

Can Data Protection Law Shield the Human Voice? Evaluating the Role of the UK GDPR in Regulating Artificial Intelligence Voice Cloning

Minna Liu

*The University of Law, London Bloomsbury Campus, London, UK
mliu513@hotmail.com*

Abstract. Neural speech synthesis can generate a voice recognisable to a human from seconds of reference audio, and the harm extends beyond fraud to a loss of control over a socially trusted identifier. No United Kingdom statute protects the voice in isolation: copyright, performers' rights and passing off protect a recording, a performance or commercial goodwill, not the speaker. The UK GDPR thus seems the most broadly applicable instrument, but its application remains unclarified. This paper makes the case that 'voice cloning' hides a series of legally different actions, so that each artefact must be classified separately: a recording of an identifiable speaker is, as a rule, personal data; a speaker embedding obtained through a specific technical process is the strongest candidate for biometric data; and general model weights should not be assumed to be biometric data. Whether a given cloning process is 'for the purpose of uniquely identifying a natural person' has never been answered, and the paper argues for an objective, stage-based approach to Article 9, under which target enrolment, identity conditioning and similarity testing trigger its application but generic synthesis does not. It then demonstrates that the United Kingdom, not bound by the EU AI Act but with regulatory guidance that aligns closely with it, can provide considerable protection today through Articles 5, 6, 9, 22A to 22D, 25, 32, 35 and 82, though no domestic duty to mark synthetic audio exists. Data-protection law can protect the voice as data, but not as persona.

Keywords: UK GDPR, voice cloning, special category data, generative AI, privacy

1. Introduction

AI voice cloning is now able to replicate a familiar voice from just a few seconds of sampled audio. Fraud dominates the policy debate: in the United Kingdom, fraud losses are estimated at approximately £1.28 billion in 2025, including £576.4 million in authorised push-payment fraud [1], and organised criminal groups are assessed as using voice cloning for payment diversion [2]. These statistics refer to an attack surface, not the prevalence of cloning, and evidence is less robust for voice cloning, with 28% of surveyed United Kingdom adults saying they had been targeted by a voice-cloning scam, though the figure is self-reported and was produced as part of a public-awareness campaign [3]. Enforcement systems record the downstream act rather than the upstream

capture and modelling of a person's acoustic features, and it is that upstream stage which a data-protection analysis addresses.

The technical barrier has also fallen. Representative neural systems build a fixed-dimensional speaker representation from a few seconds of reference speech and use this to condition a speech synthesiser, or adapt a multi-speaker model from a small number of target speech samples [4, 5].¹ What matters legally is that such a short recording can be transformed into a computational representation useful both for speaker recognition and for the reproduction of a specific speaker. The damage is not just monetary; a familiar voice is a social identifier, by means of which listeners assign speech, intent and responsibility to a person, and synthetic reproduction disturbs the speaker's control over how they are presented to listeners [6].

United Kingdom law responds to the harm in an uneven way. Copyright and performers' rights protect a recording or a performance, not the sound of a voice [7], and passing off was successful in *Irvine v Talksport Ltd* because the unauthorised use implied commercial endorsement and exploited goodwill [8]. Defamation, misuse of private information and the fraud and communications offences all operate after the damaging use, and not at the point at which the voice is extracted. The United Kingdom General Data Protection Regulation (UK GDPR) seems to be, by contrast, rather more general [9], in that a recording linked to an identifiable living speaker will typically constitute personal data under Article 4(1); specific technical processing that extracts characteristics allowing or confirming unique identification may produce biometric data under Article 4(14); and biometric data processed for the purpose of uniquely identifying a natural person is covered by Article 9(1). These are separate thresholds and not a continuous path from 'voice' to 'special category data'. The same breadth is a weakness, as a broad understanding of personal data lends itself to data protection becoming 'the law of everything' [10].

This paper holds that each artefact in a cloning pipeline must be classified separately; that Article 9(1) should be read objectively and stage by stage, so that target enrolment, identity conditioning and similarity testing engage it while generic synthesis does not; and that safeguards similar to those of the EU AI Act, to which the United Kingdom is not bound, exist under the UK GDPR. It is a doctrinal approach, with the UK GDPR being retained and amended by the Data Protection Act 2018 [11] and by the Data (Use and Access) Act 2025 [12], whose data-protection provisions were in force by 19 June 2026 [13],² and with the guidance of the Information Commissioner's Office (ICO) being considered authoritative regulatory guidance, rather than a substitute for the text of the legislation.

2. Classifying voice data under the UK GDPR

2.1. From personal data to biometric data

Article 4(14) classifies personal data resulting from a particular kind of technical processing, so the object of analysis is not the voice as such, but each artefact (recording, speaker representation, adapted model or synthetic output), and each has to pass the thresholds in turn. Under Article 4(1),

¹ Two architectures should be distinguished. A speaker-encoder system converts reference audio into a fixed-dimensional embedding used to condition a synthesiser, whereas an adaptation system fine-tunes part of a generative model on target samples. The analysis below does not assume that every commercial service uses one of them.

² Article 11A, inserted by section 74 of the Data (Use and Access) Act 2025, enables the Secretary of State to create new categories of special category data by regulation. It did not create a category of vocal identity.

personal data is any information relating to an identified or identifiable natural person; in *Nowak v Data Protection Commissioner*, information may relate to an identifiable natural person by content, purpose or effect [14]; and *Breyer v Bundesrepublik Deutschland* considers identifiability to be contextual [15]. Thus a recording of a known speaker is generally personal data, but an anonymous corpus, a generic voice that refers to no one, or a composite voice that cannot reasonably be attributed to a living person may not be. Being publicly available does not change this classification, as it is new processing to reuse a public recording, and access is not permission. A target-specific synthetic output can also sometimes be considered personal data where it is deployed so that listeners would associate it with a living person: purpose or effect may then make it information relating to that person, and falsity may intensify rather than remove that effect.

The second threshold is Article 4(14): biometric data is personal data obtained through specific technical processing of physical, physiological or behavioural characteristics which allows or confirms the unique identification of the person. The requirements are cumulative: being recognised by a human listener is not enough. The category is engaged when software processes the vocal characteristics and converts them into a representation that can be matched against a new sample, paradigmatically a speaker embedding, a fixed-length representation that carries speaker-related information [4]. Cloning can be done with an identity-bearing template, as such representations are reused to condition synthesis, even though their product is new speech. The definition is technology-neutral, the test is functional, and general model weights cannot be considered the biometric data of all the training speakers absent evidence that any identity-bearing information is retained in those weights.

The third threshold is Article 9(1), which adds the words 'for the purpose of uniquely identifying a natural person'. As a result, the ICO distinguishes between the features which make information biometric data and the purposes for which it is used, which make it special category data [16]. A cloning system uses an identity-capable representation to replicate a known target's voice, and whether preserving that identity is itself a purpose of unique identification has not been decided. The purpose language of Article 9 cannot be replaced by capability under Article 4(14), but a controller should not escape Article 9 by describing identity-preserving processing as 'generation'.

2.2. The activity being regulated and the scope of the regulation

The scope is not dependent on the label used on a file, but rather on the operation. The definition of processing in Article 4(2) is broad enough to encompass the collection, alteration, retrieval and use of the data, and because a classification can vary across a single pipeline, even a short-lived feature vector is considered to be processed. The actor needs to be identified alongside the operation, since a foundation-model developer, a cloning platform and an end user may each be a controller, joint controller or processor for different operations, and the purpose must be attributed to whoever determines it, not to 'the AI system'. Two limits follow. Article 9 is not the gateway to the UK GDPR as a whole: an operation which processes identifiable recordings without an identification purpose is still subject to Article 6 and to the principles and rights, so the level of protection is different but the regulatory relationship remains. The domestic-purposes exclusion, on the other hand, applies to the specific actor: a commercial platform cannot inherit a private user's exclusion, and territorial scope must be determined for the specific processing and not for the platform as a whole.

3. The limits of the UK GDPR across the voice-cloning lifecycle

3.1. Objectively reading Article 9, stage by stage

It is best to evaluate legal consequences on a stage-by-stage basis. Article 9 should attach to target enrolment, identity conditioning and similarity testing, that is, to any system using a deliberate isolation, storage and testing of the vocal identity of a living target, and not to voice recordings, general model parameters or synthetic audio at large.

At the collection stage, accessibility does not change the status of a recording, and the manifestly public condition in Article 9(2)(e) should be read narrowly: speaking in public discloses a word or a performance, but it does not create a licence to build a reusable template. Covert collection also undermines reliance on legitimate interests and, according to the ICO, web-scraping for generative AI falls under the category of invisible processing, which struggles to meet the balancing test [17]. A general permission for 'service improvement' does not extend to the development of a reusable target voice; Article 9 requires more, since where a service is established in order to create an identity-bearing voice of a living person, explicit consent should generally be required.

The disputed question arises directly at the stage of representation. A narrow interpretation of Article 9(1) would seem to require an authentication or search function; on that view the system already knows the target and merely determines how to produce a sound, so Article 9 is not engaged. That reading is consistent with the additional purpose wording, but it is too restrictive for a target-oriented pipeline. Identification is an accurate singling out and does not necessarily require the person's civil name, and the ICO considers that processing carried out for the purpose of identifying a specific person with accuracy is biometric recognition processing [16]. The commercial objective is reproduction, but in the process of reproduction the selection of one particular representation of the target is a constituent purpose, and not an alternative operation, and a controller cannot avoid Article 9 by adopting a higher level of product description.

Similarity testing provides the best evidence at the synthesis stage. In a service where the audio generated is compared with the enrolled representation, the comparison is closely similar to biometric verification, and both the comparison and the data it produces should be placed under Article 9 too. The presence of a score should not be a determining factor for engagement, as technical documentation, evaluation metrics and user interface design reveal whether accurate target selection is an operational purpose. That evidence-based inquiry allows formal evasion to be avoided, while leaving general speech research, accessibility tools and fictional or composite voices within the ordinary regime.

At the output and use stage, purpose limitation ensures that a recording made for broadcast is not automatically available for enrolment [17]. Functional persistence, and not the file label, must be the basis for data minimisation, because if a reusable identity template is still in existence after the raw audio has been deleted, little has been achieved. This is not a matter of acoustic fidelity but of accuracy: a convincing clone may be inaccurate in law, because it can be found to have put words in the mouth of someone who never spoke them, and labelling will not turn unlawful upstream processing into lawful processing.

3.2. Where the protection ceases

The UK GDPR achieves a great deal, regulating collection and derivation, retention and deployment, before the material is used for fraud or publication. It cannot confer ownership of a voice or a veto on synthetic resemblance, and it cannot hold where identifiability, controllership or a

territorial connection is hard to establish. On collection, the rules are adequate on paper, so the issue is one of enforcement and not of doctrine: as Section 5.1 demonstrates, covert scraping by overseas services is hard to detect and hard to reach, and ICO guidance is not binding law.

The issue at representation is classification, and the threshold of purposes, as long criticised, determines the level of protection by reference to the controller's purpose and not the nature of the data [18]. This has led to what is called a protection paradox. An operator that presents its service as entertainment is likely to attract the lighter regime, although its output remains recognisable enough to deceive, while the candid operator attracts the stricter one. Once identity-bearing information has passed into the parameters of the model, rights must be contested against the model itself. The ICO does not agree that the obligation is extinguished because of technical difficulty, and the European Data Protection Board requires case-by-case assessment of model anonymity [17, 19], but neither guarantees erasure from a trained model.

The balance tips back the other way on synthesis, as the UK GDPR is applicable to each generation event but not to the system that generates the speech as such: it lays down no requirement for the marking or watermarking of synthetic speech, or for making it detectable. That system-level duty is provided for by Article 50 of the EU AI Act, which does not itself create a legal basis for the underlying processing, so that both regimes are complementary and parallel [20, 21]. The difference is geographical: a service used only by persons in the United Kingdom might fall outside the Act, and there is no similar marking obligation in the United Kingdom.

Conceptually, the mismatch is to be found at output. A waveform may be personal data where it is named or attributed to the target, less certainly if the voice is a composite; and an output can deceive, or injure the relatives of a deceased speaker, without containing personal data at all, since a fictional voice exploits the realism of speech and the deceased is not a data subject. This is further exacerbated by the United Kingdom's regulator-led approach, under which existing regulators apply the principles of safety, transparency, fairness, accountability and contestability [22] within their own sectors. In addition, the House of Lords Library noted in 2026 that the United Kingdom has yet to establish general AI legislation or an AI regulator [23]. To broaden the scope of personal data into a general law of synthetic identity is no solution, and will only re-open the debate about data protection becoming 'the law of everything' [10]. The challenge is to coordinate.

4. Coordinated safeguarding based on the UK GDPR

4.1. The EU AI Act as a mirror (and not a source of obligations)

The status of the EU AI Act in a United Kingdom analysis needs to be clearly defined, since its structure can be easily copied and its force quietly assumed. Regulation (EU) 2024/1689 is not a part of domestic law and does not apply in the United Kingdom. It reaches a United Kingdom actor only where that actor places a system on the Union market or puts it into service in the Union, or where the output of the system is used in the Union [20].³ It does not apply to a provider serving only users in the United Kingdom, and no domestic instrument provides for a similar requirement for synthetic audio to be marked or disclosed, so transplanting Article 50 into United Kingdom practice would assert an obligation that does not exist here.

That does not mean the Act is irrelevant. The concrete AI-governance measures adopted by the ICO all track it closely. The ICO's response to the generative-AI consultation series assigns

³ The extraterritorial limb in Article 2 turns on the output being used in the Union, not on the mere global accessibility of a service.

responsibility across the supply chain in terms that track the Act's provider and deployer roles, and rejects the notion that a developer becomes a passive processor merely because a customer selects the immediate use [17]. Its guidance on AI and data protection structures expectations in terms of risk assessment, transparency, accuracy, human oversight and bias testing, which is also the governance vocabulary of the Act [24]. Its biometric guidance draws a distinction between the qualities of the data and the purpose of the use, and the Act does the same in separating system capability from intended purpose [16].

The Act therefore serves as a mirror. It illustrates the requirements of a fully developed system-level regime: documentation of provenance, risk management, technical robustness, human supervision and transparency of outputs. What it cannot provide is the legal force, which needs to be derived from the UK GDPR, the Data Protection Act 2018, the Data (Use and Access) Act 2025 and the ICO guidance arising from those. Unless Parliament passes a United Kingdom AI Act, each one of the measures below has to be stated as a duty supported by an existing provision of the UK GDPR, and Section 4.2 accordingly pairs each idea drawn from the Act or from ICO guidance with the provision that makes it enforceable here. That anchoring will be decisive for the question whether a proposal is only a recommendation to be directed to Parliament or an obligation which a controller must fulfil in any event.

4.2. A multi-tiered structure of data, model, output and redress governance

The following four safeguards are not sequential, but rather parallel, each addressed to the actor with effective control. The provider controls the architecture, the design of the enrolment, the retention defaults, the marking capability and the technical abuse controls; the deployer controls the script, the context, the audience and the decision to publish; and an end user who processes a target voice beyond a purely personal activity may itself become a controller. Even if responsibilities are documented in a contract, that documentation does not replace the factual test [17].

4.2.1. Provenance and consent at the point of collection

The data-governance and provenance requirements of the Act are reflected in the following provisions of the UK GDPR: Articles 5(1)(a) and (b), 6, 9(2)(a) and 12 to 14. Before constructing any target representation, a provider should always seek proof of source, proof of target and proof of the user's authority to make the representation, whether through direct enrolment by the speaker, an evidenced mandate or a licence covering synthetic use; an unverified tick-box declaration provides no basis. In the case of a large-scale corpus collection undertaken by a general model developer, the corresponding responsibility is documentary: a record of the sources, the methods employed, the intended speech uses, and the reasons for any special-category processing.

Notices to be given under Articles 12 to 14 must be layered. At the point of collection, a single, prominent disclosure should include the service where the voice will be used, the purposes of the processing and the cooperating entities, together with a standing entry point at which the speaker can withdraw consent; a fuller second layer could then cover retention, recipients, model effects, rights and safeguards. Consent must be specific to enrolment, public release, commercial licensing and further training, and refusal of an optional reuse must not result in denial of an unrelated core service. Article 7(3) (withdrawal of consent) and Article 21 (objection to legitimate-interests processing) are separate avenues, both of which require a working interface.

4.2.2. Governance of identity-bearing model components

The Act's stipulations on technical documentation, robustness and risk management have a direct domestic equivalent in Article 25, which renders architecture a legal issue. Wherever it is technically possible, the general speech model should be separated from the target-specific representations, and the provider should keep versioned records of which data created or updated each target component. That modularity makes access, rectification, erasure and audit possible, and challenges the notion that one voice is inseparable from a foundation model; it moves the European Data Protection Board's call for case-by-case assessment of model anonymity, and the ICO's expectation of rights mechanisms for both training data and models, from aspiration to operational control [17, 19].

Article 32 carries the security half of the same safeguard. A stolen target representation can enable unlimited new speech, which is a different risk from the disclosure of a single recording, and appropriate technical and organisational measures should include restricting the export of target adapters, rate limiting, anomaly detection and step-up verification for large-scale generation. Article 5(1)(c) stipulates that redundant representations of identities and raw enrolment audio should not be stored unless there is a documented need to do so. Article 35 supplies the assessment, as target-specific cloning is a paradigm case of processing likely to result in high risk, and because it is necessary to record the residual risk, the identity-bearing components should be treated as high-impact assets. Article 5(1)(a) makes fairness a processing obligation, so the assessment should also test whether enrolment and synthesis perform differently across accents, languages, speech impairments, ages and gender presentations, since fairness also means non-discrimination [25].

4.2.3. Traceability and transparency of synthetic output

The AI Act splits the output duty in two under Article 50. Providers need to employ technical solutions that are 'effective, interoperable, robust and reliable' to ensure that synthetic outputs can be detected so far as is technically feasible, and professional deployers must explicitly indicate qualifying deepfakes [20]. The Commission's code of practice on transparency of AI-generated content adopts a similar two-part division [26].

A similar division could be provided under the UK GDPR outside the scope of the Act. Article 5(1)(a) makes fair and transparent processing an obligation owed to the speaker, and Article 25 requires that obligation to be built into the design of the processing as opposed to merely announced afterwards. That is not speculation: as demonstrated in Section 5.1, a European supervisory authority has already held that the absence of any measure making such an alteration perceptible was itself a breach of data protection by design. Transparency should also differentiate between the target and the audience: the collecting controller owes the target information about the collection, the modelling and the rights, and the deployer owes the audience, at the point of dissemination, notice of the fact that the sound is synthetic and of who has published it.

4.2.4. Dynamic withdrawal, automated decisions and redress

An enrolled voice supports repeated generation, so a decision made at the point of collection cannot govern each subsequent use, hence the need for dynamic withdrawal through the standing entry point established there. In response to an Article 7(3), 17 or 21 request, the provider should stop the creation of new generations, revoke exposed credentials and initiate a documented review of stored recordings, embeddings, adapters and model influences, and should also inform downstream professional deployers, who should then discontinue their distribution where it is under their control.

A central voice register would establish a new resource for biometric surveillance, so provider-level suppression records, resting on a lawful basis and strictly purpose-limited, are the safer design.

The Data (Use and Access) Act 2025 replaced Article 22 with Articles 22A to 22D, which now provide safeguards for significant decisions made by solely automated means [13]. Normally, synthetic speech is not a decision about the target, but an output, and therefore not a significant decision. The safeguards engage where synthetic or inferred voice data drives an automated decision on employment, credit, policing or access, typically where a cloned voice is used to defeat a voice-authentication system. But a gap remains, as a deeply damaging impersonation could be completely automated without being covered by the protections.

A system should exist for target reports and rights requests, with the option for a provider to disable the profile and to stop distribution until authority and risk are determined, and then to provide a final response stating what was done and how that decision may be challenged [13]. The external routes are: compensation for material or non-material damage under Article 82, which is given effect by sections 168 and 169 of the Data Protection Act 2018; enforcement and penalty notices issued by the ICO under sections 142 to 155 of that Act, exercising the corrective powers of Article 58(2) and the fines of Article 83; and the offence under section 170 of that Act of knowingly or recklessly obtaining or disclosing personal data without the controller's consent, alongside fraud where a cloned voice is dishonestly deployed. The ICO should therefore produce specific guidance mapping, for each liable actor, the route from complaint to internal redress, regulatory enforcement and judicial remedy, so that ex ante controls operate prior to the occurrence of harmful publication.

5. Conclusion

5.1. The attitude of the judiciary and regulators

Whether the enrolment of a target speaker and the conditioning of a synthesiser on the target speaker's representation amount to the processing of biometric data 'for the purpose of uniquely identifying a natural person' is not a question which is answered by a judgment in the United Kingdom or in Europe. That silence is itself a finding, and an attitude must be reconstructed from the authority that exists. There are two decisions that are important: one judicial and one regulatory. In *Information Commissioner v Clearview AI Incorporated* [27] the Upper Tribunal overturned the First-tier Tribunal's decision that the processing carried out by a company which scraped public images into facial vectors and sold search access abroad fell outside the material scope of the UK GDPR. The processing at issue was the conversion of public content into an identity-capable template, which is the vocal pipeline in another medium; the issue was scope instead of sensitivity, two tribunals arrived at different answers to that question, and no final relief has yet issued.⁴ A similar instability would follow a claim against a foreign cloning provider. In *R (Bridges) v Chief Constable of South Wales Police*, it was the surrounding discretion and the impact assessment that were criticised, not the technology [28].

The order of the Italian Garante of 23 July 2026, on the use of a synthetic voice to dub studio footage of a journalist for a satirical programme, was, as reported, issued under Articles 5(1)(a) and 25 of the European Union General Data Protection Regulation [29]. Four findings are relevant to Section 4: a paid licence over the footage does not permit alteration to the detriment of the person

⁴ Permission to appeal to the Court of Appeal was granted in December 2025, so the scope question remains open at the date of writing.

filmed, who is not a party to that contract; professional prominence creates no expectation of a synthetic performance; the derogation for satire covered the processing but not the manner in which it was carried out; and disclosure was assessed against an average or inattentive viewer, so that the imperceptibility of the alteration was itself a breach of data protection by design. The authority imposed a prohibition under Article 58(2)(f) and a reprimand under Article 58(2)(b), but no fine, because the question was novel. The reported annulment by the Court of Rome of a Garante fine imposed on a generative-AI provider [30] illustrates the uncertainty of enforcement. Supervisory authorities will treat the processing behind synthetic identity content as unlawful; they do not thereby confer on the speaker any right in the voice itself.

Article 9 is interpreted broadly and its derogations narrowly. *Meta Platforms Inc v Bundeskartellamt* ruled that the processing of data capable of revealing a protected characteristic is prohibited even if the controller does not intend to use it in that way, and that information is not manifestly made public merely because someone used the service, but only when the data subject makes it accessible to the general public through an explicit act [31], a position confirmed in *Schrems v Meta Platforms Ireland Ltd* [32]. A speaker who is enrolled has not manifestly made the resulting technical representation public, whatever the speaker has made public by speaking, and therefore Article 9(2)(e) cannot be extended to cover target enrolment. None of these judgments addresses the question on which this paper turns, since each concerns data revealing a sensitive attribute rather than processing which preserves an identity.

5.2. The voice as data and the voice as persona

Data-protection law can help to protect the voice as data, but it cannot help to protect the voice as identity or persona. It protects the voice when it is captured, when it is converted to an identity-bearing representation, and when it is stored and used again, and it does this for the ordinary speaker just as much as for the speaker who has commercial value. The UK GDPR does not contain a right to a voice, nor any general right to object to imitation or false attribution, and it is silent where identifiability, controllership or a territorial connection is absent.

The problem is that two failures are intertwined and not alternative. Classification uncertainty reduces the protection available: identifiability is relative and is now the subject of legislative revision; a composite or generic voice can cause harm without reaching the threshold set in Article 4(1); and a description of its own service by a controller can shift the same processing from the ordinary regime to the special-category regime and back, which is the protection paradox of Section 3.2. Even the best protection will fall short where there is conceptual mismatch, since access, objection and erasure attach to retained artefacts rather than to resemblance. The two failures are cumulative. A claimant must win the undecided purpose argument under Article 9(1) before reaching a remedy that was never designed to answer the injury complained of; and because that remedy is known to be ill-fitting, there is little incentive to litigate the classification, and thus neither question is authoritatively resolved. Enforcement is precarious, as the jurisdictional instability in *Clearview* and the lack of any penalty in the Italian order both show. The injury suffered by a speaker is an injury to social attribution, while the only available instrument measures informational control.

5.3. A transitional bridge, not a destination

Data protection is a bridge, not a destination. The classification provides a legal floor and the Section 4 safeguards a ceiling, and both can be put into effect now, without waiting for a United

Kingdom AI Act.

The durable solution is likely to lie in a personality right, in an expansion of performers' rights, or in horizontal AI legislation, each of which reaches the interest that data protection cannot: the claim of the speaker to be the only one speaking as that person. There are three limitations: that the analysis is doctrinal and is up to date as of 12 August 2026; that the reading of Article 9 offered here has not been tested in court; and that the empirical base is weak, as enforcement systems record downstream fraud and impersonation and not upstream modelling. A voice is more than just information about a person, and the difficulty this paper has traced is that data-protection law can, at best, treat it as though it were.

References

- [1] UK Finance. (2026, June 15). *Fraud remains a national security threat as criminals steal almost £1.3 billion*. <https://www.ukfinance.org.uk/news-and-insight/press-release/fraud-report-2026-press-release>
- [2] Home Office. (2026, March 9). National assessment centre fraud assessment 2025. *GOV.UK*. <https://www.gov.uk/government/publications/national-assessment-centre-fraud-assessment-2025>
- [3] Starling Bank. (2024, September 18). *AI voice cloning scams could catch millions out: With over a quarter of UK adults targeted in the past year*. <https://www.starlingbank.com/news/starling-bank-launches-safe-phrases-campaign/>
- [4] Jia, Y., Zhang, Y., Weiss, R. J., Wang, Q., Shen, J., Ren, F., Chen, Z., Nguyen, P., Pang, R., Lopez-Moreno, I., & Wu, Y. (2018). Transfer learning from speaker verification to multispeaker text-to-speech synthesis. *Advances in Neural Information Processing Systems*, 31, 4485-4495. <https://proceedings.neurips.cc/paper/2018/hash/6832a7b24bc06775d02b7406880b93fc-Abstract.html>
- [5] Arik, S. O., Chen, J., Peng, K., Ping, W., & Zhou, Y. (2018). Neural voice cloning with a few samples. *Advances in Neural Information Processing Systems*, 31, 10040-10050. https://papers.nips.cc/paper_files/paper/2018/hash/4559912e7a94a9c32b09d894f2bc3c82-Abstract.html
- [6] McGettigan, C., Bloch, S., Bowles, C., Dinkar, T., Lavan, N., Reus, J. C., & Rosi, V. (2025). Voice conversion and cloning: Psychological and ethical implications of intentionally synthesising familiar voice identities. *Journal of the British Academy*, 13(3), Article a31. <https://doi.org/10.5871/jba/013.a31>
- [7] Copyright, Designs and Patents Act 1988, c. 48 (UK). <https://www.legislation.gov.uk/ukpga/1988/48/contents>
- [8] *Irvine v Talksport Ltd* [2003] EWCA Civ 423. <https://www.bailii.org/ew/cases/EWCA/Civ/2003/423.html>
- [9] Regulation (EU) 2016/679 (General Data Protection Regulation), as retained and amended in UK law. <https://www.legislation.gov.uk/eur/2016/679/contents>
- [10] Purtova, N. (2018). The law of everything: Broad concept of personal data and future of EU data protection law. *Law, Innovation and Technology*, 10(1), 40-81. <https://doi.org/10.1080/17579961.2018.1452176>
- [11] Data Protection Act 2018, c. 12 (UK). <https://www.legislation.gov.uk/ukpga/2018/12/contents>
- [12] Data (Use and Access) Act 2025, c. 18 (UK). <https://www.legislation.gov.uk/ukpga/2025/18/contents>
- [13] Information Commissioner's Office. (2026, June 19). The data use and access act 2025 (DUAA): Summary of the changes to data protection law. *Information Commissioner's Office*. <https://ico.org.uk/about-the-ico/what-we-do/legislation-we-cover/data-use-and-access-act-2025/the-data-use-and-access-act-2025-duaa-summary-of-the-changes/>
- [14] *Nowak v Data Protection Commissioner*, Case C-434/16, ECLI:EU:C:2017:994 (2017). <https://curia.europa.eu/juris/liste.jsf?num=C-434/16>
- [15] *Breyer v Bundesrepublik Deutschland*, Case C-582/14, ECLI:EU:C:2016:779 (2016). <https://curia.europa.eu/juris/liste.jsf?num=C-582/14>
- [16] Information Commissioner's Office. (n.d.). Key data protection concepts. *Information Commissioner's Office*. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/biometric-data-guidance-biometric-recognition/key-data-protection-concepts/>
- [17] Information Commissioner's Office. (n.d.). Information Commissioner's Office response to the consultation series on generative AI. *Information Commissioner's Office*. <https://ico.org.uk/about-the-ico/what-we-do/our-work-on-artificial-intelligence/response-to-the-consultation-series-on-generative-ai/>
- [18] Kindt, E. J. (2018). Having yes, using no? About the new legal regime for biometric data. *Computer Law & Security Review*, 34(3), 523-538. <https://doi.org/10.1016/j.clsr.2017.11.004>

- [19] European Data Protection Board. (2024, December 18). Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models. https://www.edpb.europa.eu/documents/opinion-of-the-board-art-64/opinion-282024-on-certain-data-protection-aspects-related-to_en
- [20] Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). <https://eur-lex.europa.eu/eli/reg/2024/1689/oj?locale=en>
- [21] Holtz, H. M., & Ledendal, J. (2026). AI data governance: Overlaps between the AI Act and the GDPR. *Law, Innovation and Technology*, 18(1), 380-409. <https://doi.org/10.1080/17579961.2026.2633677>
- [22] Department for Science, Innovation and Technology. (2024, February 6). A pro-innovation approach to AI regulation: Government response. *GOV.UK*. <https://www.gov.uk/government/consultations/ai-regulation-a-pro-innovation-approach-policy-proposals/outcome/a-pro-innovation-approach-to-ai-regulation-government-response>
- [23] Ostridge, H. (2026, May 20). AI regulation in the UK: Debate on the need for cross-sector legislation. *House of Lords Library*. <https://lordslibrary.parliament.uk/ai-regulation-in-the-uk-debate-on-the-need-for-cross-sector-legislation/>
- [24] Information Commissioner's Office. (n.d.). *Guidance on AI and data protection*. <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/>
- [25] Information Commissioner's Office. (n.d.). *What about fairness, bias and discrimination?* <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/guidance-on-ai-and-data-protection/how-do-we-ensure-fairness-in-ai/what-about-fairness-bias-and-discrimination/>
- [26] European Commission. (2026, June 10). *Code of practice on transparency of AI-generated content*. <https://digital-strategy.ec.europa.eu/en/policies/code-practice-ai-generated-content>
- [27] Information Commissioner v Clearview AI Incorporated [2025] UKUT 319 (AAC). <https://caselaw.nationalarchives.gov.uk/ukut/aac/2025/319>
- [28] R (Bridges) v Chief Constable of South Wales Police [2020] EWCA Civ 1058. <https://www.judiciary.uk/judgments/r-bridges-v-cc-south-wales/>
- [29] Garante per la protezione dei dati personali. (2026, July 23). Provvedimento n. 577 (doc. web n. 10281021), reported in Rijo, L. (2026, August 7), Mediaset faces 30-day deadline after Italy bans Mentana deepfakes. *PPC Land*. <https://ppc.land/mediaset-faces-30-day-deadline-after-italy-bans-mentana-deepfakes/>
- [30] Cross-Border Data Forum. (2026, March 19). *Generative AI and GDPR enforcement in Europe: A lot of noise, one fine, zero survivors*. <https://www.crossborderdataforum.org/generative-ai-and-gdpr-enforcement-in-europe-a-lot-of-noise-one-fine-zero-survivors/>
- [31] Meta Platforms Inc v Bundeskartellamt, Case C-252/21, ECLI:EU:C:2023:537 (2023). <https://curia.europa.eu/juris/liste.jsf?num=C-252/21>
- [32] Schrems v Meta Platforms Ireland Ltd, Case C-446/21, ECLI:EU:C:2024:834 (2024). <https://curia.europa.eu/juris/liste.jsf?num=C-446/21>