced by Audacy's Pineapple Street Studios) are likewise continuously available across major public-radio and podcast platforms. The series has been continuously available since its 2023 release, and Lacour's interview appearances have been continuously available since their respective publications, in each case for periods that precede NVIDIA's training of the foundational voice synthesis models at issue in this Complaint and continue thereafter.

117. On information and belief, Lacour's voice recordings were among the audio that NVIDIA ingested to train its foundational voice synthesis models, and voiceprints derived from his recordings are encoded in those models' parameters. Lacour's catalog matches the profile of training audio NVIDIA's documentation identifies as optimal — long-form, single-speaker, studio-quality, professionally produced — and his Pulitzer-honored audio work was, by virtue of its public recognition, a readily identifiable high-value source of training audio. His recordings are distributed on the same podcast platforms —

Apple Podcasts, Spotify, YouTube, and others — from which NVIDIA, on information and belief, sourced training data, including through its Granary pipeline drawn substantially from YouTube-derived corpora. NVIDIA has refused to disclose the sources of the training corpus powering its commercial voice products, placing the specific records of training data ingestion under NVIDIA's exclusive control.

118. Lacour never created an account with NVIDIA, downloaded an NVIDIA voice model, invoked an NVIDIA voice AI through any API, or accepted any NVIDIA terms of service related to its voice products. He never received notice that NVIDIA had collected his voiceprint, never received any disclosure of the purpose or duration of that collection, and never executed a written release.

119. Lacour's injury is concrete and particular to him. NVIDIA's Magpie TTS Flow, the voice cloning model NVIDIA itself markets as "ideal for studio dubbing and podcast narration," generates podcast-format audio from a target speaker's voice with "speaker similarity of over 70%" from a five-second audio sample. That product is directly in the format and the market of *You Didn't See Nothin'* and Lacour's other audio journalism. The technology NVIDIA built using Lacour's voiceprint now operates in the same long-form audio investigative-journalism market in which Lacour earns his livelihood, with the capabilities of his own voice contributing to the products that compete against him.

Victoria Nassif

120. Victoria Nassif is a first-generation Lebanese-Palestinian American actor, audiobook narrator, voiceover artist, and intimacy director based in Illinois. Nassif is a trained mezzo-soprano singer whose vocal performance skills extend beyond spoken narration to include musical and singing capabilities developed through classical acting training. Nassif's audiobook narration work has been commercially released by Penguin Random House (Random House Audio), Hachette Book Group (Little, Brown Young Readers), and Simon & Schuster, and her on-camera work includes multiple episodes of NBC's *Chicago PD* (Season 12) and nationally broadcast commercials. Her notable audiobook narrations include *The Next New Syrian Girl* by Ream Shukairy (Hachette/Little, Brown Young Readers), in which she serves as the solo narrator performing multiple characters with authentic Levantine Arabic accents; *The Skin and Its Girl* by Sarah Cypher (Random House Audio), a novel featuring a queer Palestinian American protagonist that was shortlisted for the Ursula K. Le Guin Prize and named a *them* magazine Best Book of the Year; *Gulf* by Mo Ogrodnik (Simon & Schuster); *The Jasad Crown* by Sara Hashem; and *Every Moment is a Life*, a bilingual Arabic-English anthology compiled by bestselling author Susan Abulhawa featuring stories from Palestinian writers. Nassif possesses a professionally recognized range of accents and dialects, including General American, British (Received Pronunciation and Cockney), Persian, Levantine Arabic, and American Southern.

121. Nassif's audiobook catalog is continuously distributed across Audible, Apple Books, Spotify, Libro.fm, and other major audiobook platforms, with metadata identifying Nassif as the narrator, the publisher, and the title. Her on-camera work is continuously available through NBC and other broadcast distribution channels. Her audiobook catalog has been continuously available through these platforms for years preceding NVIDIA's training of the foundational voice synthesis models alleged in this Complaint.

122. On information and belief, Nassif's voice recordings were among the audio NVIDIA ingested to train its foundational voice synthesis models, and voiceprints derived from her recordings are encoded in the parameters of those models. Nassif's audiobook catalog matches the profile of training audio NVIDIA's documentation identifies as optimal, and her authentic Levantine Arabic-accented narrations are a distinctive and readily identifiable category of voice content particularly valuable for a multilingual voice synthesis system. NVIDIA's released Magpie TTS Multilingual model, by NVIDIA's own description, generates speech with "preference alignment" training across "23 of 25 languages," and NVIDIA's published Granary speech corpus — used to train NVIDIA's Canary-1b-v2 and Parakeet-tdt-0.6b-v3 multilingual speech models — was assembled in substantial part from open-source corpora that include LibriVox audiobooks. As a first-generation Lebanese-Palestinian American with native cultural and linguistic familiarity with Arabic-language and Levantine speech patterns, Nassif brings to her narration of works featuring Middle Eastern characters and settings an authenticity that

represents a category of training-data value NVIDIA's foundational voice models depend on capturing for their cross-language and cross-accent generative capability. NVIDIA has refused to disclose whether commercial-publisher audiobook content, multilingual voice content, or culturally-specific narration was ingested into the foundational models that power its commercial voice products, placing the specific records of training-data ingestion under NVIDIA's exclusive control.

123. Nassif never created an account with NVIDIA, downloaded an NVIDIA voice model, invoked an NVIDIA voice AI through any API, or accepted any NVIDIA terms of service related to its voice products. She never received notice that NVIDIA had collected her voiceprint, never received any disclosure of the purpose or duration of that collection, and never executed a written release.

124. Nassif's injury is concrete and particular to her. The cultural and linguistic specificity of Nassif's professional voice work — authentic Arabic-accented narration of works by Arab and Palestinian American authors — is itself the basis of her distinctive market position. NVIDIA's Magpie TTS Multilingual and related multilingual voice synthesis products, on information and belief, can now generate Arabic-accented English narration and Levantine Arabic-accented voice content that competes specifically with Nassif's professional niche, and that competitive capability exists because voice characteristics encoded in NVIDIA's foundational models include the characteristics of speakers like Nassif whose recordings supplied the

multilingual and cross-accent training data on which the models were built. The technology NVIDIA built using Nassif's voiceprint now operates in the precise market, culturally authentic Arabic-accented narration, in which Nassif has built her career, in a competitive position Nassif neither chose nor authorized.

NVIDIA Acted Willfully and Recklessly

125. NVIDIA's collection, retention, commercial exploitation, and dissemination of Plaintiffs' voiceprints without notice or consent was not the result of inadvertence or unfamiliarity with BIPA. NVIDIA acted with knowledge of, or at a minimum, reckless disregard for, its obligations under Illinois law.

NVIDIA knew its obligations under BIPA

126. By the time NVIDIA released its most capable commercial voice cloning products — Magpie TTS ZeroShot and Magpie TTS Flow in October 2025, PersonaPlex in January 2026, and Nemotron 3 VoiceChat in 2026 — BIPA had been the law of Illinois for more than fifteen years and had generated some of the largest privacy settlements in American history. Those settlements include *In re Facebook Biometric Information Privacy Litigation*, No. 3:15-cv-03747 (N.D. Cal.) (approximately \$650 million); *Rivera v. Google LLC*, Cook County Cir. Ct. No. 2019-CH-00990 (approximately \$100 million); and *In re TikTok, Inc., Consumer Privacy Litigation*, No. 1:20-cv-04699 (N.D. Ill.) (approximately \$92 million). In late 2025, Google paid \$1.375 billion to the State of Texas to resolve parallel claims under Texas's biometric identifier statute, following Texas's \$1.4 billion settlement with Meta in 2024 over

biometric data collection — two of the largest privacy settlements any state has ever obtained. By any measure, NVIDIA was on notice of the biometric-privacy obligations that govern the extraction of biometric identifiers from human source material throughout the period in which it ingested Plaintiffs' voiceprints into its foundational voice synthesis models.

127. Voice-AI-specific litigation further sharpened the notice. In May 2024, voice actors filed *Lehrman v. Lovo, Inc.*, No. 1:24-cv-03770 (S.D.N.Y. filed May 16, 2024), alleging that an AI voice company created and commercialized unauthorized voice clones trained on the plaintiffs' recordings. In August 2024, additional voice actors filed *Vacker v. Eleven Labs, Inc.*, No. 1:24-cv-00987 (D. Del. filed Aug. 29, 2024), asserting analogous claims against another AI voice company. By the time NVIDIA released Magpie TTS Zeroshot, Magpie TTS Flow, and PersonaPlex, AI voice companies were active defendants in litigation alleging the same conduct: building commercial voice synthesis products on voice recordings ingested without the speakers' consent.

128. NVIDIA's corporate disclosures recognized that biometric data collection is a relevant concern for its AI models. NVIDIA's Trustworthy AI terms, incorporated by reference into the licenses under which NVIDIA distributes several of its open-weight voice models, including the NVIDIA AI Open Model License and the NVIDIA Nemotron Open Model License, expressly prohibit use of NVIDIA's models for "unlawful biometric data collection." At the same time, NVIDIA's general Privacy Policy categorically states that NVIDIA "does not use, sell, or share 'sensitive personal information' as defined by

California law", a category that expressly includes biometric information. A company that recognizes the biometric implications of its technology in one public document while categorically denying them in another, and proceeds without implementing the notice, consent, retention, and destruction measures Illinois law requires, has, at a minimum, acted with reckless disregard.

NVIDIA knew that voice training extracts biometric identifiers

129. NVIDIA's published research and product documentation describe the voice synthesis pipeline in terms that align with BIPA's biometric definitions. NVIDIA's Magpie TTS Flow model card describes the model as analyzing "a speaker's voice" to "replicate voice qualities such as pitch, timbre and speech rate" and as "[m]aintain[ing] the original characteristics that capture unique voice audio signature."

130. NVIDIA's PersonaPlex model card describes the system as conditioned on "a voice prompt" consisting of "audio tokens that establish the target vocal characteristics and speaking style."

131. NVIDIA's July 2025 Riva TTS blog post describes Magpie TTS training as using "speaker verification (SV) models within a reward system to create a preference dataset" — that is, computational comparison of speaker acoustic signatures during training.

132. NVIDIA has not pleaded ignorance, and could not credibly plead ignorance, of the proposition that its voice synthesis pipeline extracts and stores speaker-specific representations capable of identifying the source

speakers. NVIDIA built the pipeline, published the research, and shipped the product on that very capability.

133. NVIDIA has confirmed its understanding that voice synthesis training implicates the legal rights of source speakers in two distinct ways. First, by building safeguards against the downstream misuse of its own voice synthesis products. Second, by placing on its enterprise customers the upstream-rights obligation that NVIDIA failed to discharge with respect to its own training data.

134. In November 2024, NVIDIA announced FUGATTO, a 2.5-billion-parameter foundational audio generative model trained on more than 50,000 hours of audio and capable of generating and transforming human voices, including by modifying accent and emotion. NVIDIA did not release FUGATTO publicly. NVIDIA cited concerns about misuse, including voice impersonation and deepfake generation.

135. In October 2025, when NVIDIA released Magpie TTS Zeroshot and Magpie TTS Flow — voice cloning models that can replicate a target speaker's voice from a five-second audio sample — NVIDIA simultaneously announced a collaboration with Pindrop, a voice-fraud detection company, to develop deepfake detection capabilities against NVIDIA's voice models.

136. In January 2026, when NVIDIA released PersonaPlex, NVIDIA architecturally restricted that model to a fixed set of eighteen pre-packaged voice embeddings rather than supporting arbitrary voice cloning from user-supplied audio. Each of these measures was implemented at the downstream

end of NVIDIA's commercial pipeline, to constrain how end users use NVIDIA's products. None addressed the upstream issue: that NVIDIA's foundational voice models were trained on voiceprints extracted without the knowledge or consent of the speakers whose voice recordings were ingested at training time. NVIDIA built each of these safeguards against the downstream misuse of its products. It built none of them for the people whose voices it took to train those products in the first place.

137. NVIDIA's Product Specific Terms for AI Products define a "Custom TTS Application" as "a custom text-to-speech application that enables the creation of synthetic voice output that resembles an input voice," and require NVIDIA's enterprise customers to "have sufficient rights and licenses for content used with a Custom TTS Application to generate new content."

138. NVIDIA's October 2025 Magpie TTS launch documentation advised users to use voice cloning only with "target speakers who have consented to such use."

139. Each of these contractual and advisory provisions is NVIDIA's acknowledgment that the commercial use of an individual's voice to generate synthetic speech requires that individual's consent. NVIDIA imposed that obligation on its downstream enterprise customers. NVIDIA did not discharge the same obligation, in symmetric form, with respect to its own ingestion of training data, even though the foundational voice models on which the customer applications depend were built by ingesting voice recordings at training time without the speakers' notice, consent, or written release.

NVIDIA chose to apply consent selectively

140. The pattern of NVIDIA's conduct between 2024 and 2026 reflects a deliberate, contemporaneous choice. NVIDIA built and publicly deployed customer-facing consent and rights-acknowledgment infrastructure in connection with its voice cloning products. It built none of the corresponding upstream consent infrastructure for the speakers whose voices populated the training corpora that those same products depend on.

141. NVIDIA's Magpie TTS Flow training corpus, by NVIDIA's own disclosure, comprises approximately 70,000 hours of paired speech data; its FUGATTO training corpus comprises more than 50,000 hours of audio; and the Granary speech corpus assembled through NVIDIA's NeMo Speech Data Processor pipeline comprises approximately one million hours of audio.

142. The same engineering organization, ADLR, developed both the customer-facing voice cloning products with their downstream consent and rights-acknowledgment infrastructure and the foundational voice models trained on the unconsented training corpus. The same legal department drafted the customer-facing terms and the model licenses. The decision to apply consent and rights-acknowledgment to the downstream uses NVIDIA monetizes through its commercial customers, and not to the upstream voices NVIDIA took to build those products in the first place, was a contemporaneous, institutionally deliberate choice made with the consent infrastructure already in place.

143. The decision was a commercial calculation. BIPA-compliant collection of voice training data would have required NVIDIA to identify the source speakers, provide written notice of the specific purpose and duration of the collection, and obtain a written release from each speaker before ingesting the recording into the training pipeline. With foundational training corpora measured in the tens of thousands and hundreds of thousands of hours, comprising, on information and belief, the recordings of hundreds of thousands of distinct speakers and likely more, that compliance burden would have constrained the speed and scale of NVIDIA's voice AI development. NVIDIA chose speed and scale over compliance. The decision is documented in the contrast between NVIDIA's robust consent and rights-acknowledgment infrastructure for the downstream uses it monetizes and the complete absence of corresponding consent infrastructure for the upstream voices it took. That was not an oversight. It was a business decision.

NVIDIA's training-data acquisition practices confirm deliberateness

144. NVIDIA's training-data acquisition practices, taken on the public record, confirm that NVIDIA's noncompliance was a deliberate institutional choice rather than inadvertent oversight. As alleged at ¶ 57, NVIDIA is presently a defendant in two pending putative class actions, *Nazemian v. NVIDIA Corp.* and *Ted Entertainment, Inc. v. NVIDIA Corp.*, alleging non-consensual ingestion of pirated e-books and scraped YouTube videos, respectively, into the training corpora of NVIDIA's foundational language and video models. Together with the internal NVIDIA communications described at

¶ 56, documenting a company-wide initiative that processed approximately "a human lifetime" of video content per day, those actions establish a pattern of non-consensual ingestion of identifiable source material at a massive scale across NVIDIA's foundational model families. These actions concern language and video models, not voice models. But the same pattern, applied to a different category of source material, is what this Complaint alleges regarding NVIDIA's foundational voice synthesis models.

145. NVIDIA's August 2025 release of the Granary dataset confirms that NVIDIA operates the infrastructure required to ingest voice audio at the scale alleged here. NVIDIA publicly described Granary as an approximately one-million-hour corpus of speech audio assembled by processing pre-existing open-source corpora — including YODAS and YouTube-Commons (both YouTube-derived), VoxPopuli (European Parliament proceedings), and LibriLight (LibriVox audiobooks) — through NVIDIA's NeMo Speech Data Processor pipeline.

146. NVIDIA used Granary to train its Canary-1b-v2 and Parakeet-tdt-0.6b-v3 speech models, which NVIDIA distributes as open-weight models on Hugging Face.

147. The Granary pipeline demonstrates, in NVIDIA's own public documentation, that NVIDIA maintains infrastructure for the large-scale ingestion and processing of human voice audio from publicly accessible platforms; that NVIDIA willingly ingests identifiable voice recordings from those platforms without obtaining the consent of the speakers whose voices appear in

the recordings; and that NVIDIA's voice-audio processing pipeline operates at a scale, approximately one million hours, well in excess of what is necessary to produce the commercial voice synthesis and voice cloning models at issue in this case.

148. Each of the violations alleged in this Complaint was therefore committed by NVIDIA with knowledge of, or in reckless disregard for, BIPA's requirements. Plaintiffs are entitled to liquidated damages of \$5,000 per violation under 740 ILCS 14/20(2), or, in the alternative, \$1,000 per violation under 740 ILCS 14/20(1), recoverable on a per-person, per-subsection basis where multiple distinct provisions of § 15 are violated, consistent with the statute as amended and as construed in *Clay v. Union Pacific Railroad Co.*, No. 25-2185 (7th Cir. Apr. 1, 2026).

CLASS ACTION ALLEGATIONS

149. Plaintiffs bring this action individually and on behalf of all others similarly situated, pursuant to Federal Rules of Civil Procedure 23(b)(2) and 23(b)(3), as the following Class: All natural persons whose voice recordings were produced or recorded in Illinois, and from whose recordings NVIDIA extracted, derived, or otherwise obtained voiceprints or biometric information in connection with the development, training, fine-tuning, evaluation, or operation of NVIDIA's foundational voice synthesis models, or commercial voice products derived from those models, during the Class Period.

150. The Class Period runs from the earlier of (a) the date NVIDIA first ingested any voice recording into the training pipeline for any of its

foundational voice synthesis models, or (b) January 1, 2018, through the date of judgment in this action. Discovery will establish the operative start date of the Class Period.

151. Excluded from the Class are: (i) NVIDIA Corporation and each of its parents, subsidiaries, affiliates, and controlled entities, including without limitation NVIDIA International, Inc., NVIDIA Singapore Pte. Ltd., and Mellanox Technologies, Ltd.; (ii) all current and former officers and directors of NVIDIA; (iii) NVIDIA's employees, contractors, agents, and counsel; (iv) the Court, the Court's staff, and any jurors assigned to this action; (v) the immediate family members of any person excluded above; and (vi) any person who executed a written release authorizing NVIDIA's use of their voice recordings to train, fine-tune, or operate NVIDIA's foundational voice synthesis models, in compliance with 740 ILCS 14/15(b).

152. Plaintiffs reserve the right to amend or refine the Class definition based on facts learned through discovery. Nothing in the Class definition limits or disclaims claims or remedies available under any statute or theory asserted in this Complaint.

153. *Ascertainability*. Class membership is defined by objective criteria and can be determined from records that exist or will be produced in discovery. Whether a particular voice recording entered NVIDIA's training pipeline is a binary factual question — the recording is either in the pipeline's ingestion logs or it is not — and the records that answer the question are within NVIDIA's exclusive control. NVIDIA itself has publicly described, in connection with its

August 2025 release of the Granary dataset, the operation of a NeMo Speech Data Processor pipeline that ingests, processes, and tracks the provenance of approximately one million hours of speech audio. The same category of pipeline metadata — training-data manifests, ingestion logs, source-URL records, dataset-version records, and associated speaker- and file-level identifiers — exists for the voice training data on which NVIDIA built Magpie TTS Multilingual, Magpie TTS Zeroshot, Magpie TTS Flow, FUGATTO, PersonaPlex, Nemotron 3 VoiceChat, Canary, and Parakeet. Class membership can be further confirmed through publicly available distribution metadata for voice recordings (audiobook platform records, podcast directory records, streaming service catalogs, broadcast archive records), through speaker-identification technology applied to NVIDIA's training corpus, and through voice-matching analysis comparing NVIDIA's voice model outputs against publicly available recordings of class members.

154. *Numerosity*. Joinder of all members of the Class is impracticable. Fed. R. Civ. P. 23(a)(1). NVIDIA has disclosed that Magpie TTS Flow was trained on approximately 70,000 hours of paired speech data, that FUGATTO was trained on more than 50,000 hours of audio, and that the Granary speech corpus assembled through NVIDIA's NeMo Speech Data Processor pipeline comprises approximately one million hours of audio. The number of distinct individuals whose voice recordings were produced or recorded in Illinois and ingested into the training pipelines that produced NVIDIA's commercial voice products plainly satisfies the threshold for numerosity. The Class includes not

only the professional broadcast journalists, audiobook narrators, podcasters, voice actors, and other voice professionals who have produced work in Illinois during the Class Period — themselves a population numbering in the thousands — but also the interviewees, guests, panelists, witnesses, callers, public officials, and other persons whose voices were captured in Illinois-produced or Illinois-recorded broadcast, podcast, audiobook, archival, and other publicly distributed audio content, and whose recordings are accessible on the publicly distributed audio platforms from which NVIDIA, on information and belief, sourced training data. On information and belief, the Class population numbers in the hundreds of thousands or more. The exact number is within NVIDIA's exclusive control and will be established through discovery.

155. *Commonality.* Common questions of law and fact apply to every member of the Class. Fed. R. Civ. P. 23(a)(2). NVIDIA did not engage in any individualized notice, consent, retention-policy disclosure, release, or biometric-data-protection process with respect to any non-user whose voice recordings were ingested into NVIDIA's foundational voice training pipelines. NVIDIA's conduct was uniform: the same training pipelines ingested the same categories of voice recordings under the same absent-consent posture, applied to all class members through the same automated and standardized process. The questions whether NVIDIA complied with BIPA's notice, consent, retention-policy, profiting, dissemination, and privacy-protection requirements before and after extracting the voiceprints of class members can be answered

classwide because NVIDIA's compliance, or noncompliance, was identical as to every class member.

156. Common questions of law and fact include, without limitation:

(a) whether the computational representations of vocal characteristics that NVIDIA extracts during voice model training, including speaker embeddings, audio tokens, voice prompts, and the speaker-acoustic-signature representations described in NVIDIA's model documentation, constitute "voiceprints" or "biometric information" within the meaning of BIPA;

(b) whether NVIDIA informed class members in writing that their biometric identifiers were being collected, of the specific purpose and duration of collection, and obtained a written release, as 740 ILCS 14/15(b) requires;

(c) whether NVIDIA developed and made publicly available a written retention and destruction policy applicable to class members' biometric identifiers, as 740 ILCS 14/15(a) requires;

(d) whether NVIDIA sold, leased, traded, or otherwise profited from class members' biometric identifiers in violation of 740 ILCS 14/15(c);

(e) whether NVIDIA disclosed or disseminated class members' biometric identifiers to third parties — including through the open-weight release of voice synthesis models on Hugging Face under the NVIDIA AI Open Model License and the NVIDIA Nemotron Open Model License — without consent and outside any enumerated exception, in violation of 740 ILCS 14/15(d);

(f) whether NVIDIA stored, transmitted, and protected class members' biometric identifiers using the reasonable standard of care required by 740 ILCS 14/15(e)(1) and in a manner equally or more protective than its protection of other confidential and sensitive information as required by 740 ILCS 14/15(e)(2), including in light of NVIDIA's open-weight release of voice models that encode the biometric characteristics extracted from class members' recordings;

(g) whether NVIDIA's conduct was willful or reckless within the meaning of 740 ILCS 14/20(2);

(h) whether NVIDIA used class members' voices and identities for commercial purposes without prior written consent in violation of the Illinois Right of Publicity Act;

(i) whether NVIDIA's voice products generate, distribute, or make available unauthorized digital replicas within the meaning of 765 ILCS 1075/30(c), and NVIDIA materially contributes to or facilitates their distribution;

(j) whether NVIDIA's conduct constitutes deceptive trade practices likely to cause confusion in violation of the Illinois Uniform Deceptive Trade Practices Act, 815 ILCS 510/2(a)(2) and 510/2(a)(3);

(k) whether NVIDIA was unjustly enriched by its unauthorized use of class members' voice data; and

(l) the appropriate measures of damages, restitution, and injunctive relief.

157. *Typicality*. The claims of the named Plaintiffs are typical of the claims of the Class. Fed. R. Civ. P. 23(a)(3). Each named Plaintiff produced voice recordings in Illinois — broadcast journalism, audio reporting, audiobook narration, podcast production, and related professional voice work, recorded in Chicago studios, broadcast facilities, and home recording studios in this District. NVIDIA, on information and belief, ingested those recordings into the training pipelines for its foundational voice synthesis models and extracted voiceprints from them without notice, consent, or written release. The voiceprints are now embedded in NVIDIA's commercial voice models — including Magpie TTS Multilingual, Magpie TTS Zeroshot, Magpie TTS Flow, FUGATTO, PersonaPlex, and Nemotron 3 VoiceChat — and reproduced in the audio those models generate. The legal theories asserted on behalf of the Class — that NVIDIA's extraction, retention, commercial exploitation, and dissemination of voiceprints from Illinois-recorded voice work without BIPA-compliant consent violates 740 ILCS 14/15(a)–(e), and that NVIDIA's commercial use of voices and identities without consent violates the Illinois Right of Publicity Act — apply with equal force to each named Plaintiff and to every other class member. Each named Plaintiff has the same interest as every other class member in establishing NVIDIA's liability and obtaining the relief sought in this Complaint.

158. *Adequacy*. Plaintiffs will fairly and adequately protect the interests of the Class. Fed. R. Civ. P. 23(a)(4). Plaintiffs' interests are aligned with, and not antagonistic to, the interests of the absent class members; each named

Plaintiff has the same incentive as every other class member to maximize recovery and to obtain comprehensive injunctive relief addressing the unlawful extraction of voiceprints. Plaintiffs are represented by counsel experienced in complex class action litigation, privacy litigation, and BIPA litigation, with the resources to prosecute this action vigorously on behalf of the Class.

159. *Rule 23(b)(2) Certification.* Certification under Rule 23(b)(2) is appropriate because NVIDIA has acted on grounds generally applicable to the Class, such that final injunctive and corresponding declaratory relief is appropriate as to the Class as a whole. NVIDIA's training pipelines operated uniformly across every class member's voice recordings; NVIDIA's failure to obtain BIPA-compliant consent was uniform; NVIDIA's failure to publish a retention and destruction policy applicable to non-user training data subjects was uniform; and NVIDIA's continuing possession and commercial exploitation of class members' voiceprints in commercial models — and ongoing distribution of those voiceprints through open-weight model releases on Hugging Face — is uniform. Plaintiffs seek classwide injunctive relief under 740 ILCS 14/20, 815 ILCS 510/3 (IUDTPA, which authorizes only injunctive relief), and the equitable jurisdiction of this Court — including the destruction or retraining of the foundational voice synthesis models in which class members' voiceprints are encoded — that necessarily applies on the same terms to every class member.

160. *Rule 23(b)(3) Certification.* Certification under Rule 23(b)(3) is appropriate for the Class on the damages and restitutionary claims asserted in

this Complaint, including the claims under BIPA, IRPA, ICFA, IUDTPA, and Illinois common law for unjust enrichment, because common questions of law and fact predominate over questions affecting only individual members of the Class, and a class action is superior to other available methods for fairly and efficiently adjudicating the controversy.

161. *Predominance.* The questions that drive this litigation are common to the Class and predominate over individual questions. Whether NVIDIA's voice model training pipelines extract voiceprints within the meaning of BIPA is a common technical question with a common answer — NVIDIA's model documentation describes the pipelines as analyzing a speaker's voice to "replicate voice qualities such as pitch, timbre and speech rate" and to "[m]aintain[] the original characteristics that capture unique voice audio signature," and as conditioned on "voice prompts" consisting of "audio tokens that establish the target vocal characteristics and speaking style." Whether NVIDIA complied with BIPA's notice, consent, retention, profiting, and dissemination requirements is a common legal question with a common answer — NVIDIA did not, with respect to any non-user whose voice recordings were ingested into the training pipelines. Whether NVIDIA's conduct was willful or reckless turns on NVIDIA's institutional knowledge, decision-making, and conduct, all of which are common to the Class. The principal individual question — whether a specific class member's voice recordings entered the training pipelines — is binary and resolvable from NVIDIA's own records, which include the training-data manifests, ingestion logs, NeMo Speech Data

Processor pipeline metadata, and source-identifier records in NVIDIA's exclusive control. Individual damages calculations under BIPA's per-person, per-subsection liquidated-damages framework, 740 ILCS 14/20, as amended and as construed in *Clay v. Union Pacific Railroad Co.*, No. 25-2185 (7th Cir. Apr. 1, 2026), do not predominate over the common liability questions because the per-violation amounts are statutorily fixed and do not require individualized proof of actual damages.

162. *Superiority.* A class action is the superior method for adjudicating these claims:

(a) *Class members' interest in individual control.* Class members are voice professionals — journalists, audiobook narrators, podcasters, voiceover artists — and incidental speakers whose voiceprints were extracted without their knowledge. Many class members are unaware their biometric identifiers were ever taken. Even those who become aware face the prospect of individual litigation against the largest publicly traded technology company in the world, with statutory damages amounts that, while meaningful in the aggregate, are likely too modest in individual cases to justify the cost and burden of independent representation. A class action is the only realistic vehicle for redressing the violations alleged in this Complaint.

(b) *Existing related litigation.* Plaintiffs are unaware of any other action asserting these claims against NVIDIA on behalf of persons whose voice recordings were produced or recorded in Illinois and used to train NVIDIA's foundational voice synthesis models. Two related actions are pending against

NVIDIA in the Northern District of California, both arising from related training-data practices within NVIDIA's foundational model families: *Nazemian v. NVIDIA Corp.*, No. 4:24-cv-01454-JST (N.D. Cal. filed Mar. 8, 2024), in which authors allege NVIDIA trained its NeMo Megatron large language models on the Books3 dataset of pirated e-books; and *Ted Entertainment, Inc. v. NVIDIA Corp.*, No. 5:25-cv-10287 (N.D. Cal. filed Nov. 26, 2025), in which YouTube content creators allege NVIDIA bypassed YouTube's technical protections to scrape millions of videos to train its Cosmos foundational video model. Both actions arise from training-data practices within the same NVIDIA research enterprise but assert claims on behalf of authors and copyright holders, not the voice professionals — broadcast journalists, podcasters, audiobook narrators, voiceover artists, and other speakers — whose voices were ingested into NVIDIA's foundational voice synthesis models. Both actions are pending in a different district. Neither asserts claims under BIPA or any analogous biometric-privacy statute.

(c) *Desirability of concentration in this forum.* Concentrating this litigation in this District is appropriate. Plaintiffs are Illinois residents whose recordings were produced or recorded in Illinois. The claims arise under Illinois statutes. The injuries were suffered in Illinois. NVIDIA conducts substantial commercial business in Illinois, including through enterprise licensing of its voice AI products to Illinois-based customers. This Court is well-suited to adjudicate BIPA and other Illinois statutory claims arising from NVIDIA's commercial conduct in Illinois.

(d) *Manageability*. The case is manageable as a class action. NVIDIA's conduct was automated, uniform, and standardized; the common questions identified above are susceptible to common proof; class membership can be determined from NVIDIA's records, supplemented as needed by publicly available distribution metadata and voice-matching analysis; and BIPA's per-person, per-subsection liquidated-damages framework eliminates the need for individualized damages calculations on the principal claim. No unusual management difficulties are anticipated.

163. To the extent any portion of the Class Period predates the limitations period applicable to any claim asserted in this action, Plaintiffs allege that the limitations periods are equitably tolled by NVIDIA's concealment of its training-data sources, by NVIDIA's failure to provide any notice of its collection of biometric identifiers, by Plaintiffs' inability through reasonable diligence to discover that NVIDIA had ingested their recordings into its training pipelines, and by the continuing nature of NVIDIA's violations. Independently, each retention of the unlawfully obtained biometric data, each operation of the foundational audio models in which the unlawfully obtained biometric data is encoded, each disclosure or transmission of that biometric data, and each public release of model parameters encoding that biometric data is a separate violation of BIPA under *Cothron v. White Castle System, Inc.*, 2023 IL 128004 (Ill. 2023), each accruing a separate limitations period from the date of the discrete violative act.

CLAIMS FOR RELIEF

Count I

**Violation of the Illinois Biometric Information Privacy Act,
740 ILCS 14/15(b)**

Brought on behalf of the Class

164. Plaintiffs reallege and incorporate by reference all allegations in this Complaint.

165. Plaintiffs bring this Count individually and on behalf of the Class.

166. BIPA defines "biometric identifier" to include a "voiceprint" and "biometric information" to include "any information, regardless of how it is captured, converted, stored, or shared, based on an individual's biometric identifier used to identify an individual." 740 ILCS 14/10. Section 15(b) prohibits a private entity from collecting, capturing, purchasing, receiving through trade, or otherwise obtaining a person's biometric identifier or biometric information unless the entity first informs the subject in writing that biometric data is being collected or stored, informs the subject in writing of the specific purpose and length of term of collection, and receives a written release executed by the subject. 740 ILCS 14/15(b).

167. NVIDIA is a "private entity" within the meaning of BIPA. 740 ILCS 14/10.

168. NVIDIA collected, captured, and otherwise obtained voiceprints and biometric information from the voice recordings of Plaintiffs and Class members by ingesting their voice recordings into its training pipelines and extracting from those recordings computational representations capable of

identifying the speakers, as alleged in detail at ¶¶ 33-41. The resulting representations — variously denominated as speaker embeddings, audio tokens, voice prompts, and speaker-acoustic-signature representations in NVIDIA's technical and product documentation — are voiceprints and biometric information within the meaning of BIPA.

169. NVIDIA did not, before extracting Plaintiffs' or Class members' voiceprints, inform any Plaintiff or Class member in writing that their biometric identifiers were being collected or stored, did not inform any of them in writing of the specific purpose or length of term of collection, and did not receive a written release executed by any of them. NVIDIA obtained no consent of any kind, in any form, from any Plaintiff or Class member.

170. NVIDIA's violations of § 15(b) were intentional or reckless, as alleged at ¶¶ 125-148. In the alternative, NVIDIA's violations were negligent.

171. Plaintiffs and the Class seek all relief available under 740 ILCS 14/20, including, for each Class member, the greater of liquidated damages or actual damages on a per-person, per-subsection basis consistent with the statute as amended and as construed in *Clay v. Union Pacific Railroad Co.*, No. 25-2185 (7th Cir. Apr. 1, 2026), in the amount of \$1,000 per negligent violation or \$5,000 per intentional or reckless violation; injunctive relief; and reasonable attorneys' fees, costs, and any other relief the Court deems just and proper.

Count II

**Violation of the Illinois Biometric Information Privacy Act,
740 ILCS 14/15(a)**

Brought on behalf of the Class

172. Plaintiffs reallege and incorporate by reference all allegations in this Complaint.

173. Plaintiffs bring this Count individually and on behalf of the Class.

174. Section 15(a) of BIPA requires a private entity in possession of biometric identifiers to develop a written policy, made available to the public, establishing a retention schedule and guidelines for permanently destroying biometric identifiers when the initial purpose for collecting or obtaining such identifiers has been satisfied or within three years of the individual's last interaction with the private entity, whichever occurs first. 740 ILCS 14/15(a). Because Class members never interacted with NVIDIA in connection with the collection of their biometric data, the operative destruction prong for Class members is that biometric identifiers be destroyed when the initial purpose for their collection has been satisfied.

175. NVIDIA has been, and remains, in possession of voiceprints extracted from recordings of Plaintiffs and Class members.

176. NVIDIA has not developed and made publicly available a retention and destruction policy applicable to voiceprints extracted from non-user training data and embedded in the parameters of NVIDIA's foundational voice synthesis models, as alleged at ¶ 60. NVIDIA's general Privacy Policy does not address the retention or destruction of biometric data extracted from non-user

voice recordings, and applies only to persons with an existing relationship with NVIDIA. NVIDIA has not provided, and on information and belief does not maintain, any mechanism by which Plaintiffs or other non-user training-data subjects can request access to, correction of, or deletion of their biometric data.

177. NVIDIA's violations of § 15(a) were intentional or reckless, as alleged at ¶¶ 125-148. In the alternative, NVIDIA's violations were negligent.

178. Plaintiffs and the Class seek all relief available under 740 ILCS 14/20 on the terms set forth in Count I.

Count III

Violation of the Illinois Biometric Information Privacy Act, 740 ILCS 14/15(c)

Brought on behalf of the Class

179. Plaintiffs reallege and incorporate by reference all allegations in this Complaint.

180. Plaintiffs bring this Count individually and on behalf of the Class.

181. Section 15(c) of BIPA provides that no private entity in possession of a biometric identifier may sell, lease, trade, or otherwise profit from a person's biometric identifier. 740 ILCS 14/15(c). The phrase "otherwise profit from" is a statutory catch-all that extends beyond the enumerated forms of selling, leasing, and trading.

182. NVIDIA has profited and continues to profit from Plaintiffs' and Class members' voiceprints by using them to develop, train, and commercially operate the voice synthesis models that power NVIDIA's commercial voice products, and by monetizing those products through the integrated commercial

channels alleged at ¶¶ 73-80, including (i) NVIDIA AI Enterprise subscription fees; (ii) NVIDIA API Catalog cloud-hosted inference fees; (iii) NVIDIA NIM containerized microservice deployments; (iv) the strategic and demonstration value supporting NVIDIA's GPU and AI-hardware sales; (v) developer-ecosystem and commercial-moat value generated through open-weight model distribution on Hugging Face; and (vi) strategic deployment value across NVIDIA's broader product portfolio. The voice quality, expressiveness, speaker similarity, and multilingual capability that NVIDIA sells through these channels exists because of the voiceprints encoded in NVIDIA's foundational voice synthesis models. NVIDIA's commercial exploitation of Plaintiffs' and Class members' voiceprints does not fall within any exception to § 15(c).

183. NVIDIA's violations of § 15(c) were intentional or reckless, as alleged at ¶¶ 125-148. In the alternative, NVIDIA's violations were negligent.

184. Plaintiffs and the Class seek all relief available under 740 ILCS 14/20 on the terms set forth in Count I.

Count IV

Violation of the Illinois Biometric Information Privacy Act, 740 ILCS 14/15(d)

Brought on behalf of the Class

185. Plaintiffs reallege and incorporate by reference all allegations in this Complaint.

186. Plaintiffs bring this Count individually and on behalf of the Class.

187. Section 15(d) of BIPA provides that no private entity in possession of a biometric identifier may disclose, redisclose, or otherwise disseminate a

person's biometric identifier unless an enumerated exception applies. 740 ILCS 14/15(d).

188. NVIDIA has disclosed, redisclosed, and otherwise disseminated Plaintiffs' and Class members' voiceprints in three categories of conduct, each independently sufficient to violate § 15(d) and as alleged at ¶¶ 77-80.

189. *First*, NVIDIA has, on information and belief, transferred the foundational voice synthesis models and the voiceprints encoded within them, among NVIDIA's affiliated corporate entities, including NVIDIA International, Inc., NVIDIA Singapore Pte. Ltd., and Mellanox Technologies, Ltd.

190. *Second*, NVIDIA has, on information and belief, transmitted the foundational voice synthesis models and the voiceprints encoded within them to vendors, service providers, and cloud-infrastructure partners in the ordinary course of training, evaluating, deploying, and operating its commercial voice products.

191. *Third*, NVIDIA has publicly released the parameters of multiple voice synthesis models, including PersonaPlex, Canary, and Parakeet, as open-weight releases on Hugging Face under the NVIDIA AI Open Model License and the NVIDIA Nemotron Open Model License, making the encoded biometric characteristics available for download by any party worldwide. The open-weight release is a unilateral, irreversible act of dissemination by NVIDIA. Plaintiffs and Class members did not consent to any of these disseminations. No enumerated exception under § 15(d) applies.

192. NVIDIA's violations of § 15(d) were intentional or reckless, as alleged at ¶¶ 125-148. In the alternative, NVIDIA's violations were negligent.

193. Plaintiffs and the Class seek all relief available under 740 ILCS 14/20 on the terms set forth in Count I.

Count V

Violation of the Illinois Biometric Information Privacy Act, 740 ILCS 14/15(e)

Brought on behalf of the Class

194. Plaintiffs reallege and incorporate by reference all allegations in this Complaint.

195. Plaintiffs bring this Count individually and on behalf of the Class.

196. Section 15(e) of BIPA requires a private entity in possession of biometric identifiers to store, transmit, and protect the identifiers from disclosure (1) using the reasonable standard of care within the entity's industry, and (2) in a manner that is the same as or more protective than the manner in which the entity stores, transmits, and protects other confidential and sensitive information. 740 ILCS 14/15(e). Section 15(e)(2) imposes an asymmetry test: a private entity satisfies it only if its protection of biometric data is at least equal to its protection of its own other confidential and sensitive information. A generally adequate security posture does not satisfy § 15(e)(2) if the entity protects its other confidential information with greater care than it protects biometric data.

197. NVIDIA fails the § 15(e)(2) asymmetry test. NVIDIA's established institutional practice with respect to a substantial portion of its voice and

speech synthesis models is open-weight release on Hugging Face, while NVIDIA simultaneously protects its other confidential and sensitive information under strict closed-system controls.

198. NVIDIA's open-weight release of voice and speech synthesis models is sustained, public, and intentional. NVIDIA distributes its PersonaPlex conversational voice model, released January 2026, as an open-weight model on Hugging Face under the NVIDIA AI Open Model License. NVIDIA distributes its Canary-1b-v2 and Parakeet-tdt-0.6b-v3 speech models, trained on the approximately one-million-hour Granary speech corpus, as open-weight models on Hugging Face. NVIDIA's NeMo open-source toolkit, distributed on GitHub under the Apache 2.0 license, packages the underlying voice and speech model architectures. NVIDIA's broader Nemotron family of open-weight foundational models is distributed under the NVIDIA Nemotron Open Model License. NVIDIA's AI Open Model License and NVIDIA Nemotron Open Model License are both, in NVIDIA's own description, intended to make NVIDIA's foundational AI models broadly available to the developer ecosystem.

199. The voice synthesis and speech models NVIDIA has distributed as open-weight releases encode the biometric characteristics extracted from training data. Once a model's weights have been published on Hugging Face under a permissive license, the encoded biometric characteristics persist in every downloaded copy. NVIDIA cannot recall the downloaded copies. NVIDIA cannot prevent third parties from redistributing, modifying, or fine-tuning them

on additional data while preserving the underlying biometric encodings. The open-weight release is, by design, an irreversible publication.

200. NVIDIA contemporaneously protects its other confidential and sensitive information under closed-system controls that are materially more protective than open release. NVIDIA maintains the source code for its core commercial products as closed-source proprietary information. NVIDIA protects its proprietary GPU architecture and integrated-circuit designs, its internal financial records, its employee personnel files, its customer data, and its business communications under strict access controls, encryption, internal classification systems, and information-security regimes. NVIDIA protects technology subject to United States export controls administered by the Department of Commerce under additional regulatory protective layers. NVIDIA's decision to withhold FUGATTO from public release in November 2024, based on NVIDIA's recognition of misuse risks associated with the model's voice generation and voice transformation capabilities, is one example of the protective measures NVIDIA applies to its own model weights when it perceives reputational or legal exposure.

201. The asymmetry between NVIDIA's open-weight release of voice and speech synthesis model parameters, including, on information and belief, model parameters derived from training data containing voiceprints and biometric information from non-consenting individuals, and NVIDIA's closed-system protection of its other confidential and sensitive information violates § 15(e)(2)'s requirement that biometric data be protected "in a manner that is

the same as or more protective than" NVIDIA's protection of its other confidential and sensitive information. Open-weight release is publication, not protection. NVIDIA's institutional pattern of open-weight release applies to and has been applied to speech synthesis models that ingested Class members' voice recordings; at minimum, those production models are not protected with the same level of care that NVIDIA applies to its other confidential and sensitive information, because NVIDIA's established public practice with respect to comparable voice and speech models is open-weight release rather than closed protection.

202. Independently, NVIDIA fails the § 15(e)(1) reasonable-standard-of-care test as to voiceprints extracted from non-user training data. The reasonable standard of care within the AI and machine learning industry for the handling of biometric training data includes documented provenance for training data, access controls, audit logging, encryption, a documented retention and destruction schedule, and a mechanism through which data subjects may request access to, correction of, or deletion of their biometric data. As alleged at ¶¶ 49-61, NVIDIA has not published documented provenance for the bulk of the voice training data used to build its commercial voice synthesis models; has not published any access-control protocol, audit-logging policy, encryption standard, or retention and destruction schedule specifically applicable to voice training data or to the biometric characteristics encoded in its voice models; and has not published any mechanism through which non-user data subjects may request access to, correction of, or deletion

of their biometric data from NVIDIA's training datasets, model parameters, or commercial inference pipelines. A private entity cannot adequately protect biometric data belonging to individuals to whom it has never disclosed possession of their data, because those individuals have no ability to monitor, verify, or challenge how their data is stored, transmitted, or protected. The reasonable industry standard of care for biometric data presupposes some mechanism by which the data subject can verify the integrity of the protection; NVIDIA has provided none.

203. NVIDIA's violations of § 15(e) were intentional or reckless, as alleged at ¶¶ 125-148. In the alternative, NVIDIA's violations were negligent.

204. Plaintiffs and the Class seek all relief available under 740 ILCS 14/20 on the terms set forth in Count I.

Count VI

Violation of the Illinois Right of Publicity Act, 765 ILCS 1075/1 et seq.

Brought on behalf of the Class

205. Plaintiffs reallege and incorporate by reference all allegations in this Complaint.

206. Plaintiffs bring this Count individually and on behalf of the Class.

207. The Illinois Right of Publicity Act protects an individual's right to control commercial use of their identity. 765 ILCS 1075/30(a) provides that a person may not use an individual's identity for commercial purposes during the individual's lifetime without having obtained previous written consent from the individual. IRPA defines "identity" to include "any attribute of an individual that

serves to identify that individual to an ordinary, reasonable viewer or listener," and expressly enumerates "voice" among the protected attributes. 765 ILCS 1075/5.

208. Each Plaintiff and each Class member is an individual whose distinctive voice — including timbre, tone, cadence, phrasing, accent, and stylistic vocal expression — is part of the individual's identity within the meaning of IRPA.

209. NVIDIA used Plaintiffs' and Class members' identities for commercial purposes within the meaning of IRPA, by extracting and modeling the distinctive vocal characteristics embodied in their recordings and using those characteristics to develop, train, and operate the commercial voice synthesis products that NVIDIA monetizes through the integrated commercial channels alleged at ¶¶ 73-80. NVIDIA holds out its products' ability to generate realistic, expressive, and human-sounding voices — and, in the case of Magpie TTS Zeroshot and Magpie TTS Flow, their ability to achieve "speaker similarity of over 70%" from a five-second voice prompt — as a core commercial feature of those products, a capability built from the identities of the individuals whose recordings were used to train the foundational models. NVIDIA did not obtain written consent from any Plaintiff or Class member to use their identity, including their voice, for any commercial purpose.

210. Independently, the Illinois Right of Publicity Act, as amended by P.A. 103-836 effective January 1, 2025, prohibits knowingly distributing, transmitting, or making available to the general public a sound recording or

audiovisual work containing an "unauthorized digital replica" of an individual without the individual's consent or the consent of an authorized representative. 765 ILCS 1075/30(b). The amended IRPA further imposes liability on any person who "materially contributes to, induces, or facilitates" such distribution. *Id.* On information and belief, NVIDIA's commercial voice products generate, distribute, and make available to the public voice outputs that constitute unauthorized digital replicas of Plaintiffs and Class members within the meaning of the amended IRPA, and NVIDIA materially contributes to and facilitates the distribution of such replicas through its integrated commercial channels — including NVIDIA AI Enterprise, NVIDIA API Catalog cloud-hosted inference services, NVIDIA NIM microservice deployments, and open-weight model releases on Hugging Face that enable any downloader to generate unauthorized digital replicas.

211. NVIDIA's violations of IRPA were willful and knowing, as alleged at ¶¶ 125-148. NVIDIA's Product Specific Terms for AI Products require its enterprise customers to "have sufficient rights and licenses for content used with a Custom TTS Application to generate new content," and NVIDIA's October 2025 Magpie TTS launch documentation advised users that voice cloning should be used only with "target speakers who have consented to such use." NVIDIA imposed those obligations on its downstream customers and end users specifically because NVIDIA understood that the commercial use of an individual's voice requires that individual's consent. NVIDIA chose not to apply

the same consent infrastructure to Plaintiffs and Class members, whose voices NVIDIA itself extracted to train the foundational models.

212. Plaintiffs and the Class have suffered concrete injury from NVIDIA's IRPA violations, including loss of control over the commercial use of their identities and voices, dilution and commodification of their distinctive voices, and economic harms including the diversion of licensing value and the diminished demand for authentic vocal performances in the markets where Plaintiffs and Class members earn their livelihoods, as alleged at ¶¶ 81-88.

213. Plaintiffs and the Class seek all relief available under 765 ILCS 1075/40, including actual damages, profits attributable to NVIDIA's unauthorized use of Plaintiffs' and Class members' identities, the statutory minimum of \$1,000 per violation under 765 ILCS 1075/40 where actual damages are below that amount, punitive damages, injunctive relief, attorneys' fees pursuant to 765 ILCS 1075/55, and such other relief as the Court deems just and proper.

Count VII

Violation of the Illinois Consumer Fraud and Deceptive Business Practices Act, 815 ILCS 505/1 et seq.

Brought on behalf of the Class

214. Plaintiffs reallege and incorporate by reference all allegations in this Complaint.

215. Plaintiffs bring this Count individually and on behalf of the Class.

216. The Illinois Consumer Fraud and Deceptive Business Practices Act prohibits "unfair or deceptive acts or practices ... in the conduct of any trade or commerce." 815 ILCS 505/2. ICFA's prohibition of unfair practices is broader than its prohibition of deceptive practices and reaches conduct that, even absent affirmative misrepresentation, offends Illinois public policy, is immoral, unethical, oppressive, or unscrupulous, or causes substantial injury to consumers that consumers could not reasonably have avoided. *Robinson v. Toyota Motor Credit Corp.*, 201 Ill. 2d 403, 417 (2002).

217. NVIDIA engaged in trade and commerce in Illinois within the meaning of ICFA by marketing, offering, and distributing voice synthesis services to Illinois customers, and by collecting, extracting, and commercially exploiting the biometric identifiers of Illinois persons in connection with that commerce.

218. NVIDIA's conduct is unfair within the meaning of ICFA. NVIDIA's collection and commercial exploitation of the biometric identifiers of Illinois persons without their knowledge or consent offends Illinois public policy as expressed in BIPA — the Illinois statute enacted specifically to address the unauthorized collection of biometric identifiers, 740 ILCS 14/5 — and as expressed in IRPA — the Illinois statute enacted specifically to address the unauthorized commercial use of personal identity, 765 ILCS 1075/30. NVIDIA's conduct caused substantial injury to Plaintiffs and Class members, including the diversion of licensing income they would otherwise have earned for the use of their voices and the displacement of their professional voice work

in the markets where they earn their livelihoods, as alleged at ¶¶ 81-88. Plaintiffs and Class members could not reasonably have avoided this injury because they had no knowledge that NVIDIA was collecting their biometric data.

219. NVIDIA's conduct is also independently deceptive within the meaning of ICFA. NVIDIA's general Privacy Policy categorically states that NVIDIA "does not use, sell, or share 'sensitive personal information' as defined by California law," even though biometric information is expressly included within that definition under California law. That representation is materially misleading because, on information and belief, NVIDIA was simultaneously ingesting tens of thousands of hours of voice recordings from non-users and extracting biometric voice representations from those recordings without the speakers' knowledge. NVIDIA maintained total opacity about the sources of its voice training data while commercially exploiting that data at scale.

220. NVIDIA's unfair and deceptive conduct proximately caused actual injury to Plaintiffs and Class members, including lost and diminished licensing income, suppressed voiceover and narration rates, diverted opportunities, and loss of control over their biometric data and professional identities. These injuries flow directly from NVIDIA's decision to collect and commercially exploit Plaintiffs' biometric data without authorization, and not merely from the existence of competing AI products.

221. NVIDIA's conduct was willful, knowing, and in reckless disregard of the rights and interests of Plaintiffs and Class members, as alleged at ¶¶ 125-148.

222. Plaintiffs and the Class seek all relief available under ICFA, including actual damages, punitive damages under 815 ILCS 505/10a, injunctive relief, attorneys' fees and costs, and such other relief as the Court deems just and proper.

Count VIII

Violation of the Illinois Uniform Deceptive Trade Practices Act, 815 ILCS 510/1 et seq.

Brought on behalf of the Class (injunctive relief only)

223. Plaintiffs reallege and incorporate by reference all allegations in this Complaint.

224. Plaintiffs bring this Count individually and on behalf of the Class, seeking injunctive relief under 815 ILCS 510/3.

225. The Illinois Uniform Deceptive Trade Practices Act prohibits a person from engaging in conduct that, in the course of business, "causes likelihood of confusion or of misunderstanding as to the source, sponsorship, approval, or certification of goods or services," 815 ILCS 510/2(a)(2), or "causes likelihood of confusion or of misunderstanding as to affiliation, connection, or association with or certification by another," 815 ILCS 510/2(a)(3).

226. NVIDIA's commercial voice products generate voice outputs that sound like real human speakers, including voice outputs that, on information and belief, replicate or closely simulate the distinctive vocal characteristics of

Plaintiffs and Class members. Once generated, these voice outputs can be downloaded, shared, and commercially exploited by NVIDIA's enterprise customers and end users without consumer-facing disclosure that the voice was AI-generated, that the model that generated the voice was built using voiceprints collected without consent, or that the individual whose vocal characteristics are reproduced has not authorized the use. The absence of disclosure creates likelihood of confusion or misunderstanding about whether real persons created, endorsed, sponsored, approved, or have any affiliation with the AI-generated voice content.

227. Plaintiffs and Class members are persons likely to be damaged by NVIDIA's deceptive practices within the meaning of 815 ILCS 510/3. NVIDIA's conduct diverts demand from licensed human voice performances and impairs source attribution and authorization-status disclosure in the voice services markets where Plaintiffs and Class members earn their livelihoods. Actual confusion and actual damages need not be shown to obtain injunctive relief.

228. Plaintiffs and the Class seek preliminary and permanent injunctive relief under 815 ILCS 510/3, including injunctive relief requiring NVIDIA to provide adequate consumer-facing disclosure that voice outputs are AI-generated and that the individuals whose vocal characteristics are reproduced did not authorize the use. If the Court finds that NVIDIA has willfully engaged in deceptive trade practices within the meaning of 815 ILCS 510/3, Plaintiffs also seek reasonable attorneys' fees.

Count IX

Unjust Enrichment (Illinois Common Law)

Brought on behalf of the Class

229. Plaintiffs reallege and incorporate by reference all allegations in this Complaint.

230. Plaintiffs bring this Count individually and on behalf of the Class. This Count is pled in the alternative pursuant to Fed. R. Civ. P. 8(d)(2)–(3). Plaintiffs do not seek duplicative recovery.

231. Under Illinois common law, a defendant is liable for unjust enrichment when the defendant has unjustly retained a benefit to the plaintiff's detriment, and the defendant's retention of that benefit violates fundamental principles of justice, equity, and good conscience.

232. NVIDIA obtained substantial benefits from Plaintiffs' and Class members' voice recordings, voiceprints, and identity attributes without permission and without compensation. These benefits include the avoided costs of licensing or obtaining consent for the voice recordings used to train NVIDIA's foundational voice synthesis models; the product capability and competitive advantage NVIDIA captured by training its models on a diverse corpus of professional human voices, including the voices of Plaintiffs and Class members; and the revenue NVIDIA generates and continues to generate through the commercial exploitation of those models, as alleged at ¶¶ 73-80.

233. NVIDIA obtained these benefits at Plaintiffs' and Class members' expense. Plaintiffs and Class members invested time, talent, training, and

resources to develop their voices and create professional recordings. NVIDIA's unauthorized extraction of voiceprints from those recordings diverted economic value from Plaintiffs and Class members to NVIDIA, and NVIDIA's commercial deployment of products built on those voiceprints continues to compete with and displace Plaintiffs and Class members in the markets where they earn their livelihoods.

234. NVIDIA's retention of these benefits is unjust. NVIDIA's conduct violated the Illinois statutory protections for biometric data (BIPA) and for personal identity (IRPA), each of which expresses the public policy of Illinois that the unauthorized commercial extraction of biometric and identity-related personal attributes is unlawful. NVIDIA's institutional practice confirms the inequity of its retention: NVIDIA built and deployed customer-facing rights-acknowledgment infrastructure for the downstream uses of its voice products that NVIDIA monetizes through enterprise customers, but did not apply equivalent infrastructure to the voices of Plaintiffs and Class members from whom NVIDIA itself took voiceprints without consent. NVIDIA could have pursued a lawful licensing path with respect to Plaintiffs and Class members; it chose not to.

235. Plaintiffs and the Class seek restitution of the benefits NVIDIA has unjustly retained, including disgorgement of profits NVIDIA earned from the unauthorized exploitation of Plaintiffs' and Class members' voiceprints and identities, an accounting of those benefits, and such other equitable relief as the Court deems just and proper.

PRAYER FOR RELIEF

WHEREFORE, Plaintiffs, individually and on behalf of all others similarly situated, respectfully request that this Court enter judgment against Defendant NVIDIA Corporation and award the following relief:

A. *Injunctive Relief.* Permanent and, where appropriate, preliminary injunctive relief requiring NVIDIA to:

(1) Cease the collection, capture, purchase, receipt through trade, or other obtaining of biometric identifiers from voice recordings produced or recorded in Illinois without first providing the written notice, disclosure of specific purpose and duration of collection, and written release that 740 ILCS 14/15(b) requires;

(2) Cease the commercial use, sale, lease, trade, profiting from, or dissemination of voiceprints and biometric information of Plaintiffs and Class members that NVIDIA has already collected without BIPA-compliant consent;

(3) Identify and disclose, by name or other identifying information, the sources of all voice training data used to develop, train, fine-tune, evaluate, or operate NVIDIA's foundational voice synthesis models, including without limitation Magpie TTS Multilingual, Magpie TTS Zeroshot, Magpie TTS Flow, FUGATTO, PersonaPlex, Nemotron 3 VoiceChat, Canary, and Parakeet;

(4) Develop and make publicly available a written retention and destruction policy applicable to voiceprints and biometric information of non-users, in compliance with 740 ILCS 14/15(a);

(5) Destroy all voiceprints and biometric information of Plaintiffs and Class members that NVIDIA obtained without BIPA-compliant consent, and certify the destruction;

(6) Destroy or retrain, without the unlawfully obtained voiceprints and biometric information, the foundational voice synthesis models — and the downstream commercial products built on those models — in which the unlawfully obtained voiceprints and biometric information are encoded, including without limitation Magpie TTS Multilingual, Magpie TTS Zeroshot, Magpie TTS Flow, FUGATTO, PersonaPlex, Nemotron 3 VoiceChat, Canary, and Parakeet, together with any iterative or successor models trained from those models or derived from voiceprints encoded in them; and as to NVIDIA's open-weight model releases on Hugging Face and other public distribution platforms, take all reasonable measures to recall and remove the unlawfully encoded model parameters from public availability, including by issuing takedown requests, ceasing further distribution under permissive licenses, and ceasing further support of derivative releases;

(7) Cease the use of Plaintiffs' and Class members' identities, including their voices, for commercial purposes without prior written consent, in violation of 765 ILCS 1075/30(a), and cease the distribution or making

available of sound recordings or audiovisual works containing unauthorized digital replicas within the meaning of 765 ILCS 1075/30(b); and

(8) Provide adequate consumer-facing disclosure that voice outputs generated by NVIDIA's commercial voice products are AI-generated, are derived from foundational voice synthesis models built using voice recordings of unidentified speakers, and have not been authorized by the individuals whose vocal characteristics are reproduced.

B. *BIPA Damages.* Award each Class member, on each statutory subsection for which Defendant is found liable, the greater of liquidated damages or actual damages on a per-person, per-subsection basis consistent with 740 ILCS 14/20 as amended by P.A. 103-769 (effective August 2, 2024) and as construed in *Clay v. Union Pacific Railroad Co.*, No. 25-2185 (7th Cir. Apr. 1, 2026), in the amount of \$1,000 per negligent violation or \$5,000 per intentional or reckless violation.

C. *IRPA Damages.* Award the actual damages Plaintiffs and Class members sustained as a result of NVIDIA's unauthorized commercial use of their identities and voices, the profits NVIDIA earned from that unauthorized use to the extent not taken into account in computing actual damages, statutory minimum damages where applicable, and punitive damages under 765 ILCS 1075/40.

D. *ICFA Damages.* Award the actual economic damages Plaintiffs and Class members sustained as a result of NVIDIA's unfair and deceptive practices, and punitive damages under 815 ILCS 505/10a(c).

E. *Restitution and Disgorgement.* Order restitution of the benefits NVIDIA has unjustly obtained from the unauthorized collection and commercial exploitation of Plaintiffs' and Class members' voiceprints and identities; disgorgement of the profits NVIDIA has earned through the integrated commercial chain alleged in this Complaint, including without limitation revenue from (i) NVIDIA AI Enterprise subscription fees; (ii) NVIDIA API Catalog cloud-hosted inference fees; (iii) NVIDIA NIM containerized microservice deployments; (iv) GPU and AI-hardware sales supported by the demonstrated commercial capability of NVIDIA's voice synthesis products; (v) developer-ecosystem and commercial-moat value generated through open-weight model distribution on Hugging Face under the NVIDIA AI Open Model License and the NVIDIA Nemotron Open Model License; and (vi) strategic deployment value of voice synthesis capability across NVIDIA's broader commercial product portfolio; and an accounting of all such revenues and benefits.

F. *Non-Duplication.* Plaintiffs do not seek duplicative recovery across Counts. To the extent a particular harm is compensated under one Count, Plaintiffs do not seek to recover the same harm under another Count. Plaintiffs' claims for the same conduct under multiple statutes and theories are pleaded in the alternative pursuant to Federal Rule of Civil Procedure 8(d)(2)–(3); the Court may award the relief that, in its judgment, makes Plaintiffs and Class members whole and addresses the unlawful conduct.

G. *Class Certification.* Certify the Class as defined in this Complaint pursuant to Federal Rules of Civil Procedure 23(b)(2) and 23(b)(3); appoint Plaintiffs as Class Representatives; and appoint Plaintiffs' counsel as Class Counsel.

H. *Judgment.* Enter judgment in favor of Plaintiffs and all Class members and against Defendant NVIDIA Corporation on all Counts on which Defendant is found liable.

I. *Attorneys' Fees and Costs.* Award reasonable attorneys' fees, costs, and litigation expenses under 740 ILCS 14/20 (BIPA), 765 ILCS 1075/55 (IRPA), 815 ILCS 505/10a (ICFA), 815 ILCS 510/3 (IUDTPA, upon a finding of willfulness), and any other applicable fee-shifting provision.

J. *Interest.* Award pre-judgment and post-judgment interest at the maximum rate permitted by law, including the rate available under 815 ILCS 205/2 for pre-judgment interest where applicable.

K. *Further Relief.* Award such other and further relief as the Court deems just, equitable, and proper.

JURY TRIAL REQUESTED

Plaintiffs, individually and on behalf of all other Class members, request a trial by jury on all claims so triable.

Dated: May 12, 2026

LOEVY & LOEVY

/s/ Ross Kimbarovsky

Ross Kimbarovsky (6229590)

ross@loevy.com

Jon Loevy (6218524)

jon@loevy.com

Michael Kanovitz (6275233)

mike@loevy.com

Matthew Topic (6290923)

matt@loevy.com

Aaron Tucek (98624)

aaron@loevy.com

LOEVY & LOEVY

311 North Aberdeen, 3rd Floor

Chicago, IL 60607

312.243.5900 (phone)

312.243.5902 (fax)

*Attorneys for Plaintiffs Philip Rogers, Carol
Marin, Alison Flowers, Robin Amer, Lindsey
Dorcus, Yohance Lacour, and Victoria
Nassif.*