

Can you erase your data after an AI model has learned from it?

Akintunde Idowu

akintundeidowu.com

Keywords

personal-data erasure; AI training; machine unlearning; data protection; accountability

Introduction

Suppose a Nigerian voice actor uploads recordings to ElevenLabs to create a Professional Voice Clone. The company describes this as a dedicated model trained on a larger collection of voice data. Once created, the clone can read words she never recorded while reproducing the sound and manner of her voice.

¹[ElevenLabs, ‘Voice Cloning’](#)

Months later, she decides that she no longer wants the service to hold or use her voice. She selects “Delete voice”. ElevenLabs’ documentation says that deleting a voice is permanent.

²[ElevenLabs, deleting voices](#)

But the word “voice” could refer to several things: the recordings she uploaded, the dedicated clone produced from them, generated audio already created with the clone, or the influence of her recordings on another model.

This is not an allegation that ElevenLabs has failed to delete any of those things. Its public product documents simply provide a useful way to see why erasure becomes more difficult once personal data have entered an AI system. Deleting a file, deleting a user-facing feature and removing what a model learned are separate technical actions. An organisation responding to an erasure request needs to know which one it has carried out.

What does “Delete voice” remove?

Deleting the recordings would remove the source material supplied by the voice actor. That exercise should cover the files visible in her account as well as working copies, cleaned or segmented audio prepared for training, copies held by processors and any backup from which the recordings could return to an active system.

Deleting the clone addresses something else. In ElevenLabs’ terminology, a Professional Voice Clone is a dedicated model created through training and [Additional training that adapts an existing model to particular data, behaviour or a specialised task.](#) It remains in the user’s account and can be used to generate new speech until it is deleted or otherwise made unavailable. The company’s terms state that a user may request deletion of a User Voice Model through the account.

³[ElevenLabs Terms, cl 4\(b\)](#)

Generated audio presents a further issue.

A person may already have downloaded an advertisement, audiobook passage or video narration made with the clone. Removing the clone from the service will prevent later use through that account, but it will not retrieve every copy of an audio file already distributed elsewhere. The organisation needs to separate what it controls from what another person now holds, rather than describe both as a single deletion.

The answer displayed to a user after pressing “Delete voice” should reflect these different layers. A confirmation that the clone has disappeared from the account does not, without further information, explain what happened to the source recordings, internal copies or related model training.

The recording, the clone and the underlying model

ElevenLabs’ own explanation of its two cloning products shows why the technology used affects the erasure analysis. Its Instant Voice Cloning product uses a short recording as a [Input supplied to a model while it generates a result, used to guide that result without changing the model’s trained numerical values.](#) conditioning

signal when speech is generated. The company says this process does not train a custom model or alter the underlying model's weights. Professional Voice Cloning works differently: it trains a dedicated model using a much larger set of recordings.⁴ [ElevenLabs, how voice cloning works](#fn-4) data-citation="ElevenLabs, how voice cloning works">4

If the user deletes an Instant Voice Clone, the main technical question is likely to concern the reference recording and any representation created from it. With a Professional Voice Clone, the organisation must also deal with a model deliberately adapted to reproduce that person's voice. The same button may therefore sit above two different technical processes.

There may also be a larger model used across the service. ElevenLabs' non-EEA terms allow a user to opt out of future use of content for training. They state that, once processed, the content will no longer be used to improve the services or models, except where needed to provide the service, but that the opt-out does not affect earlier uses or materials resulting from them.⁵ [ElevenLabs Terms, cl 4\(i\)](#fn-5) data-citation="ElevenLabs Terms, cl 4(i)">5

Whatever the legal effect of those terms in a particular case, the product design illustrates a useful point: stopping future training and deleting a voice clone are not the same request.

The Nigerian right to erasure

Section 34 of the Nigeria Data Protection Act 2023 gives a data subject the right to obtain erasure of personal data without undue delay. A controller must erase personal data where they are no longer necessary for the purpose for which they were collected or processed, or where the controller has no other lawful basis to retain them.⁶ [Nigeria Data Protection Act 2023, s 34](#fn-6) data-citation="Nigeria Data Protection Act 2023, s 34">6

The Act therefore requires more than a search for the original upload. The organisation must identify the purpose for which each form of the data is still held and the lawful basis on which any continued processing rests. A recording retained to operate an active clone, the trained clone itself, a security log and a copy stored temporarily in a backup do not necessarily serve the same purpose or follow the same retention process.

Withdrawal of consent under section 35 and objection under section 36 may change what processing can continue, but they are separate rights with their own conditions.⁷ [Nigeria Data Protection Act 2023, ss 35-36](#fn-7) data-citation="Nigeria Data Protection Act 2023, ss 35-36">7

An organisation should establish what the person has asked it to do, the lawful basis relied upon and whether any basis remains. It should not treat deletion from an account, withdrawal from future training and erasure of personal data as three descriptions of the same event.

Suppliers remain part of the response. Section 29 requires a controller engaging a processor to ensure that the processor assists it, through appropriate technical and organisational measures, in honouring data-subject rights. Those arrangements must be recorded in writing.⁸ [Nigeria Data Protection Act 2023, s 29](#fn-8) data-citation="Nigeria Data Protection Act 2023, s 29">8

A controller cannot give a complete answer if it does not know which provider stored the recordings, prepared the training data or hosts the resulting model.

When personal data remain in the model

Deleting every accessible copy of the uploaded recording still leaves the trained model to be considered. During training, a model adjusts its [parameters](#) in response to the examples it receives. Removing an example from the training dataset afterwards does not ordinarily undo those adjustments.

That does not mean that every model contains a retrievable copy of every item used to train it. Nor does describing a model as a collection of numbers establish that it is anonymous. The relevant question is what can still be obtained from the model and linked to the person.

The contrast between a general speech-recognition system and the voice actor's Professional Voice Clone helps here. The first may have learned broad language and acoustic patterns from thousands of speakers. The second exists to reproduce the characteristics of one identifiable person's voice. The European Data

Protection Board uses a generative model fine-tuned on recordings of one person to mimic that person's voice as an example of a model designed to provide information about an individual.⁹ The EDPB's opinion is not Nigerian law, but its technical test is useful. It says that a model trained on personal data should not automatically be treated as anonymous. The assessment should consider whether personal data can be extracted directly from the model or obtained through queries, using means reasonably likely to be employed.¹⁰ For a voice clone, the ordinary output may already answer part of that enquiry. The product is expected to reproduce recognisable features of the person's voice in new speech. For a model trained across a much larger population, the organisation may need more deliberate testing to establish whether a person's data can be reproduced, whether membership of the training set can be inferred, and whether access to the model or other available information changes that risk.

Unlearning, retraining and output controls

The appropriate response depends on what the organisation finds. Removing the recordings from active and supplier-held datasets deals with the source data. Excluding them from later training prevents further use. Deleting a dedicated voice clone prevents it from producing new speech through the service. Each step has a defined effect, but none proves that every other layer has been addressed. An output control can block a particular voice, phrase or result from being returned in ordinary use. It may be necessary where immediate protection is required, but it does not establish that the underlying information has been removed. Someone using a different prompt, interface or level of access may still be able to obtain it.

A technique for selectively removing the influence of specified training data from an already trained machine-learning model. Machine unlearning aims to remove the influence of selected training data from a trained model. NIST notes that some approximate methods may avoid retraining the model from the beginning.¹¹ The organisation would still need to test whether the method sufficiently removed the relevant influence. Calling a process "unlearning" does not show that erasure has been achieved.

Retraining without the recordings may provide clearer evidence, although it will not be necessary in every case. The UK Information Commissioner's Office states that removing personal data from a training set does not automatically require every model trained on the data to be erased. Where personal data are contained in a model by design, however, rectification or erasure may require retraining without the data or deleting the model.¹² This is comparative guidance, but it captures why the general speech model and the dedicated voice clone should not receive an identical response.

What the organisation can honestly tell the user

The voice actor should receive more than a message saying that her voice has been deleted. The organisation should be able to explain whether it removed her uploaded recordings and prepared files, obtained deletion confirmation from its processors, deleted the dedicated clone, prevented later training, and assessed any model in which her data may remain. It should also explain the treatment of backups and any relevant retention period.

Doing that requires **A record of where data came from, how they were changed, where they moved and which systems or model versions used them.** data lineage connecting source recordings to prepared datasets, suppliers and model versions. Sections 24 and 28 of the Nigerian Act support that preparation: personal data must be limited to what is necessary and retained for no longer than necessary, while processing likely to create a high risk requires a data privacy impact assessment before it begins.¹³

If the

organisation can only say, “we deleted the recording from your account”, it has answered a question about the visible file. The erasure request asks a wider question: where does this person’s data still exist, and what processing continues because of it? An AI system does not need to forget in the way a person does. The organisation does need to know what it removed, test what remains and describe both accurately.

Footnotes

- ¹ElevenLabs, ‘Voice Cloning’ (accessed 24 September 2026). !©
- ²ElevenLabs, ‘How do I delete voices from My Voices?’ (accessed 24 September 2026). !©
- ³ElevenLabs, ‘Terms of Service’ (31 March 2026) cl 4(b) (accessed 24 September 2026). !©
- ⁴ElevenLabs, ‘Voice cloning: how it works’ (accessed 24 September 2026). !©
- ⁵ElevenLabs, ‘Terms of Service’ (31 March 2026) cl 4(i) (accessed 24 September 2026). !©
- ⁶Nigeria Data Protection Act 2023, s 34(1)(d) and (2), official text (accessed 24 September 2026). !©
- ⁷Nigeria Data Protection Act 2023, ss 35 and 36, official text (accessed 24 September 2026). !©
- ⁸Nigeria Data Protection Act 2023, s 29(1)(b) and (2), official text (accessed 24 September 2026). !©
- ⁹European Data Protection Board, Opinion 28/2024 on Certain Data Protection Aspects Related to the Processing of Personal Data in the Context of AI Models (18 December 2024) para 29 (accessed 24 September 2026). !©
- ¹⁰European Data Protection Board, Opinion 28/2024 on Certain Data Protection Aspects Related to the Processing of Personal Data in the Context of AI Models (18 December 2024) paras 31, 33-45 (accessed 24 September 2026). !©
- ¹¹National Institute of Standards and Technology, ‘Machine unlearning’, sourced from A Vassilev and others, Adversarial Machine Learning: A Taxonomy and Terminology of Attacks and Mitigations (NIST AI 100-2 E2025, March 2025) (accessed 24 September 2026). !©
- ¹²Information Commissioner’s Office, ‘How do we ensure individual rights in our AI systems?’ (accessed 24 September 2026). !©

href="#fnref-12" aria-label="Return to footnote reference 12">!© <li id="fn-13">Nigeria Data Protection Act 2023, ss 24(1)(c)-(d) and 28, official text (accessed 24 September 2026). !©