

How To Think About End-To-End Encryption and AI: Training, Processing, Disclosure, and Consent

Mallory Knodel¹, Andrés Fábrega², Daniella Ferrari¹, Jacob Leiken¹, Betty Li Hou¹,
Derek Yen¹, Sam de Alfaro¹, Kyunghyun Cho¹, and Sunoo Park¹

¹New York University

²Cornell University

December 2024

Abstract

End-to-end encryption (E2EE) has become the gold standard for securing communications, bringing strong confidentiality and privacy guarantees to billions of users worldwide. However, the current push towards widespread integration of artificial intelligence (AI) models, including in E2EE systems, raises some serious security concerns.

This work performs a critical examination of the (in)compatibility of AI models and E2EE applications. We explore this on two fronts: (1) the integration of AI “assistants” within E2EE applications, and (2) the use of E2EE data for training AI models. We analyze the potential *security implications* of each, and identify conflicts with the security guarantees of E2EE. Then, we analyze *legal implications* of integrating AI models in E2EE applications, given how AI integration can undermine the confidentiality that E2EE promises. Finally, we offer a list of detailed *recommendations* based on our technical and legal analyses, including: technical design choices that must be prioritized to uphold E2EE security; how service providers must accurately represent E2EE security; and best practices for the default behavior of AI features and for requesting user consent. We hope this paper catalyzes an informed conversation on the tensions that arise between the brisk deployment of AI and the security offered by E2EE, and guides the responsible development of new AI features.

Contents

1	Introduction	2
2	Background I: Encryption, AI, and Trusted Hardware	4
3	Our Definitions and Taxonomy	16
4	Technical Implications of Integrating AI with E2EE	19
5	Background II: Consent, Privacy, Users, and Areas of Law	23
6	Legal Implications of Integrating AI with E2EE	26
7	Recommendations	32
8	Discussion	36
9	Conclusion	38
A	E2EE Features Evaluation Framework	47

1 Introduction

Widespread adoption of end-to-end encryption (E2EE)¹ has improved the confidentiality and integrity of data in various contexts, including messaging, video, and voice communication. Leading applications—such as iMessage [8], WhatsApp [142], and Signal [121]—have integrated E2EE into their systems by default, bringing strong confidentiality and privacy to billions of users’ communications. This trend appears poised to expand further, with an increasing number of services committing to the adoption of E2EE in order to ensure “safer, more secure and private service” to users [60, 52].

In an era where so many forms of communication have moved online, widespread E2EE has become a critical underpinning for the security of communications, providing critical protections for journalists, dissidents, and activists, as well as the everyday family and social communications of billions. As underscored by independent experts from the UN Office of the High Commissioner on Human Rights, strong encryption is essential to fundamental human liberties including free expression, the right to hold opinions, the right to privacy, and the right to assembly [86]. Beyond individual privacy, the ubiquitous nature of E2EE offers *systemic* protections for particular at-risk users, preventing their activities from appearing exceptional: widespread deployment of E2EE has made privacy the norm. This normalization of secure communications, in turn, makes it harder to distinguish whether individuals are engaged in frivolous or sensitive activities, and deters the targeting of individuals merely for using privacy technologies, resulting in a critical “network effect” [73] that creates the conditions necessary for journalists, human rights defenders, activists, and others to communicate securely with less fear of surveillance and repercussions.

The movement towards widespread E2EE communication stayed strong alongside the “big data” trends of recent decades: a recognition of the value of keeping private messaging data out of increasingly vast repositories of personal information, even as data was called the new oil by some [135]. Privacy technologies *gained* momentum as prominent examples of government and corporate intrusions on individual privacy came to light [72, 25, 70], and people became increasingly aware of their digital activity being tracked and profiled at large scale [17]. Those providing E2EE services were popularly seen as protectors of consumer privacy and freedom of speech [26]. Both corporate and government surveillance were key issues, but corporations often stepped into the “protector” role and deployed E2EE technologies that would permit nobody other than the users participating in a communication to access private communication content: a decision that made business sense at the time (e.g., [25]).

Following remarkable recent advances and an explosion of interest in large language models and generative artificial intelligence (AI) more broadly, however, we observe three trends that raise alarms for E2EE security.

First, the way people interact with AI models is changing. While they originally served as standalone applications, AI models are now increasingly incorporated into other everyday applications, including messaging applications, in the form of AI “assistants.” Interacting with these assistants is often baked into the user experience by default, made readily available as part of the application client (e.g., within a messaging app). Such integration creates systemic data flows at scale between previously separate systems, and accordingly raises security and privacy considerations not limited to the realm of E2EE.

Second, high-quality training data is becoming scarce. This has created a race between model developers, with tech companies under increasing pressure to tap any potential source of human-written content [139]. That which is publicly accessible online has already been harnessed, so privately held data is naturally the next resort, and indeed, companies have been quietly changing their terms of service to enable training AI models on more and more of the user data they hold [134]—a practice about which the U.S. Federal Trade Commission has already raised concerns [5]. At the time of writing, E2EE messaging data on major existing platforms remains off limits for AI training, but companies may be incentivized to reconsider these policies in the near future. Apple has very recently launched AI assistant features that can process E2EE messaging content (see Section 2.5), and other companies may be considering similar options.

Third, recent trends in AI appear to have shifted business incentives. Over the last decade, E2EE became a central part of the business model of online intermediaries. Business incentives to protect proprietary and user data were aligned with human rights protections. For example, in 2016, Apple and Facebook were lauded for their public resistance to government requests for access to E2EE data, including their refusal to comply with government requests to build systemic weaknesses into their encryption systems which would allow law enforcement access to messaging data [25]. However, in the last couple of years we have seen a shift towards

¹We use “E2EE” to mean either “end-to-end encryption” or “end-to-end encrypted,” depending on context.

prioritizing AI features as a necessary component of modern applications, and a corresponding search for sources of data to power these, sometimes even at the cost of existing product features and profits [71, 93, 63]. At the same time, interestingly, research and reporting shows that consumers themselves can be much more skeptical of AI-powered products [87, 38].

Against the backdrop of these three concerning trends, the motivating question of our paper is as follows:

Is processing of E2EE content by integrated AI models compatible with end-to-end encryption?
(If so, to what extent and under what circumstances?)

1.1 Summary of contributions

We systematically analyze the above question from both a technological and a legal perspective, combining the expertise of an multidisciplinary research team across cryptography and security, AI and LLMs, and technology law.

We identify the key confidentiality and integrity properties provided by E2EE, for both *individual* and *systemic* privacy considerations, which might be adversely impacted by feeding E2EE content into AI models. This requires a novel conceptualization of key properties of E2EE, ranging from the strictly definitional, to realities of E2EE deployments in practice and their divergences from the strict definitions, to essential properties that only hold when E2EE is used at scale, to many important properties that are widely associated with E2EE but not technically within the scope of formal definitions.

We then examine a wide range of detailed technical configurations that could fall under the broad umbrella of “feeding E2EE content to AI models,” taking into consideration the state of the art in cryptography and privacy technologies, as well as the latest technical developments in AI. We distinguish the use of E2EE content for *inference*, for *training*, or for both. We assess the capacity of each technical configuration—that is, of each approach to integrating AI features into an initially E2EE system—to uphold the key guarantees of E2EE. We conclude that some configurations *cannot* uphold these guarantees, while some others can.

Next, we overview relevant areas of law, and provide a detailed analysis of the circumstances under which E2EE service providers are likely to be able to offer AI features which use E2EE content, consistent with their legal obligations under current US and EU law alongside a brief note on other jurisdictions. We highlight areas of legal uncertainty, and provide a detailed exposition of pending legal (or quasi-legal) processes where relevant. In addition, looking beyond the limits of existing law, we discuss relevant critiques of current law and diverse academic theories of consent and privacy (legal and otherwise).

Putting together all of the above, our analysis yields the following key recommendations summarized here and elaborated in Section 7. Key terms are in italics; we provide detailed definitions in Section 3.

1. **Training.** Using end-to-end encrypted content to train *shared AI models* is not compatible with E2EE.
2. **Processing.** Processing *E2EE content* for AI features (such as inference or training) may be compatible with end-to-end encryption only if the following recommendations are upheld:
 - (a) Prioritize *endpoint-local* processing whenever possible.
 - (b) If processing E2EE content for *non-endpoint-local* models,
 - (i) No *third party* can see or use any *E2EE content*² without breaking encryption,³ and
 - (ii) A user’s *E2EE content* is exclusively used to fulfill that user’s requests.
3. **Disclosure.** Messaging providers should not make unqualified representations that they provide E2EE if the default for any conversation is that *E2EE content* is used (e.g., for AI inference or training) by any *third party*.
4. **Opt-in consent.** AI assistant features, if offered in E2EE systems, should generally be off by default and only activated via opt-in consent. Obtaining meaningful consent is complex, and requires careful consideration including but not limited to: scope and granularity of opt-in/out, ease and clarity of opt-in/out, group consent, and management of consent over time.

²Or any derivative of E2EE content (i.e., any non-trivial function of content). See Section 3.1 for full definitions.

³This may sound paradoxical; see § 4.2 for detailed discussion of technical configurations that may achieve this condition.

Our analysis indicates that there would be serious risks—security risks and legal risks—involved in diverging from these recommendations and thus undermining the established operation of the current E2EE ecosystem according to privacy-by-default practices. Perhaps the greatest risk is in eroding E2EE as a systemic guarantee—an assumed condition of any application or service that is held to be E2EE—and turning it instead into merely a bonus feature that is routinely compromised in exchange for convenience.

Currently E2EE offers essential protections for billions of users communicating confidentially on WhatsApp, iMessage, Signal, and other platforms daily. Moreover, the widespread adoption of E2EE today offers critical *systemic* protections (“network effects”) for users because of its everyday use for more and less sensitive communications alike—protections that a special-purpose platform just for top-secret communication would not be able to provide. These “network effects” currently protect privacy and freedom of expression for at-risk groups and everyday users around the world.

E2EE is a context in which confidentiality is the core guarantee to be held above all others—confidentiality is the essence of the product. Though E2EE may be relatively unusual in this respect, we believe our inquiry in this paper resonates with themes in the broader debate about privacy and trust in AI.

Roadmap. In Section 2 (Background), we overview the properties of E2EE (§ 2.1), how AI assistants work (§ 2.2), how AI assistants can process E2EE data (§ 2.3), and some background on trusted hardware (§ 2.4). In Section 2.5 (Apple Case Study), we provide a case study of Apple’s recent deployment of Apple Intelligence, which can process data from E2EE environments: the only publicly announced scheme of this type, to date. Section 3 provides definitions of key terms that we rely on throughout the paper (§ 3.1), including a taxonomy from “strict E2EE” to “not E2EE” (§ 3.2).

In Section 4 (Technical Implications), we analyze and present a range of technical configurations in which AI assistants might process E2EE data, and associated security/confidentiality implications. In Section 6 (Legal Implications), we overview different areas of US and EU law and how they may constrain the use of E2EE data as input for AI models. In Section 7, we offer our detailed recommendations on how and when messaging data can be fed into AI models consistent with E2EE, taking into careful consideration both technical possibilities and limitations, and existing legal frameworks. Section 8 provides discussion and Section 9 concludes (and includes a detailed note on author contributions).

2 Background I: Encryption, AI, and Trusted Hardware

This section covers technical background, focusing on end-to-end encryption (Section 2.1), AI assistants (Section 2.2), how it is possible for AI assistants to process E2EE content (Section 2.3), and trusted hardware (Section 2.4). Lastly, Section 2.5 gives a brief overview of Apple’s recently deployed system in which E2EE content can be processed by AI assistants.

2.1 Background on End-to-End Encryption (E2EE)

Online messaging systems that allow communication between users (such as iMessage, WhatsApp, and Signal) are generally intermediaries to every communication. Although the user experience may make it seem like messages go directly from a sender S to a recipient R , in reality every message goes from S to the platform P and then from P to R . This means that, absent any additional guarantees, the platform has access to all communications of users on the platform.

End-to-end encryption (E2EE) is a standard of security for communication in which only the sender and an intended recipient can read communications between them. In particular, even the platform P cannot read these communications, despite P still acting as the intermediary for communication: that is, communications still go from S to P then P to R .

The key tool used in E2EE is *encryption*. Encryption is a method for a sender S to “scramble” their communications in a manner which is easy for their intended recipient R (and S) to unscramble, but practically impossible for anyone else. Using encryption, S will scramble (or *encrypt*) each message for R before sending it to platform P , which only handles messages in scrambled form (and does not have the ability to unscramble messages). When R receives a scrambled message from S via P , they are able to unscramble the message (or *decrypt* it) and read the original message that S wrote.




Figure 1: A sender S and receiver R communicate using an end-to-end encrypted application hosted by a company (middle). Solid lines represent plaintexts, and dashed lines represent ciphertexts. S and R can read their messages on their devices; however, while a message is “in transit” between their devices, it is encrypted so that it is not readable to the intermediary platform P handling it on its servers (or indeed to anyone but the intended recipient, whether a network eavesdropper, an employee at P , a hacker who compromises P , or a wrong recipient to whom the message was delivered by mistake).

The additional security that E2EE provides is an important protection for the privacy of the communicating users, as well as a key protection against misuse of the message contents by or at the platform P : e.g., the messaging data is safely inaccessible if P suffers a data breach, if P is asked to share data with other organizations (e.g., business partners or the government), or if an insider at P wants to read the data.

We conclude this section with some terminology and follow it with a slightly more detailed and technically accurate explanation of how E2EE works.⁴

Plaintext and Ciphertext. Message content that is unencrypted is described as *plaintext*. Normal, readable messages that you type into iMessage, Signal, or WhatsApp are plaintexts. Despite the name, plaintexts can be any type of content: audio files, images, spreadsheets, and any other digital information. Encryption provide a technique to convert (or scramble) plaintexts into *ciphertexts*, which are unreadable except to the designated recipient (and the sender who wrote it).⁵

Secret Key. Each user in an E2EE system has a *secret key*, a secret piece of information that should be known only to them (much like a password). A user’s secret key is necessary to decrypt messages intended for them. Secret keys are created and stored on users’ devices: platforms (and other users) do not have access to them.

Encryption and Decryption. The process of converting a plaintext into a ciphertext is called *encryption*, and the process of converting a ciphertext back into a plaintext is called *decryption*. Only the intended recipient R should be able to perform decryption to recover plaintexts from ciphertexts; accordingly, decryption requires the use of some secret information which only R knows, namely, R ’s *secret key*.

Putting it together. Equipped with more formal vocabulary, we restate the basic mechanism of E2EE. In E2EE messaging, each sending user S ’s device encrypts each plaintext message, then sends the resulting ciphertext to the platform P . The platform P passes the ciphertext on to the recipient user R ’s device, which uses its secret key to decrypt the ciphertext and get back the original plaintext message that S wrote.

⁴A complete exposition of how E2EE works is beyond the scope of this paper.

⁵The terms *plaintext* and *ciphertext* can each be used either as a noun or as an adjective.

2.1.1 The E2EE Confidentiality Guarantee

The core privacy guarantee of an E2EE system is called *confidentiality*.⁶ Informally, E2EE confidentiality states that plaintext messages handled by an E2EE system can only be seen by the sender and intended recipients: no one else, including law enforcement or the service provider, can learn anything about the plaintext messages [89]. Importantly, this guarantee is stronger than saying that nobody else can *learn the plaintext messages*: E2EE confidentiality requires that nothing can be learned about the content of an E2EE message, including the message itself, but also any derivative of the message content.⁷ Thus, E2EE confidentiality means that anyone besides the sender and intended recipient must not be able to learn, for example, what letter of the alphabet a message begins with, or whether the message contains a given word.

E2EE confidentiality refers to a specific technical definition of confidentiality for messages, based on encryption—E2EE confidentiality (or E2EE security) is a term of art with a precise meaning, and does not ensure perfect “confidentiality,” “privacy,” or “security” for every colloquial meaning of the terms. In general use, these terms are very broad and mean different things in different contexts.⁸

Consider the context of messaging. The E2EE confidentiality guarantee is both stricter than one might expect, in some senses, and less strict in other respects. If strictly enforced, it is incompatible with many features common in messaging apps.

Strictness of the E2EE confidentiality guarantee. A strict reading of the E2EE confidentiality guarantee would be incompatible common features such as *link previewing*, where an app displays a “preview” of a link that is sent by E2EE messaging. In this case, the link itself is E2EE message content, and in order to fetch and display the preview, a third party must process that link. Some E2EE applications only display link previews upon explicit user opt-in, but some have it on by default.

Permissiveness of the E2EE confidentiality guarantee. The E2EE confidentiality guarantee covers only encryption of *message contents*, as noted above. For any other auxiliary information related to messaging, such as who is messaging whom and at what times, E2EE does not impose any requirements or provide any guarantees. To give a more complex example, in 2021, the Government of India introduced legislation mandating that messaging platforms provide message “traceability,” or methods for determining which user first shared particular content [129]. Which user first shared content is not the content itself, and this feature does not violate E2EE confidentiality.

Many academic definitions of E2EE messaging consider a simple model where the only activity that occurs is literal messaging—sending messages between senders and receivers and within group chats. In practice, E2EE messaging apps—like any modern messaging app—are more richly featured (e.g., including options to back up your data and see link previews). How such additional features process message contents, and whether such processing respects E2EE confidentiality, has been a subject of long debate. Some common features of E2EE applications *enhance* confidentiality despite not being part of the basic definition of E2EE; other common features of E2EE applications may *limit* confidentiality (such as link previewing).⁹

Next, in Sections 2.1.2–2.1.3, we present a range of technical and non-technical properties often associated with E2EE applications in practice, but which are distinct from the basic E2EE confidentiality guarantee.

⁶We will hereafter use the term “E2EE confidentiality” in this paper to stress that we are referring to this particular technical property, and not the common use of the word.

⁷The formal mathematical definition of security of encryption states that no adversary can any function of the message content, except for the message length in bits.

⁸For example, E2EE technologies do not prevent your friend from choosing to disclose the sensitive information you shared with them to others, even though this might seem to “violate confidentiality” or “violate your privacy” in a broad sense. Indeed, no technology can prevent this. See also Section 5 for more discussion.

⁹Other additional features include, for example: user reporting, user blocking, importing contacts, stories, message backups, message forwarding, blurring explicit images, autocorrect, accessibility features (such as text-to-voice), and hands-free integration or carplay. It is an instructive exercise to consider how each of these do or do not have implications for confidentiality and E2EE confidentiality (depending on how they are implemented). While outside the scope of our work, there is a delicate balance (which has been the source of much debate [[empty citation](#)]) to be struck between the deployment of usability features that increase adoption of E2EE systems (and, thus, help enhance the systemic guarantees discussed in Section 2.1.2) on the one hand and strict E2EE confidentiality on the other hand.

2.1.2 Systemic Properties of E2EE Systems

Even if an E2EE application provides strong technical privacy guarantees (including and beyond E2EE confidentiality), the strength of the privacy guarantee it offers will also depend on its usage in practice.

Messaging applications that are widely used for diverse purposes offer additional protection because their use is unremarkable, and hence, the use of the application in itself does not draw attention. In economic terms, this is a *network effect* where the (privacy-related) utility of the application for any given user depends on other users. For example, suppose that a particular E2EE messaging app A happens to have been developed by a developer who is an outspoken member of a particular political community and has primarily advertised A in the social circles and publications of that community. Accordingly, A grows such that the userbase disproportionately represents members of this political community. Even if the application has no flaws whatsoever in providing E2EE confidentiality (and other properties), knowing that a particular person uses application A can divulge their political views, if only by association.

This type of example motivates the intuition that to realize the full privacy potential that E2EE can provide, the use of the application should be commonly used and unremarkable in the relevant context. This “unremarkableness” is not a technical property: it is instead an emergent social and contextual property of the system’s use as a whole considered over the entire userbase (relevant to a context of use¹⁰). Existing security and privacy literature has discussed similar network effects in privacy technologies not limited to encrypted messaging, and noted their critical importance to engineering effective privacy technologies [59].

In summary, there is a key privacy property of E2EE systems which cannot be achieved by technical means alone, but *is* attained by popular E2EE applications such as iMessage and WhatsApp today—namely, the network effect gained from wide usage, so that the use of the system in itself is unremarkable.

2.1.3 Auxiliary Security Properties of E2EE Systems

Besides the core guarantees mentioned so far in this section, there are additional (cryptographic) security properties that E2EE applications often provide in practice. While distinct from basic E2EE confidentiality, these properties have become very commonly associated with, or even expected of, modern E2EE applications. Notable examples include:

- *Forward secrecy*: An adversary who has compromised the communication channel (e.g., by stealing the victims’ secret session key) cannot recover messages that were transmitted before the session was compromised.
- *Post-compromise security*: An adversary that has compromised the communicate channel cannot recover messages that are transmitted after the protocol “heals” itself (e.g., after session keys are rotated).
- *Deniability*: No party can prove that a message was authored by any particular user of the system, even if the prover was one of the intended recipients of the message.
- *Unlinkability*: Knowledge of the fact that a user authored a particular message reveals no information about whether they authored other messages.

State-of-the-art E2EE messaging protocols, such as the Signal protocol [121], meet these (and many additional) security properties.

2.1.4 Necessary Conditions for E2EE Confidentiality

End-to-end encryption provides its promised confidentiality guarantee only under certain conditions: namely, when (1) the underlying encryption method is secure and (2) that encryption method is securely implemented in the application being used.

Cryptographic assumptions. Every practical encryption scheme relies upon a *hardness assumption*: an assumption that a particular mathematical problem is difficult for computers to solve. The hardness of the mathematical problems underlying today’s widely used encryption schemes has held strong in the face of heavy scrutiny over decades by mathematicians, computer scientists, and others, leading to a strong community belief that it is acceptable to use cryptography that relies on these assumptions in security-critical contexts. When new encryption schemes that rely on new, less-tested assumptions are proposed, they are

¹⁰E.g, widespread use in country X may not be helpful for a user in country Y .

not considered suitable for deployment until they have been vetted in a multi-year process of scrutiny by mathematicians and computer scientists worldwide.

Practical assumptions. Even if the math underlying an encryption scheme is perfect in theory, the process of deploying the encryption scheme in software or hardware introduces avenues for security failures, and the use of the encryption scheme by fallible humans in practice introduces further potential for security failures.

In order for E2EE confidentiality to hold: (1) the code implementing the E2EE application must be correct, (2) the hardware device that the application runs on must be secure, and (3) the users using the application must use it correctly and take adequate security precautions.¹¹ Undermining any one of these three could undermine E2EE confidentiality: for example, (1) the software might fail to perform the encryption correctly; (2) the hardware device might fail to generate a secret key correctly, or it might be infected by malware; and (3) a user might fail to keep their secret key secret. Errors can be introduced in myriad ways: through simple mischance, human mistake, or intentional sabotage.

Typically, in current practice, a messaging application is implemented by an intermediary platform that provides the associated messaging service. Thus, the party with the most control over and visibility into the implementation process is the platform. Users therefore rely on platforms for (1), i.e., that the code implements E2EE correctly. Users rely on device manufacturers (and their suppliers) for (2), i.e., that the device running the app is secure. Finally, the users themselves determine (3), i.e., their usage of the application; that said, users often lack the expertise to confidently take adequate security precautions.

2.2 Background on AI Assistants

AI assistants are programs designed to interpret everyday language and perform computational tasks. Today’s AI assistants, such as ChatGPT, Gemini, and Llama, are able to handle a wide range of tasks, including text analysis, content creation, code generation, language translation, and more. At the core of these technologies are statistical models—programs trained on data to identify patterns. Specifically, AI assistants typically rely on large language models (LLMs), which are trained on vast datasets of text to generate outputs based on predictions of the next word. Due to the size of their training datasets and sophisticated architectures, LLMs can handle complex and previously unseen user inputs (*queries*) and provide contextually relevant responses.

2.2.1 Training and Inference

AI assistants undergo two key phases: *training*, when the model is built to identify patterns, relationships, and rules; and *inference*, when real-time responses are generated based on user inputs. Training itself often involves multiple stages, typically referred to as *pretraining* and *fine-tuning*. In pretraining, the model takes in vast, general datasets to learn structures and patterns in the data. Fine-tuning then refines this pretrained model by training it on additional, more specialized datasets, improving its accuracy and relevance in particular contexts. We describe these processes in further detail below.

Training. Models have configurable *parameters* which are numerical values that determine how outputs are generated. These parameters are used in calculations to predict outputs when the model is queried with an input. During training, the model processes large datasets—such as books, websites, code repositories, and conversational datasets in the case of LLMs—to adjust (or “tune”) its parameters. This tuning happens over many iterations of training, during which the model’s parameters capture more and more refined patterns in the data. The finalized set of parameters constitutes the trained model, which can then be *prompted* with inputs (e.g., questions) to generate outputs (e.g., answers). Models at the scale of commercial AI assistants contain billions or even trillions of parameters trained over the course of months on high-performance computers.

Inference. Once trained, models can perform inference, during which they generate real-time responses to user inputs. When a user inputs a query to the model, it is processed through many steps using the model parameters to calculate an output. While inference is less resource-intensive than training, it still requires significant computational power and is typically handled by centralized servers.

¹¹In fact, these are essential considerations for the proper functioning of any software product, not limited to encryption.

Deployed models undergo continuous improvement through updates that may involve additional training on more diverse or recent datasets, reinforcement learning with human feedback, where collected user feedback is used to improve response quality; patches to address issues such as specific problematic outputs, and model refinement with new techniques and architectures. As such, inference and training often overlap in a *feedback loop*, where data collected during inference can be fed back into the training pipeline to fine-tune the model further. This process is shown in Figure 2.




Figure 2: Inference and training continuously happen with an AI assistant as a feedback loop. As users continue to query the AI assistant, this generates data that is stored and can be used to then train and improve the model.

2.2.2 Data Collection

Once the AI assistant is deployed, various types of data may be collected. This can include *chat logs*, which record user inputs (text, voice, or image data) and the system’s responses. These historical records, even if kept by the platform, may or may not be visible to users.¹² Additionally, *metadata*, such as timestamps, device type, language, and geographic region, as well as user preferences and settings, may also be collected and stored by the platform. The collection of this data serves multiple purposes, which can be broadly categorized as follows:

1. **Ease of use:** The data is used to enable certain features, such as maintaining conversation histories, or maintaining continuity within a session or across sessions, for an individual user.
2. **Personalization:** The data is used to identify usage patterns, such as query frequency and commonly accessed features, to tailor the system’s functionality for an individual user.
3. **Analytics:** The data and usage patterns is aggregated across users to steer further improvements and developments of the model.
4. **Model refinement:** The data is used as training data to fine-tune the model and continuously improve its performance.

These categories are not mutually exclusive, as data may be collected to serve any or all of these purposes. The collection of user data is commonly enabled by default with collected user data transmitted and stored on centralized cloud servers. In some cases, some or all user data may remain local on the user’s device, eliminating external transmission.

¹²Search/query history can be surprisingly revealing when aggregated over time, involving family concerns, health issues, religious beliefs, and more. Research has shown users can be reluctant to disclose all their search history even for better personalized search results [117].

2.2.3 Architecture of AI Assistants in Messaging Platforms

Next, we overview the architecture of AI assistants integrated in messaging platforms. The architecture is how the tool is structured and has implications on how and where data is processed. First, it is important to distinguish between the messaging application, the AI assistant, and the model. The messaging application is the application installed on the user’s device which facilitates communication between users and provides the interface for interacting with the AI assistant, e.g., WhatsApp, Messenger, and Signal. The AI assistant is the system that performs tasks such as answering queries, generating responses, or providing recommendations. At its core is the model—in this case, an LLM—the underlying computational engine that processes inputs and generates outputs based on patterns learned during training.

When the messaging application gets an input from the user, this input is sent to the model in the form of a *request*. Once the request is processed and an output has been generated by the model, the output is sent back to the platform as a *response*. As such, while the messaging application itself will be installed on the user’s device, the model need not also be directly on the device; network connections can facilitate this communication between the application and model.

Here we differentiate between two types of architectures: *cloud-based* and *local*, illustrated in Figure 3.

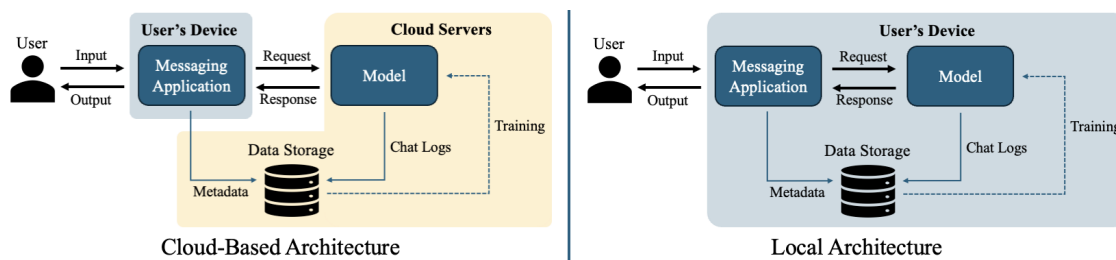


Figure 3: Generating an output from a user’s input, differentiating between what is on-device and on cloud servers between a cloud-based and local architecture.

Cloud-Based. Cloud servers are designed to handle large computation and storage demands. Given the demands of an LLM—recall that there may be trillions of parameters to be processed simultaneously, requiring very high levels of computation power as well as memory to store the parameters and data—typically, models are hosted on cloud servers. Cloud servers can carry thousands of high-performance computers (Graphics Processing Units or Tensor Processing Units), exceeding the capabilities of a typical laptop by many magnitudes.

Local. In limited scenarios, computations can occur *locally*—that is, directly on the user’s device using lightweight versions of the model or specific components. Local processing ensures that data does not leave the device, but is only possible with smaller, less powerful models due to reduced processing power and storage capacity of local hardware.

2.2.4 Availability in E2EE Applications

AI assistants are already available in major E2EE applications as a standalone tools outside of conversations between users. For example, WhatsApp features a search bar that can be used both to search through conversations and to “Ask Meta AI”. When the user inputs a search query, the results show matches from their conversation history (via a local search) as well as an option to invoke Meta AI with the search query. In this case, we do not describe Meta AI as being integrated with E2EE content, as only explicit Meta AI queries initiated by the user are processed by Meta AI, without any E2EE content. The availability of such non-integrated features in E2EE applications is outside the scope of our concern in this paper, which is focused on AI models processing E2EE content.

In contrast, integrating AI assistants as a feature within E2EE conversations would typically require the AI assistant to process E2EE content. For example, this could be in the form of allowing users to invoke the AI assistant within an ongoing conversation to provide contextually relevant support or responses.

2.3 How Is It Technically Possible for E2EE Data to Be Fed Into AI Assistants?

The definition of E2EE may make it sound like it should be impossible to feed E2EE content into a (cloud-based) AI assistant. How can an AI model have access to plaintext data, when E2EE guarantees that only the sender and recipient can access E2EE content in plaintext?

Certainly, the cloud-based service provider has access to the encrypted message—the ciphertext—and could input that ciphertext to an AI assistant. In this case, however, the AI assistant would not be able to read any meaningful information from the ciphertext, and thus would not be able to produce a useful answer, either.¹³

The current discourse around using E2EE data in AI assistants, therefore, is focused primarily on the idea of feeding readable *plaintext* versions of E2EE messaging data into AI assistants. Since only the sender and recipient devices in an E2EE system can read the plaintext messages, the sender or recipient devices would be the only ones who could provide the plaintext data to the AI assistants. This possibility appears to be primary focus of the ongoing discourse: that users’ devices would be set up to send E2EE content to cloud-based AI assistants for processing.

Nothing about E2EE technically prevents this from happening—indeed, no technical measure would both be able to allow senders and recipients to read their own messages, and prevent them from showing them to anyone else. This is much like no matter how secretly you encode and deliver your hand-written letter personally to your friend, nothing you can do can stop your friend from showing the decoded letter to other people once they have it.

But then, what is the point of end-to-end encryption?. We have just explained that E2EE is a technology that allows transmission of messages in encrypted form so that nobody but the sender and intended recipient can read them, not even intermediary platforms. We have also explained that E2EE does not stop the sender and recipient devices from sending plaintext copies of the messages to anybody — and that devices can do this in a way that is very hard to detect for average users.

Then, *is there any point to using E2EE?* After all, our devices could be routinely sending plaintext copies of all our messages *to the messaging platforms themselves*, and E2EE would not stop this — yet such a practice would undermine precisely the benefit that E2EE purports to provide.

E2EE *does* still have a point, for a number of reasons. First, it provides important protections when service providers *choose* not to be able to read your messages, and implement E2EE in order to achieve this goal. Such a choice might be motivated by various reasons: it might provide a competitive edge, provide protection in the event of data breaches, or reduce compliance costs for companies. Once such a choice is made, the platform moreover has a reputational interest in making sure E2EE is achieving the stated goal that the platform not be able to read messages. Second, E2EE provides stronger protections in the case of popular messaging platforms that are widely used and therefore widely scrutinized by technical experts, including security and privacy experts. While subversions like sending plaintext copies of messages to the platform would be hard to detect for an average user, they would be much more easily detected under such scrutiny by experts — and from the experts’ scrutiny, average users benefit too. Third, E2EE provides even stronger protections in the case of popular messaging platforms that choose to openly disclose the details of how their system and security features work, for similar reasons. Balsa et al. [21] provide more detailed discussion of these and related issues around platforms implementing E2EE and other cryptography.

2.4 Background on Trusted Execution Environments (TEEs)

Trusted hardware refers to physical devices that are designed to ensure the confidentiality and integrity of computations, even if the owner of the hardware may be compromised. Trusted hardware is designed to create a secure area of computation, in which no information about the computation is released outside of the trusted hardware during execution of code, and that code will run exactly as specified. There are many types of technologies that fall under the broad umbrella of trusted hardware—such as hardware security modules (HSMs), trusted platform modules (TPMs), and trusted execution environments (TEEs)—which offer different capabilities. Of these technologies, TEEs are the most relevant for the implementation of

¹³This is mathematically guaranteed by the formal definition of secure encryption.

AI functionalities, as they are designed to run arbitrary programs; other technologies currently have more limited, bespoke use cases, such as key generation and storage. We thus focus on TEEs in this subsection, and in the rest of the paper.

We can conceptualize a TEE as a physical “lockbox”, which is designed to execute arbitrary code such that no outside entities (including, for example, the operating system, other device processes, and third parties) can observe nor manipulate the state of the code as it is running. During initial setup of the TEE, a computer program—call it Π —is loaded inside the TEE. The TEE then has a single input/output interface, through it which receives inputs x from outside the TEE, computes the program Π on the inputs, and outputs the result $\Pi(x)$ from the TEE. During execution of Π , the TEE maintains some internal ephemeral state in memory, which stores temporary values during computation. TEEs are also often equipped with persistent storage, on which they can store data longer-term. See [120] for a formal treatment of TEE abstractions.

TEE security goals. TEEs are intended to provide two core security guarantees. The first is *confidentiality*, which means that the internal state of the TEE should not be visible to any parties during execution of Π ; only x and $\Pi(x)$ are revealed to third parties. Similarly, the persistent storage of the TEE should always be kept secret. The second goal of TEEs is *integrity*, which means that Π should be the exact program run by the TEE, which cannot be undetectably tampered with nor replaced by another program. Besides these two core guarantees, some TEE implementations have a third security goal, *remote attestation*, which means that the TEE should be able provide convincing evidence that a given program Π is the code that is loaded inside the TEE. Critically, these guarantees should hold even in the presence of a compromised TEE owner (e.g., a cloud provider): using a number of hardware-level protections, TEEs are designed to be tamper-resistant (for example, TEE chips often have physical seals that are designed to detect attempts to open them, upon which the TEE’s storage is immediately erased).

The goal of TEE design is to enable outsourcing of computations over user data to potentially untrusted (but more computationally powerful) environments, simply by loading the program inside a server-side TEE, and processing user data inside of it. As described so far, the inputs to and outputs from the TEE can be seen by the owner of the TEE. However, inputs and outputs can be hidden from the TEE owner with an additional step: encrypt inputs and outputs so that they can only be decrypted using secret key that is stored inside the TEE. Then, computations inside the TEE can still happen in plaintext, within the “shield” of the hardware protections provided by the TEE.

TEE performance. Modern TEEs (such as Intel SGX [51] and AMD SEV [85]) are capable of supporting complex computations with practical overheads. For example, TEE-based security has been deployed in Signal’s contact discovery mechanism [98], and even entire blockchains [110]. However, efficiency is not the only engineering consideration for TEEs: their design of strict isolation imposes important limitations on the programs that TEEs can run, since they cannot have non-trivial interfaces to the outside world. So, for example, TEEs generally cannot make Internet requests, and cannot rely on external sources of state.

2.4.1 Necessary conditions for TEE security

The security guarantees of TEEs naturally rely on the correct implementation of the TEE, and that the hardware has not been tampered with at any point in the TEE’s life cycle (e.g., tamper-proof seals have not been removed, backdoors have not been inserted, etc). This requires trust in various actors across the hardware supply chain, from manufacturing, to assembly, to deployment. Further, users need assurance that the application provider is actually running a legitimate TEE on their servers. While this can be implemented via the TEE’s attestation capabilities, the attestation service needs to leverage some trusted authority, such as the hardware provider.

Even if the hardware is correctly implemented according to specification, the security of TEEs additionally relies on the specification itself being free of vulnerabilities. Unfortunately, there is a history of practical attacks on various TEE implementations [138, 67]. These attacks range across a variety of threat models—including adversaries that have physical access to the TEE, can launch co-located TEEs with arbitrary programs, control the system’s boot process, etc—and the adversary’s goal is to violate the confidentiality and integrity of the TEE, i.e., recover sensitive runtime state or tamper with the data in the TEE. Notable examples of attack vectors include physical and micro-architectural side-channels, which exploit shared caches [34], branch predictors [90], page faults [144], transient execution [136], and more. This history of attacks has prevented TEEs from reaching mainstream adoption for security applications.

2.4.2 TEEs vs E2EE

While TEEs and encryption have similar security goals, there are some critical differences between the two technologies, which we review in this section. We first explain why they are not interchangeable security solutions, and then why TEEs introduce additional security concerns not present in the encryption context.

TEEs and encryption provide different security guarantees. E2EE security and TEE security each rely on a number of conditions holding true, as summarized in Section 2.4.1 just above (for TEEs) and in Section 2.1.4 (for E2EE). The practical guarantees provided by each technology depend on both the precise security goals they are designed to achieve, and the conditions (or the *threat model*) under which they are designed to achieve them. While TEEs and E2EE have overlapping security goals, the conditions necessary for E2EE security and for TEE security are *orthogonal*—qualitatively different and incomparable.

The security of E2EE and TEEs rely on two broad assumption types: (1) specification correctness, and (2) implementation correctness. The former states that the design of the technology/protocol is secure, while the latter states that a particular instance of the technology/protocol is correctly implemented according to the specification. TEEs and E2EE have important differences for both assumption types. Assumption type (1) generally hinges on the hardness of a mathematical problem for E2EE; and on a robust hardware design for TEEs. On the other hand, assumption type (2) hinges on the security of software and hardware implementations for E2EE and TEEs, respectively. These entail a completely different set of considerations between E2EE and TEEs, since the software and hardware supply chains involve different entities, vendors, and processes.

Since the security of TEEs and E2EE depend on orthogonal considerations, their resulting security guarantees are *incomparable* in a technical sense. That is, the precise type of security you get, and the conditions under which you get it, are different and not directly comparable. TEEs rely upon a distinct set of conditions for their security from E2EE, and therefore cannot be substituted into an E2EE protocol without acknowledging that the nature of the security offered has changed, away from E2EE security. Thus, again, while the security goals are similar, they are not interchangeable technologies.

TEEs raise certain security concerns that encryption does not. The difference in the type of security offered by TEEs and by encryption mean that depending on the context, one may be more suitable for use than the other. That said, when considering whether to use TEEs or encryption for a given application that might admit either, the former raise certain kinds of security concerns not present for the latter. Widely used encryption schemes today have largely withstood decades of scrutiny without critical confidentiality breaks; indeed, new encryption schemes that have not yet been subject to such scrutiny are widely considered unsuitable for use in security applications. Our understanding and scrutiny of TEE security is comparatively young, and the field has seen and continues to see a stream of discoveries of security flaws with impacts on deployed systems (e.g., [138, 67]); at the same time, TEEs enable secure computation in promising new applications, which would not be practical based on cryptography alone.

2.5 A Practical Example: Apple Intelligence

Apple introduced a new product called Apple Intelligence in late 2024, which provides new OS-integrated AI capabilities across the Apple ecosystem.¹⁴ Apple Intelligence will be able to process and perform AI tasks on E2EE content [12]. According to its online documentation, Apple’s plan does not involve *training* any models on any “users’ private personal data or user interactions,” which presumably excludes E2EE content from use for training [13].

In this section, we overview Apple’s plan for AI processing of E2EE content. We cover Apple’s plan not in order to spotlight Apple’s particular practices as opposed to those of other companies, but rather, simply because Apple’s plan is the best documented real-world example of AI processing of E2EE content on a major platform at the time of our writing. We use their plan as a convenient illustrative example for our exposition and analysis. Additional reasons, not focused primarily on Apple, that we find it compelling to study the example of Apple include: (1) in light of Apple’s proposal and the competitive AI landscape, it is likely that other major platforms are seriously considering similar programs; (2) in light of the large-scale weakening of encryption that could occur if AI assistant features in E2EE applications were deployed overly hastily,

¹⁴Apple Intelligence’s initial release was in October 2024, the ChatGPT integration was launched on December 11th, and further features (including Siri integration across third party apps) are coming soon [12] [7].

the kind of analysis in this paper is essential to conduct early; and (3) as other platforms consider similar programs, they may refer to Apple’s plan—conceivably both for technical reference and for normative and competitive justification for their own plans—as the earliest detailed documented proposal of the kind. We encourage readers to view the Apple examples throughout the paper as illustrative more than spotlighting, and to keep the broader landscape of other platforms in mind.

We commend Apple for their comparative transparency, which has facilitated our writing with reference to their concrete real-world example, while also acknowledging that despite that, their documentation leaves many critical questions unanswered (as discussed in Section 2.5.2). As we elaborate in later sections, we believe that stronger transparency and disclosure are essential when using AI to process E2EE content.

What data on Apple devices is E2EE? Applications across the iOS ecosystem differ on the confidentiality guarantees they offer. Our work focuses only on data that is protected with E2EE, so in this section we explain what data in Apple devices is E2EE. There are two broad classes of applications for Apple devices: native applications, which are managed by Apple and come built-in with the device; and third-party applications, which are downloaded from the App Store. A few native applications are E2EE by default, such as iMessage, Facetime, and Keychain. However, storage on iCloud is not E2EE by default, so most native applications that rely on cloud storage are not E2EE. For additional security, users have the option of enabling *Advanced Data Protection (ADP)* [10], a feature which enables E2EE for the majority of iCloud data, such as photos, notes, voice memos, etc. Some iCloud data is not E2EE even when ADP is activated, such as mail, contacts, and calendars.¹⁵ In this section, we refer to native applications for which ADP enables E2EE as *ADP-E2EE applications*. In addition, ADP enables E2EE for certain Siri features, such as Siri shortcuts and other information (e.g., settings and personalization). It is not clear to us whether Siri transcripts are E2EE or not, even with ADP enabled.

The confidentiality guarantees of third-party applications vary by application. Some third-party applications are E2EE, such as Signal [121], and WhatsApp [142].

What data can Siri and Apple Intelligence access? Siri and Apple Intelligence are able to process users’ content drawn from any application where the application developer has “integrated” Siri and Apple Intelligence functionalities [11], *and* the user has not deactivated this integration in their settings [14]. That is, if an application supports integration of Siri and Apple Intelligence, then integration/access is on by default unless a user turns it off in their settings.

Possible use cases involving E2EE data. We highlight two key contexts in which E2EE data is used for inference in Apple Intelligence. First, Apple Intelligence may be used to process user content produced inside an E2EE application. For example, long sequences of messages in iMessage chats may be processed by AI to produce a short “summary” (say, when a user has been away from a group chat, to help the user catch up) [12]. Using Apple Intelligence or Siri inside any third-party E2EE or (if users have ADP enabled) any ADP-E2EE application could involve interaction with E2EE data.

The second way in which Apple Intelligence may process E2EE data is through its content generation features, which may be accessible from within E2EE applications—this includes third-party E2EE and ADP-E2EE applications if ADP is enabled. For instance, image and emoji generation can be based on user-uploaded images and accessed through the iMessage app. If ADP is enabled, the generated images and emojis may be based on E2EE content.

2.5.1 Three-Tiered Approach

As users interact with Apple applications on their iPhones, iPads, or Macs, these applications may submit queries to the Apple Intelligence system. These queries, often a function of user data—and sometimes a function of user E2EE data—are then processed and responses are sent back to the user (e.g., producing short summaries of long text, or updating a contact’s information).

Apple employs a three-tiered approach to handling user queries: (I) *local models* which process simple queries entirely on-device; (II) a *server model*, which processes more complex queries in a cloud-based trusted execution environment hosted by Apple; and (III) sending data to a *third-party model*, OpenAI’s ChatGPT,

¹⁵We refer readers to [10] for a complete description of the confidentiality guarantees for various data types and native applications in iOS, both in the default case and when ADP is activated.




Figure 4:

which processes outsourced queries that could benefit from “ChatGPT’s broad world knowledge” [12]. Each time a query is sent to ChatGPT, user consent is requested.¹⁶

These three cases are coordinated by a local *orchestration layer*, which examines every query and determines how to process it: under Case I, II, or III. The orchestration layer prioritizes these cases in the order presented above: “many of the models” used to process queries are on-device [12] [7], only delegating “more sophisticated requests” to the more advanced model [65], and ChatGPT is only summoned through Siri “when helpful” [12]. As we discuss in Section 2.5.2, Apple has not clarified which models process which types of queries, nor the criteria used by the orchestration layer to decide how a given query should be processed¹⁷. We now discuss what we do know about the three cases in more detail.

Case I (Local model). Apple Intelligence’s first processing method consists of a series of small AI models stored entirely on a user’s device [13]. While some of these models are bespoke for certain tasks (e.g., code completion), their core local model consists of an approximately three-billion-parameter general-purpose language model, which can be adaptively fine-tuned for particular language tasks (e.g., summarizing text).

Case II (Server model). Tasks that are “sophisticated” enough to fall in Case II are outsourced to a larger model, which is hosted on Apple’s servers [65]. While the full details of this model have not been made public, we refer to [58] for an in-depth third-party analysis of possible architectures by Trail of Bits. Outsourcing queries to a cloud-based model places—including E2EE data—undermines E2EE confidentiality. Apple offers some discussion of this heightened risk and proposes to use a trusted execution environment (TEE)¹⁸ to perform server-side inference more securely [65].

Apple has built custom trusted execution environments in their data centers—which they call *Private Cloud Compute (PCC)*—to host (and perform inference with) their server-side model. User devices communicate directly with Apple’s PCC and supply plaintext inputs to the model, which are evaluated inside the PCC (i.e., inside Apple’s TEE). Apple employs several additional measures to enhance the security of the PCC computation nodes; some of these measures appear to be designed partly in response to security considerations such as those outlined in Section 2.4.¹⁹

Case III (Third-party model). Finally, Apple Intelligence sends certain user queries to ChatGPT. This feature is turned off by default, but can be manually enabled in settings [16]. When the feature enabled, Siri examines each request and determines if ChatGPT “might be helpful to get the information” to respond to it [15]. The types of requests for which Siri may invoke ChatGPT may include: direct prompts to ChatGPT, text generation, and summarization [15]. If either ChatGPT is directly invoked or Siri deems ChatGPT could be “helpful” [12], the user is asked in a pop-up prompt whether they would like to send the request to OpenAI for processing.²⁰ By default, this prompt occurs every time a request could be sent to ChatGPT,

¹⁶Our understanding is that permission prompts can be disabled with the exception of when files are sent to ChatGPT, in which case the user will always be asked [116].

¹⁷We know that all images generated with Image Playground are created on device [12].

¹⁸See Section 2.4 for an introduction to TEEs.

¹⁹For example: the PCC nodes leverage many of the hardware security technologies used by Apple’s end-user products (such as their Secure Enclave Processor (SEP) and Secure Boot); the nodes are stripped of most functionality (except for security guardrails and inference capabilities) to minimize potential attack vectors; and user data (which is kept exclusively inside the PCC during computation) is deleted immediately once a response is returned to the user [65]. Apple additionally relies on a number of techniques to address supply-chain attacks and other backdoors, such as publishing every PCC production build in append-only, tamper-proof logs [65], and X-ray imaging during assembly for anomaly detection [9].

²⁰For an example of what this looks like, see this example from the Apple Intelligence press release [12].

but can be disabled in settings²¹ [15]. Apple collects limited metadata about the queries sent to ChatGPT, including the number and size of requests [14].

Upon enabling the ChatGPT extension in Siri, users may choose to continue with or without logging in to their OpenAI account.²² If the user doesn't log in, query processing is consistent with Apple Intelligence's privacy policies: OpenAI is required to process queries ephemerally and must not use data from requests to train its models [16].²³ However, without logging into their OpenAI account, the user would not have access to any premium features linked to their OpenAI account [115]. If, on the other hand, the user chooses to log in to their OpenAI account, query processing is governed by ChatGPT's privacy policies and the user's OpenAI account settings [16]—not Apple's policies. As such, for logged-in users, OpenAI may log requests and use query data to train its models if the user's OpenAI account permits such processing (even if this is inconsistent with Apple's policies) [16].²⁴

2.5.2 What We Still Don't Know

Based on available documentation at the time of writing, significant details remain unclear about Apple's system. Open questions, to which the answers could be essential to independently evaluate the system's security risks and compatibility with E2EE, include the following.

- Will users be aware of each time Apple Intelligence is invoked?
- Will users be aware of how each of their queries are processed—under Case I, II, or III?²⁵
- Can users freely choose which Cases to allow? Could a user permit Case I (local processing)?
- How is it determined whether any given query will fall under Case I, II, or III?
- How will queries be processed for users who opt into Advanced Data Protection? There does not appear to be documentation regarding whether the processing of queries from these services (when Advanced Data Protection is enabled) will be handled differently by Apple Intelligence.

After our analysis and recommendations, we will return to the Apple system in Section 7.3.

3 Our Definitions and Taxonomy

Built upon the technical background of the previous section, this section presents the definitions and taxonomy that underpin our findings.

3.1 Terminology

Next, we define nine key terms we use in our recommendations and analysis. In addition to the definitions in this section, we rely throughout on key definitions introduced in Section 2 (Background), such as *encryption*, *plaintext*, *ciphertext*, *model*, *inference*, *training*, and others. While Section 2 defines basic terms from established literatures, the present section focuses on clarifying the terminology essential to make our analysis and recommendations precise within the scope of this paper. As such, our claim is not that the following are the universal definitions of these terms, but that these are the meanings we use in this paper.

Endpoint. An *end* is either the sender or a recipient of a message. An *endpoint* encompasses an “end” user and the user's device or devices (that the user has authenticated with a client application to send and receive messages²⁶).

Endpoint-local. We say that an operation, computation, or process is *endpoint-local* if it takes place entirely on the user's physical device. Endpoint-local computation is often done with a client application.

²¹See footnote 16.

²²For an example of what the log-in prompt looks like, see [116].

²³When a conversation is continued in the ChatGPT app or website, queries are processed under OpenAI's privacy policy, regardless of whether the user was initially signed in through Apple Intelligence or not [115].

²⁴ChatGPT accounts have model training on by default [114].

²⁵We know that users will be prompted before sending any query to ChatGPT, but the delegation protocol between on and off-device processing remains unclear.

²⁶Endpoint identification is an important security consideration in practice, which is largely out of scope for this paper [74].

This definition excludes any forms of implicit or virtual locality (e.g., a user-specific private cloud container) as well as any transmission of data off the user’s physical device.

Primary application content. For a given application, the *primary application content* is the content within the application that is part of the main service(s) that the application provides.

For example, in a messaging app, the primary application content consists of communication contents (including text, images, voice messages, and any other content types supported by the app). In a videoconferencing app (like Zoom), the primary application content consists of the contents of video calls (including audio and video but also other in-call content such as polls or screen-sharing).

E2EE application. We use *E2EE application* broadly to mean any application that makes claims that it offers E2EE in any part of the application functionality—for primary application content or otherwise.

Metadata. Information that is auxiliary to the content offered in an application is called *metadata*.²⁷ In E2EE applications, metadata is generally not encrypted.

For a messaging service, this includes who is communicating with whom, at what times and how frequently messages are exchanged, message lengths, and potentially many other types of data associated with a message—but not the message content.

E2EE content. We write *E2EE content* or *E2EE data* to refer to any data which a service provider explicitly or implicitly indicates is protected by E2EE in an application, and any derivatives of such data.²⁸ Many claims of E2EE by service providers are not very specific regarding which data are E2EE—in such cases, we treat a generalized claim of E2EE in an application as an implicit claim that the primary application content is end-to-end encrypted. We do not treat non-primary application content as E2EE content unless a provider has explicitly claimed that that type of content is E2EE.

For example, if a messaging provider claims that its app is E2EE, then the text, audio, images, videos, and any other message content exchanged by users—and any derivatives thereof—are E2EE content. In contrast, messaging metadata would not be E2EE content in this scenario.)

Shared model. A shared model is an AI model that may be used for inference by multiple users, and/or may be trained on multiple users’ queries. In a given E2EE context, we describe a model as *shared* if its users may include people and/or AI models other than the sender or the recipients of E2EE content.

Per-user model. A per-user model is an AI model that can be used for inference exclusively by a single user, and that is continuously improved with queries from only that user. Note that the latter implies that data from a particular user is not used to train the per-user models of other users. As further discussed in Section 4.2, this definition encompasses both “physical” and “logical” implementations, such as completely separate models for each user or a single foundational model dynamically refined during inference, respectively.

Third party. With respect to a particular piece of E2EE content, a *third party* is defined as any party which is neither the sender nor intended recipient of that content.

For example, if Alice is sending a message to Bob using a messaging platform *A*, then *A* is a third party with respect to that message. If the platform *A* outsources some processing to another company *B*, then *B* is also a third party to users’ conversations.

If Alice is querying an AI assistant Hal directly, (e.g., “Hal, what will the weather be like this evening?”) and Hal is a cloud-based AI assistant offered by platform *A*, then we do *not* consider *A* (and Hal) to be a third party with respect to that query.

²⁷The distinction between content and metadata is complex and context-specific [24].

²⁸Derivatives include any mathematically correlated function of such data (excluding bit-length, as noted in footnote 7). This could be like the examples from Section 2.1.1, whether messages are happy/sad, or aggregate statistics about message content.

3.2 A Taxonomy From “Strict E2EE” to “Not E2EE”

In this section, we present a taxonomy of applications ranging from “strict E2EE” to “not E2EE,” consisting of five categories in decreasing order of the degree to which they maintain the integrity of E2EE. As discussed in Section 2.1, E2EE applications today typically involve additional features and complexity beyond those explicitly modeled by formal definitions — and as such, they cover the full range of our first four categories, as further discussed below.

Category 1 (Strict E2EE) refers to applications where *no E2EE content is seen or used by any third party at any point in the functioning of the application, for all available configurations of application settings and features.*

No major E2EE application meets Category 1’s most stringent level of confidentiality. Certain commonplace features that enhance user experience involve the disclosure of (often small amounts of) E2EE content to third parties.²⁹ As such, even with these features, it has become common to refer to and advertise the application as “end-to-end encrypted,” both in industry and in academia, including among encryption experts. We refer to such features, which involve third-party processing of strictly limited types or quantities of plaintext-dependent content to provide a user-experience-enhancing feature, and which are commonplace in applications that provide the strongest E2EE guarantees in current practice, as *caveat features*.³⁰

Category 2 (E2EE in current practice) encompasses a range of applications that do not qualify as “Strict E2EE” due to caveat features, but otherwise provide strong E2EE guarantees. To be in Category 2, *no E2EE content may be seen or used by any third party at any point in the functioning of the application, except for caveat features,* and applications must offer each user the option of disabling all caveat features in their own use of the application. (In other words, there must exist a configuration of application settings and features in which the application operates like a Category 1 application.) We place Signal,³¹ and WhatsApp³² in Category 2.

The current and prevailing use of the term “E2EE” generally refers to Categories 1 and 2, both in industry and in academia, including among encryption experts.

Category 3 (E2EE with opt out) refers to applications which offer strong E2EE features as an option, but either these features are off by default, or the application allows users to opt out of these features. Note, in contrast, that Categories 1 and 2 definitionally preclude E2EE being off by default and/or users turning E2EE off. A Category 3 application’s E2EE features, *when enabled,* must be strong enough that *no E2EE content may be seen or used by any third party at any point in the functioning of the application, except for caveat features.* (That is, again, there must exist a configuration of settings and features in which the application operates like a Category 1 application.) We place Zoom [145] in Category 3, as it offers E2EE videoconferencing only when turned on in user settings. We consider Zoom to be in Category 3 even though a user can set E2EE as their personal default in their account settings (which means calls that that user sets up will be E2EE unless the user explicitly indicates otherwise); what we focus on is whether E2EE is off by default *for any given user account.* Notably, Category 3 applications lack the important *systemic protections* discussed in Section 2.1.2.

Category 4 (E2EE options available) encompasses applications that offer E2EE as an option for limited features, but does not offer any configuration where all primary application content is protected by E2EE. That is, Category 4 applications are not E2EE because they *always involve processing of primary application content by third parties in the functioning of the application* unless data is protected by E2EE as a narrow exception. Telegram allows users to turn on E2EE on a per-chat basis and it is always off by default [64]. At the time of writing, Meta’s Messenger is in Category 4, as it allows users to enable E2EE only in a special type of chat called “private chats” [52].

²⁹One example, already discussed in Section ??, is *link previews*.

³⁰We recognize that term “strictly limited” is ambiguous here; we do not think it is realistic or helpful to articulate a precise threshold amount of processing data that would be acceptable. That said, we would *not* consider a feature involving third-party processing of a majority of E2EE content to be “strictly limited” processing. We *would* consider the caveat features that are commonplace in the most prominent E2EE applications in current practice to involve third-party processing that is “strictly limited” in the sense we mean here.

³¹Showing the preview image of a link shared in Signal violates strict E2EE, for example. [82]

³²WhatsApp backup feature violates strict E2EE, for example. [80]

Finally, **Category 5 (Not E2EE)** describes applications that fall in none of the above categories. Category 5 does not preclude the use of encryption: for example, applications that use *encryption* but not *end-to-end* encryption fall in Category 5. As later sections elaborate, Category 5 applications are outside the scope of this paper—unless the service providers make claims that the application does offer E2EE.

4 Technical Implications of Integrating AI with E2EE

Building an AI assistant into an E2EE application involves several layers of configuration which affect the ways in which content may be used, processed, and stored. Section 4.1 discusses how integrating AI assistants into E2EE applications raises significant security concerns due to the mechanisms by which AI assistants use and process data. Section 4.2 provides an overview of existing cryptography and privacy technologies that may (seem to) mitigate some of the aforementioned security concerns, and explain how these technologies are *insufficient* to provide E2EE confidentiality for AI assistants processing E2EE content.

4.1 Security Considerations for AI assistants

As explained in Section 2, AI assistants interact with data during both training and inference. For an implementation of an AI assistant to be compatible with E2EE, it must uphold the guarantees of E2EE during both.

4.1.1 AI Assistants Interacting with E2EE Data

As discussed in Section 2.2, a model’s training data consists of vast amounts of information which can be continuously supplemented by data collected from users (e.g, queries, usage patterns) after the model is deployed and used for inference. This data would necessarily be plaintext as a model cannot be trained on encrypted text. As such, if user data is collected and incorporated into the training process, the model is inherently, itself, a derivative of E2EE content.

Additionally, models have been observed to reproduce portions of the training data in their outputs, at times producing significant portions of training data, sometimes verbatim or nearly verbatim [50, 122]. The exact reasons behind this behavior are not fully understood, but it is widespread and believed by many experts to be a fundamental limitation of current machine learning techniques, and unlikely to be eliminated in the near term or without a significant shift in the underlying approaches to AI development. As a result, violating the E2EE guarantee by training AI assistants on plaintext user data poses a tangible risk to the security of user information, as models may expose information from the training data during regular usage of the model, thereby leaking sensitive information.

The memorization and reproduction of training data in models also introduces the risk of individuals exploiting this behavior to extract certain data points and patterns. Prior work has shown that this can be done to expose information in the training data, which would additionally compromise the confidentiality of E2EE data [37]. In particular, individuals may perform *adversarial attacks*, involving using the model (often in a black-box manner) as a means to acquire private data. Adversarial attacks are an active research area, and come in various forms, including *membership inference attacks*, where an adversary can determine whether a particular data point is part of the model’s training set [128, 81]; and *training data extraction attacks*, where an adversary can recover elements of the model’s training data often by generating candidate data points and using membership inference to verify whether the items are in the training set [37].

4.1.2 Key Considerations

The risks outlined above underscore the following critical considerations.

1. *Which users can query the models?* To preserve the guarantees of E2EE, user messages should only be used to train models that can be queried by the senders and intended recipients of the messages.
2. *Which parties have access to the inputs and outputs from the model?* Intermediaries that process plaintext user queries before relaying them to the model provider may have access to inputs and outputs from the model, in the form of chat logs. In order to preserve the guarantees of E2EE, inputs




Figure 5: Evaluation framework for implementations of AI assistants

and outputs that depend on user messages should only be visible to the intended recipients of the messages.

3. *Are the inputs to the model used for anything else besides inference?* Even if the no parties and no other users can query or access logs of the model, inputs to the model may be used for other purposes—for example, the model provider may collect aggregate statistics on user queries. To be compatible with E2EE, queries should be used exclusively to fulfill users’ requests. In other words, inference should be a *stateless* process.

4.2 How Privacy Technologies Can and Cannot Help

The privacy concerns outlined in the preceding subsection have long been acknowledged by the machine learning, privacy, and security communities. As a result, there exist various techniques to enhance privacy—according to various metrics—during inference and training. In this section, we evaluate whether the types privacy protection that these techniques offer meet E2EE confidentiality.

We introduce a comprehensive framework which *identifies the key design considerations for implementing AI assistants in E2EE environments*. Our framework can be used to evaluate and compare different implementations of AI assistants (which incorporate various privacy techniques), and the degree to which they uphold E2EE security. These considerations are interdependent: at each consideration point, specific design choices can either affirm the configuration’s compliance with E2EE or breach it (thereby rendering further considerations irrelevant). Figure 5 summarizes this evaluation framework.

Based on our analysis, we evaluate the ways in which these configurations do and do not mitigate the security concerns associated with AI training and inference, and whether they are compatible with E2EE. While some of the configurations reflect existing practices, others represent novel configurations. In our analysis, we include the following privacy-enhancing techniques: *endpoint-local training and inference*, *fully homomorphic inference (FHE)*, *trusted execution environments (TEEs)*, *data sanitization*, and *differential privacy*. Our framework can be generalized beyond these examples.

Consideration #1: Where is the model located?. The first design parameter is the location where data is processed, i.e., where the model is hosted. We distinguish between *endpoint-local* and *non-endpoint-local* designs, using the definitions from Section 3.

Endpoint-local models are fully compatible with E2EE, and are the safest way to perform inference and training, since no messages are ever handled by an external server. (This holds true provided that E2EE messages are only used to train endpoint-local models of the sender and recipients, and not used to fine-tune endpoint-local models in other devices.) A further advantage of endpoint-local models is that they allow the use of unmodified plaintext messages when training, and thus can train upon full-quality data.

Using endpoint-local models is not simple, however, as most consumer grade devices are unable to support the hefty computations involved in running an AI assistant. In some circumstances, local processing is feasible: e.g., in Apple Intelligence, as discussed in Section 2.5. But even with the benefits of Apple’s custom hardware, not all AI assistant queries can be supported with the limited computational faculties of a smartphone, and certain features require sending the message in plaintext form off the device.

When non-endpoint-local models are used to answer queries, the compatibility with E2EE depends on the further design considerations #2 and #3 given subsequently.

Consideration #2: If the model is non-endpoint local, is the model “protected” or not?. For non-endpoint local models, the privacy of data depends on whether the model is stored in *plaintext* or

protected (i.e., encrypted or contained in secure hardware) form in the servers. Plaintext inference and training is not compatible with E2EE, as they would expose user data to third parties. Thus the only viable option is if data is “protected” in some way that makes it inaccessible to third parties, while still allowing certain computations (inference and training) over this data. There are various technologies that could implement this functionality—such as fully homomorphic encryption (FHE), and trusted execution environments (TEEs)—which differ on their confidentiality guarantees, performance, and other practical considerations. The compatibility of protected models with E2EE depends on the nature of the confidentiality offered by the technology being used, which is the focus of consideration #3.

Consideration #3: If the model is non-endpoint local and is protected, what are the type of the confidentiality guarantees? As discussed in Section 2.4.2, different security technologies—such as TEEs and encryption—provide different types of confidentiality under different conditions.

Established *cryptographic* techniques tend to provide the same type of security as E2EE under types of conditions, and are therefore more likely to be compatible with E2EE. Fully homomorphic encryption (FHE) provides a good example. FHE allows a party to perform arbitrary computations on encrypted data without first needing to decrypt it, with the result of the computation also obtained in encrypted form. FHE makes it possible to “outsource” a computation on a piece of private data without revealing anything about the data to the outsourcing provider. FHE would thus allow non-endpoint-local models to process user messages for inference without decrypting them—at least in theory.

FHE schemes are far less efficient than typical encryption schemes like those used in E2EE applications. Consequently, computationally intensive operations like inference are prohibitively slow under FHE, and would not result in practical turnaround times for users. For example, Zama, a company working on applications of FHE in blockchain and AI, estimates that, at present, running an “average-sized LLM” under FHE costs about \$5,000 per word [78]. While FHE is not a presently viable solution, it is an active area of research; substantial breakthroughs in the state of the art in the machine learning, hardware, and cryptography fronts will be necessary to eventually make this approach practical.

Hardware-backed security is the second major class of confidentiality guarantees, which involve using trusted hardware, such as TEEs (Section 2.4), to process user data. In this case, the AI model would be hosted inside a non-endpoint-local TEE. Processing inside the TEE would happen in plaintext, while still being shielded from outside entities by the hardware protections of the TEE.

As described in Section 2.4, hardware-backed security is different from E2EE security. Hardware solutions represent a substantial security/privacy improvement over plaintext inference, which may well be a fruitful avenue for exploration, beyond the scope of this paper. Our scope is to stress that the importance of clearly distinguishing them from E2EE and cryptographic security guarantees.

Consideration #4: If the model is non-endpoint local, is it a shared or per-user model? Regardless of all of the above, training introduces an additional critical consideration: which users can query the models trained with the E2EE data of which users? Since the model is a derivative of its training data, it is important that queries over this training data are consistent with the guarantees of E2EE.

There are two possible configurations for this design choice: *shared* or *per-user* models, as defined in Section 3. In the former, E2EE data is used to train a model that users other than the sender or the recipient can use for inference (for example, the application could maintain a single server-side model, which is trained on the E2EE data of all users and that all user devices can query). Conversely, in the latter, users can only query a model that is trained exclusively on their own E2EE data, and this data is not used to train any other models.

Shared models are not compatible with E2EE, since users would be able to receive inference outputs that are a function of other users’ E2EE data, which violates E2EE confidentiality. In practice, this means that E2EE data may be exposed to entities that can query the model (e.g., due to the possibility of leakage of training data and adversarial attacks). As we discuss in more detail in Section 4.2.1, mitigations for adversarial attacks (e.g., differential privacy) are insufficient to fully address this concern. Shared models are incompatible with training on E2EE content, irrespective of the implementation.

On the other hand, per-user models can be compatible with E2EE, provided that applications enforce strict isolation between user data, and that any components that depend on E2EE data must themselves be “protected” in the sense outlined in considerations #2 and #3. There are a number of ways an application could support per-user models. For example, a service could maintain a separate copy of a foundation model for each user in the system, which is fine-tuned with that user’s data. This approach would maintain strict

isolation, but require large amounts of resources for applications with many users. A more practical approach is to use *retrieval augmentation* techniques: the application hosts a single foundation model, which is not trained on user data, but is dynamically refined during inference using user-dependent data sources [100].³³ That is, the application would store E2EE data for every user (e.g., locally) [100], and then, during inference, the data for the querying user is retrieved and incorporated into the model in some way, which temporarily fine-tunes the model. After the query is processed, the base model is reverted back to its original state, which was not dependent on any E2EE data, resulting in a *stateless* augmentation process. While this reversion or statelessness would largely address the specific concern of Consideration #4—namely, which users can query the models trained with the E2EE data of which users—the storage and processing of the per-user E2EE data and the temporarily fine-tuned model (which is itself E2EE data) would, in turn, need to be scrutinized under Considerations #1–#3 and comply with the recommendations we put forward throughout this paper.

Implementation	Location	Protected	Guarantees	Shared or per-user
On-device, fine-tuned model	Endpoint local	No	N/A	N/A
Server-side plaintext inference without fine-tuning	Non-endpoint local	No	N/A	N/A
Server-side inference using FHE without fine-tuning	Non-endpoint local	Yes	Cryptographic	N/A
Server-side inference on separate per-user fine-tuned models on TEEs	Non-endpoint local	Yes	Hardware-backed	Per-user
Server-side inference using FHE on a global model fine-tuned with DP	Non-endpoint local	Yes	Cryptographic	Shared
Server-side inference using FHE on a global model	Non-endpoint local	Yes	Cryptographic	Shared
Apple’s proposal	Non-endpoint local	Yes	Hardware-backed	N/A

Table 1: Examples of how different implementations of AI assistants map to our design framework. **Green** tags denote an implementation compatible with E2EE, **yellow** tags denote an implementation that is not compatible with E2EE but has additional protections, and **red** tags denote implementations that are most vulnerable.

4.2.1 A Note on Privacy-Preserving Training

A number of techniques have been put forth towards mitigating adversarial attacks, such as differential privacy (DP) and data sanitization.

By definition, statistics computed on datasets reveal some information about the contents of the data. *Differential privacy* is a method for designing algorithms to compute statistics over datasets of personal information, in such a way that the results are guaranteed not to depend “too much” on any individual data entry. In contrast with encryption, which offers a binary type of security—a ciphertext must hide *all* information about the underlying plaintext—a key feature of DP is that it is parameterized by an “acceptable” amount of so-called “privacy loss” which represents dependence upon individual data entries (encoded through a parameter ϵ) [62]. DP’s configurable privacy loss parameter stands in sharp opposition to the strict confidentiality offered by encryption. Applying DP to machine learning specifically is an area of active research. Training with DP incurs computational overhead and loss in accuracy compared to a non-DP model—the less the acceptable privacy leakage, the greater the overhead and the loss [118, 123].

Data sanitization can mitigate attacks of the kind mentioned in Section 4.1.1 by removing particularly “sensitive” information from messages before adding them to a training dataset for an AI model. This of course presents the difficulty of determining what information within a message is sensitive; sensitivity depends on context and there is no generally correct heuristic for performing this procedure [36]. Even putting aside this difficulty, the focus on “sensitive” components of a message stands in sharp opposition

³³See [143] for more on retrieval augmentation.

to the strict confidentiality offered by encryption which is agnostic to the nature of message contents—and indeed provides additional security benefits by virtue of protecting mundane communications alongside sensitive ones (as discussed in Section 2.1.2).

While privacy-preserving training techniques such as the above have made important steps towards mitigating the feasibility of attacks, they do not meet the security guarantees of E2EE: they reduce information leakage in training data, but they do not prevent it completely—and their preventative measures are not based on the security of encryption, so the protection they provide is of a different nature to E2EE security.

The guarantees of private training techniques are contextual, depending on factors such as the type and structure of data, the number of data points, etc. Conversely in E2EE, confidentiality holds for any possible message under any context. Thus, the incompatibility of private training and E2EE is not a matter of computational limitations, but rather a definitional one. In other words, private training techniques create a *spectrum* of privacy—as determined by, for example, the level of noise introduced, or how many sensitive words are removed from the message—whereas E2EE is a *binary* classification.

Private training techniques themselves provide a lens to illustrate the tension between E2EE and AI training: we can frame the guarantees of E2EE as “extreme” forms of various private training techniques, which correspond to degraded, impractical instantiations of each technique. For example, the guarantees of E2EE can be cast as differential privacy with $\epsilon = 0$ (i.e., when the statistical distance between training data with or without a particular item included is 0). However, this value can only be achieved when private data is removed from training to begin with [143], at which point differential privacy is vacuous. Similarly, the guarantees of E2EE can be described as data sanitization in the case when all parts of an input are considered sensitive. Thus, sanitization techniques would remove E2EE messages entirely, which is equivalent to not training at all to begin with. As these examples show, to meet E2EE security, private training techniques essentially degrade to cases where training data is removed altogether.

5 Background II: Consent, Privacy, Users, and Areas of Law

5.1 What Users Understand About How Their Data is Used

Users are typically unaware of data practices and don’t often know what data is being collected about them, by whom, and for what use [137]. Privacy issues are often context-dependent, and decisions may be made based on how safe or trustworthy a company is perceived to be [91].

TOS and privacy policies in practice. When evaluating whether to try a new product, users cannot easily understand terms of service (TOS) and privacy policies, which are often incomprehensible in practice, exceeding reasonable standards of length and readability [96]. Research has shown that very few people read privacy policies and terms of service, and that reading these for all the apps that a person uses would take hundreds of hours a year [102]. Even if users did somehow read all the available information about how their data is being handled, it would be unrealistic to expect them to assimilate this information and act in their best interest based on it, given the vast amounts of data generated over an extended period of time [3].

Deceptive patterns. Another barrier to users’ ability to make informed decisions about their data is the use of deceptive design. These deceptive patterns (or “dark patterns”) are sets of design choices that coerce users to do things that aren’t in their best interest. Dark patterns can take the form of trick wording, comparison prevention, hidden costs, and nagging. These techniques can cause financial harm, loss of control (for instance over data usage), and discrimination [35]. In the case of privacy policies as well, dark patterns may be employed to coerce users into making choices that are beneficial to the given service, but not necessarily beneficial to the user. Deceptive patterns are often used to make users part with more data than they may be aware of or be willing to disclose [108]. Documenting and regulating dark patterns is challenging: they are deceptive, so users often are not aware of them, and users may also feel embarrassed about having been tricked, which complicates reporting [132].

Data can be surprisingly revealing. It is possible to infer surprisingly invasive information from someone’s shopping habits; even to predict a consumer’s pregnancy before close family members are made aware [61] [77]. Time-stamped record of an individual’s cell phone location may be stored by mobile network operators for extended periods of time without the user’s knowledge, and can .³⁴

³⁴*Carpenter v. United States*, 585 U.S. 296 (2018)

Data brokers. Companies that collect data from users often sell user information to data brokers who combine the data and build incredibly detailed dossiers about individuals’ profiles that are available for sale to the general public [127]. Users may not know that data brokers even exist.

The “Privacy Paradox”. Consumers tend to undervalue their privacy. This social phenomenon is captured by the “privacy paradox,” which identifies that people’s stated concerns about privacy are often not reflected in their online behaviors [4]. Thus, giving consumers a choice to opt in or out will often not effectively translate their preferences. Additionally, the sheer volume and time constraints of decision-making with regards to privacy has been known to cause “consent fatigue” [130]. Traditionally, consent has been provided in a binary format, like an on/off switch. However, as services offer wider ranges of settings and process increasingly more types of user data, privacy choices have become more granular. These consent requests compete for consumers limited time and attention, making it less likely for each request to be considered appropriately [68].

5.2 The Failure of Consent

5.2.1 Legal Consent

The concept of user *consent* plays a significant role in nearly all laws that regulate privacy [130]. From contract law to tort law to privacy statutes, notions of consent are relied upon to legitimize a wide range of data collection and processing activities that would otherwise not be permissible. Legal definitions describe consent idealistically as an individual’s voluntary agreement that is wilfully and freely manifested, in a manner that is not coerced, fraudulent, or erroneous [141]. In practice, the situations in which courts consider a user to have consented to a condition often fail to meet this idealistic description. Normatively, the legal framework of consent has been heavily criticized for insufficiently reflecting the realities of consumers and for ignoring the underlying social dynamics of data privacy.

Legal consent presupposes a prior *disclosure* or legal *notice*. In digital services, legal notice may come in the form of a privacy statement, terms of service, or banner displayed by websites and apps. In contract law, privacy notices are regarded by a great majority of courts as a valid contracts, despite the extensive research (referenced in Section 5.1) showing that these notices are not read in practice [22]. Terms of service and privacy policies are generally upheld as legally binding contracts even if users did not consent by an affirmative action like signing a document or checking a box labeled “I agree” while creating an account—American courts often treat a user’s continued use of a website or app as “implicit consent” to its terms and conditions.

The “notice and consent” (or “notice-and-choice”) approach is the prevailing framing of consent in the United States. The underlying rationale is that consumers have the choice to stop using the service entirely or to opt out of specific features if they are not in agreement with the provided terms or notice [32]. Laws like the Gramm-Leach-Bliley Act, the Children’s Online Privacy Protection Act, and the Health Insurance Portability and Accountability Act, which respectively regulate financial information, children’s data, and health data, require organizations to provide individuals with a privacy notice and an opt-out mechanism.

Another legal approach to consent is to require the “express consent” of users. This approach requires consumers to opt in by taking an affirmative action to indicate consent, such as by checking a box or clicking an “Accept” button. The European Union’s General Data Protection Regulation requires “informed,” “specific,” and “unambiguous” consent, where the term “unambiguous” is understood to reject notions of implied consent through inaction.³⁵ Due to the European Union’s regulatory strength, the so-called Brussels effect, and the GDPR’s extraterritorial scope, many privacy laws across the globe have adopted similar standards for consent [33]. In the United States, although the notice-and-choice approach is still more prevalent, some U.S. state privacy laws require express consent in specific circumstances.³⁶

5.2.2 Normative Critiques

Regardless of the consent regime adopted by privacy laws, the conditions for meaningful consent are widely regarded as fictional in normative terms. Meaningful consent would be informed, not coerced or unduly

³⁵Regulation (EU) 2016/679 (“GDPR”), Article 7 and Recital 32.

³⁶E.g., the Colorado Privacy Act requires opt-in consent for processing sensitive data. (CPA § 6-1-1308 (7))

manipulated, and require that individuals have the capability to make an appropriate risk assessment (e.g., based on a cost-benefit analysis) [130].

There are many reasons why users are unable to meaningfully consent, many of which are related to the considerations already discussed in Section 5.1. Notoriously, notice and consent is inadequate in informing individuals about an organization’s privacy practices because consumers do not read privacy notices [99].

Moreover, understanding data privacy implications often requires technical expertise. This is especially true for AI/ML systems, which are frequently opaque and referred to as “black box” models. Even the developers of many such systems do not fully understand the privacy implications of training their models on user data [76]. In general, consumers cannot be expected to understand the implications of disclosing particular pieces of their data, especially in aggregate forms. Consumers overestimate baseline legal protections [79] and underestimate the risk seemingly inconsequential disclosures pose over time [131].

Regarding the contractual relationship between consumers and online platforms, consumers do not have meaningful bargaining power or choice [41]. The monopolistic tendencies of online social networks are exacerbated by the low costs of scalability and increased value as the number of users grow, resulting in lack of real choice between services [42]. Even if alternative services exist, privacy practices and terms will often be the same among competitors. Additionally, the asymmetric nature of the relationship means that individual opt-outs require scale in order to exercise compelling power against large corporations.

Another often overlooked condition of consent is its validity across time. Viewing consent as a one-time deal does not match the ongoing nature of data processing activities. Consent may expire or may simply not be enough to capture an activity that happens too far into the future, where circumstances may have changed for either party. Attempting to solve this, some laws like the GDPR and the Colorado Privacy Act explicitly acknowledge that consent may be withdrawn at any time and that consumers have a right to request for data deletion.³⁷ Although these are technically two different user rights, in the case of recurring data processing activities, the withdrawal of consent will often require deleting personal data in order to honor the consumer’s wishes. However, considering the ongoing nature of AI/ML system-training, the deletion of data is, in general, infeasible. Thus, opting out or withdrawing consent after the data has already been used for training is rendered pointless and ineffective to stop undesired results, such as a model outputting the personal data on which it was trained.

Finally, legal consent largely elides the social dimensions of privacy. Although privacy can easily be understood to have a social connotation that highlights its value as a shared good, legal consent (in its current form) cannot, and is typically linked to an individual permission. A person is unable to consent for another unless in some legally accepted representation capacity. Multi-party messaging data raises complex issues outside the scope of existing legal consent frameworks, as legal consent does not account for the fact that shared data is extra-individual by nature. Genetic data is another example of information in this category, as it reveals information about a person’s entire family.

5.2.3 Social Dimensions of Information Privacy

In his influential book *Code and Other Laws of Cyberspace*, Larry Lessig argued that laws are only one kind of the several regulations that affect the way people behave. [92] Architecture (whether of buildings or code), markets, and social norms can regulate individuals’ behavior just as much, if not more than, law.

Defining social norms precisely may be an impossible task in its own right. Helen Nissenbaum’s theory of contextual integrity argues that the term “privacy” must be understood depending on the context. [107] The norms around privacy in a health context may be completely different than in a school context, and from group to group. For example, if a child is diagnosed with ADHD by a private psychiatrist, it would be a severe violation of the child’s privacy for the psychiatrist to disclose that information outside of the child’s family. However, if the psychiatrist is employed by the child’s school, disclosure to the child’s teacher and the school’s accommodation staff may be expected. Blanket consent fails to account for these norms and undercuts privacy—addressing the flow of data itself is the only way to fully address privacy concerns.

Linnet Taylor, Luciano Floridi, and Julie Cohen are three of many academics who argue for notions of “group privacy.” The motivating thought behind this concept is that giving each individual the last word on privacy concerning their personal data may not capture all the interests at stake. There are different

³⁷Regulation (EU) 2016/679 (“GDPR”), Article 7(3) and Article 17(1)(b) and CPA (“Colorado Privacy Act”) §6-1-1306(1)(a)(IV)(C).

flavors: Taylor and Floridi argue that informed consent can only be given by a data subject, so if the subject is a group, the consent of individuals is meaningless. [95] The definition of the relevant group greatly affects the manner in which informed consent can be given. For example, acquiring the consent of a group chat would look very different than acquiring the consent of an ethnic group. Cohen argues that reduced privacy harms the normal functioning of society in a way that manifests far beyond the individual. [43] For example, individualized search engine results distort citizens’ perception of the world and make participation in the democratic process less likely to result in compromise. Both of these approaches acknowledge that an individual’s grant or withholding of consent is unlikely to capture the full set of issues involved.

Several scholars have proposed changes to consent based on the idea of group privacy. August Bourgeois and Laurens Vandercruyse suggest a form of semi-automated consent, which would allow individuals to set their preferences accurately once and carry them throughout their online life. [31] Daniel J. Solove advocates for “murky consent,” a baseline model which would rely on data holders accepting a duty of loyalty, duty to avoid thwarting reasonable expectations, and a duty to avoid unreasonable risk. [130] Jack Balkin, and many scholars since, argues that holders of sensitive information should act as fiduciaries to data subjects [20].f

5.3 Areas of Law

We refer to several substantive areas of law that affect E2EE service providers. They are summarized below.

Contract law. Contracts are agreements between private entities. “Contract law” is a legal framework that enables enforcement of contracts by the judicial system. With few exceptions (for illegal and unconscionable conduct), contract law does not place limits on what may be included in a contract. This “freedom to contract” is especially strong in the United States.

Consumer protection law. Consumer protection law places limits on certain practices that corporations might undertake that could harm consumers. Thus, consumer protection law can be seen as one limitation on the freedom to contract. For example, one particularly voluminous subfield of consumer protection law encompasses all the safety regulations regarding car design.

Data protection law. Data protection laws are regulations that govern how personal data is handled and used. The definition of personal data may vary between jurisdictions, but generally means information that can identify individuals. These laws establish rules and guidelines for lawful processing of personal data.

Antitrust law. Antitrust laws are regulations that promote competition and prevent monopolies by regulating the conduct and organization of businesses.

6 Legal Implications of Integrating AI with E2EE

The integration of AI assistants into applications that process any form of E2EE data raises important considerations under current legal regimes. In this section, we begin with a survey of several E2EE messaging service providers’ Terms of Service and Privacy Policies to show that terms drafted by the service do not offer consumers meaningful choice in how their data is used. If a consumer aims to safeguard their data as much as possible, their only choice is to switch services. With this background in place, we overview relevant legal obligations in connection to data and consumer protection, considering how AI service providers are currently regulated in the US and in the EU. In tandem, we analyze how providers will likely be regulated if AI assistants and E2EE are integrated, highlighting aspects subject to ongoing regulation in the EU. Finally, we note regulatory trends across other jurisdictions and provide a brief note on how competition law may affect data-driven business operations.

6.1 Terms of Service and Privacy Policies of E2EE Services in Practice

As described in the prior section, when drafting terms of service and privacy policies, international companies must account for several significant regulations. In general, companies are incentivized to write overly broad terms of service that cover any possible future business development, because courts generally uphold broad contract terms, especially in the United States [99]. In the case of an E2EE messaging service, for example, a company might be incentivized to include a term in its fine print that says that it may stop offering E2EE at any time, even if all of its promotional material touts its E2EE, and even if it does currently offer strong

Service	Legible privacy policy in a single document?	E2EE Promised?	Circumstances of transfer of consumer data to third parties	Limitations to third parties on-platform mentioned in ToS
WhatsApp	Yes	“Offered”	For marketing, normal operations	None
Facebook Messenger	No - spread across complex privacy center	“On some Products, you can use end-to-end encrypted messages.”	For marketing and to advertisers on Facebook, normal operations	None
Signal	Yes	Yes	For text message verification	N/A
Telegram	Yes	For “secret chats”	For on-platform purchases, translation, and transcription	Bots can only get username, language, IP, and shared messages

Table 2: Comparing Privacy Policies of Select E2EE Messaging Services

E2EE and has no plans to cease providing it. Similarly, while E2EE messaging services may have high standards with regards to the third parties they contract with, they regularly attempt to disclaim liability for the actions of any other entity. Third party liability disclaimers are found in all four of the terms of service summarized below.³⁸ So, while E2EE messaging services may advertise lofty ideals, and even if they really do promote those ideals by deploying strong privacy technologies, their terms of service rarely make the same promises. WhatsApp’s terms of service would allow it to share user data for marketing purposes, a function that would conflict with the E2EE that it *does*, elsewhere, commit itself to providing. Such inconsistencies are common in terms of service because of the incentive to over-include.

Contractual privacy terms of select E2EE messaging services are summarized in Table 1. Note that the three services (WhatsApp, Facebook Messenger, and Telegram) that do not promise use of E2EE do not use it for messages with bots or businesses. Bots or businesses could store user messages in plaintext if they chose to—in other words, these three services have decided not to dictate third-party message handling practices, instead allowing the third-party’s privacy policy to govern these interactions. Finally, WhatsApp, Telegram, and Signal broadly market themselves as “secure” messaging platforms, and occasionally use the term “end-to-end encryption” in marketing. Facebook Messenger, on the other hand, does not mention security on its splash page, at the time of writing.

6.2 Overview of US Legal Considerations for E2EE & AI

As the sampled Terms of Service indicate, contract law is unlikely to create significant data protection to consumers in the absence of data protection and consumer protection laws.

Data protection statutes provide consumers with relatively limited protections as well. The United States does not have a federal general data protection law and sector-specific laws would only cover messaging services in the rarest edge cases. While some states, like California, have stronger general data protection laws, services can and do vary their products from state to state to comply with these laws. The GDPR may have more significant implications, although its scope is also limited, and many general data protection laws outside of the EU are modeled after the GDPR and suffer from the same limitations.

In the US, consumer protection law, a broad hammer mostly wielded by the FTC and state attorneys general, is likely to be provide stronger protections than data protection law, and has the flexibility to provide strong checks against the most egregious business practices.

³⁸This does not prevent companies from including quality assurances in contracts with vendors. In the event of a data breach, vendors can be liable to both direct customers and end consumers.[1]

6.2.1 Consumer Protection and the US Federal Trade Commission

Consumer protection law in the United States would apply to some implementations of AI in E2EE systems. At the federal level, the Federal Trade Commission Act of 1914 (15 U.S.C. § 45) created the Federal Trade Commission (FTC) and enabled it to investigate “unfair” or “deceptive” trade practices.

Deceptive trade practices—the most relevant category here—are defined as those which constitute a “a misrepresentation, omission, or other practice, that misleads the consumer acting reasonably in the circumstances, to the consumer’s detriment.” [104] The FTC investigates security-related incidents as deceptive practices. Such investigations are often prompted by data breaches and result in the FTC discovering misleading statements made to consumers about security. The FTC has discretion over which investigations to pursue—and must make choices, as it does not have the capacity to pursue every lead it encounters.

For example, increased public use and scrutiny of Zoom in 2020 led to the discovery that Zoom was not protecting all video calls with the “end-to-end, 256-bit encryption” it advertised. The company settled with the FTC, agreeing constructively to implement a comprehensive security program and cease misrepresentations related to encryption [94]. A dental software company advertising its use of “encryption,” unqualified, was found to be misleading customers because the level of encryption used did not meet industry standards [84]. A mortgage company settled with the FTC for advertising that the information it collected was encrypted, when in reality it was encrypted in transit but not once the company received it [40].

In recent years, the FTC has decided to investigate dark patterns as deceptive practices as well [57]. In 2021 it published a policy statement against negative option marketing, and in 2022, it issued a report categorizing various dark patterns and the actions the agency would take against them.

The FTC is not the only governmental agency capable of investigating deceptive practices. Many state attorneys general have similar powers granted by local versions of the FTC Act.³⁹ State attorneys general are important consumer protection enforcers because they can address smaller or more local concerns that the FTC would miss or lack the bandwidth to address. They can also be significant regulators in their own right, and collaborate between states to take on the most sophisticated corporations.

Under FTC precedent, there are clear lines that an E2EE messaging app could not cross without being found deceptive. If the app is marketed as using “end-to-end encryption,” established FTC precedents demonstrate that the app must use end-to-end encryption throughout the processing of the relevant data [94, 40], have E2EE on by default [94], and meet industry standards [84].

6.3 Overview of EU Legal Considerations for E2EE & AI

6.3.1 EU Legal Background

In many ways, legal obligations for data-driven services in the EU extend beyond the disclosure and consent model. For example, the GDPR sets out key principles that impose additional constraints upon collecting, storing, and using personal data.⁴⁰ Under the GDPR, obtaining consent is not the only basis to legitimize the processing of user data—it is just one among six. Many user data-related activities within companies can rely on the broad alternative basis of “legitimate interests,” that is, where processing is necessary for a legitimate interest pursued by the organization.⁴¹ For instance, analyzing customer data for fraud prevention and signs of criminal activity or engaging in direct marketing with preexisting clients are typically deemed activities within a company’s legitimate interests [113]. The legal basis of “legitimate interests” is often seen by organizations as more flexible and less cumbersome than obtaining user consent.⁴²

This is because the GDPR sets higher standards for user consent: it must be freely given, specific, informed, unambiguous, and may be withdrawn (or opted out) at any time. Employing dark patterns to obtain user consent or to hide mechanisms to opt out may violate the GDPR and other overlapping regulations, such as the 2022 Digital Services Act (DSA) and the Unfair Commercial Practices Directive (“UCP” Directive)[46]. The European Data Protection Board (EDPB) provided examples of deceptive

³⁹For example, the New York State Attorney General investigates deceptive practices under the Deceptive Practices Act, N.Y. Gen. Bus. Law § 349.

⁴⁰For example, the “lawfulness, fairness and transparency”, “purpose limitation” and “data minimization” principles are important guardrails enshrined in the GDPR. Regulation (EU) 2016/679 (“GDPR”), Article 5(1).

⁴¹Regulation (EU) 2016/679 (“GDPR”), Article 6(1)(f).

⁴²Despite the requirement to ensure that organizational interests do not override users’ fundamental rights (known as the “balancing test”). Regulation (EU) 2016/679 (“GDPR”), Article 6(1)(f).

techniques used by social media platforms that do not qualify as “freely-given” consent [27]. These include “forced” consent, where refusing to consent leads to a denial of the service, and “bundled” consent, when consent is bundled with the acceptance of the platform’s terms and conditions.

The GDPR is not the only data protection law in the EU. Other data protection laws do mandate obtaining user consent in certain situations. For example, the surveillance of private communications necessitates user consent.⁴³ Another example is the storing of non-essential browser cookies and other tracking technologies in user devices—a requirement that has reshaped the appearance of websites with the display of consent-collecting “cookie banners”.⁴⁴ The requirement to obtain user consent in these instances is mandated by the ePrivacy Directive (Directive (EU) 2002/58 on Privacy and Electronic Communications). Article 5 of the ePrivacy Directive establishes a general prohibition on storing information or gaining access to stored information on users’ end devices without obtaining prior consent, including listening, tapping, or other kinds of interception or surveillance of communications and related metadata. This legislation was enacted to complement the 1995 Directive on Data Protection, the GDPR’s predecessor, to secure the confidentiality of electronic communications in protection of “the right to respect for private and family life”, as set forth in the EU’s Charter of Fundamental Rights.⁴⁵

While the ePrivacy Directive clearly applies to providers of network services, it is less clear whether it applies to “over-the-top” (OTT) communication services such as WhatsApp, Telegram and Zoom. In 2017, the European Commission adopted a proposal to replace the Directive with a new ePrivacy Regulation in an effort to update the law and explicitly cover OTT services [119]. Attempts to advance the reform received significant pushback from industry players and were undercut by the passing of a temporary derogation that allows the scanning of private communications to detect online Child Sexual Abuse Material [83]. As of May 2021, the draft ePrivacy Regulation still faces complicated trilogue negotiations in the EU Parliament and is currently blocked [119]. Despite the uncertainty facing the reform proposal, the European Data Protection Board (EDPB) recommends that a new ePrivacy Regulation must under no circumstances lower the level of protection offered by the current ePrivacy Directive, as the confidentiality of electronic communications is a fundamental right that requires specific protection [30]. The EDPB further recommends that any exceptions must be narrowly tailored, and should not lead to the systematic monitoring of electronic communication content, nor allow providers or third parties to circumvent any encryption.

The EU’s recently enacted AI Act reinforces the applicability of both the ePrivacy Directive and the GDPR to AI systems, highlighting that these laws establish the basis for sustainable and responsible data processing.⁴⁶ Recital 69 of the AI Act notes that the right to privacy must be guaranteed throughout the entire lifecycle of the AI system, and that measures taken by providers to ensure compliance may include not only anonymization and encryption, but also “the use of technology that permits algorithms to be brought to the data and allows training of AI systems without the transmission between parties or copying of the raw or structured data themselves”.

6.3.2 Training AI in the EU: Ongoing Controversies

Conditions for processing user data to train AI have been a contentious topic in the European Union. Throughout 2024, tech companies like Google, Meta, X and LinkedIn have delayed or suspended plans to launch generative AI models in the EU due to regulatory inquiries into their training practices [39]. So far, these companies have been aiming to harness user-generated content to train their AI assistants or large language models and have not been openly targeting private messaging. Nevertheless, AI training practices have already been raising substantial privacy concerns in the EU.

Ireland’s Data Protection Commissioner (DPC) has actively monitored the release of several AI-powered virtual assistants fueled by EU user data. The Irish authority is the European Union’s lead regulator for many US tech companies due to their EU headquarters being based in Dublin. In June 2024, the DPC raised concerns about Meta’s privacy policy update that allowed Meta AI to be trained on public posts of EU users on Facebook and Instagram [45]. The DPC acted in response to a complaint filed by the advocacy group NOYB - European Center for Digital Rights, founded by privacy activist Max Schrems. NOYB filed

⁴³Directive (EU) 2002/58 on Privacy and Electronic Communications (“ePrivacy Directive”), Article 5(1).

⁴⁴Directive (EU) 2002/58 on Privacy and Electronic Communications (“ePrivacy Directive”), Article 5(3).

⁴⁵Charter of Fundamental Rights of the European Union (2000), Article 7.

⁴⁶Regulation (EU) 2024/1689 (“EU AI Act”), Recital 10 and Article 2(7).

complaints in eleven countries urging regulators to act against Meta’s “abuse of personal data for AI”, drawing attention to the feature’s lack of opt-in consent and hidden opt-out form [109].

Similarly, in August 2024, the DPC initiated proceedings against X (formerly Twitter) in the Irish High Court for concerns related to the training of X’s AI assistant Grok [44]. A few months prior, users noticed a default setting on the app, collecting user consent via a pre-ticked box, allowing X to use posts, interactions, inputs and results for AI-training [105]. The proceeding was dismissed after X agreed to suspend the use of publicly posted data on X to train Grok. In September, the DPC opened an inquiry into Google’s use of European Union users’ personal data to help develop its foundational AI Model, Pathways Language Model 2 (PaLM 2) [75]. The DPC also publicly confirmed that it has been in contact with Apple regarding its plans for Apple Intelligence AI [140]. At the broader industry level, an open letter coordinated by Meta and signed by executives from more than two dozen companies has criticized the EU’s regulatory approach for generating uncertainty and for blocking innovation in the region [97] [6].

Other regulators have taken slightly different approaches. In contrast to the EU, plans to launch Meta AI were resumed in the United Kingdom after Meta’s engagement with the UK Information Commissioner Office’s (ICO) [103]. Under the UK GDPR,⁴⁷ Meta confirmed it will be relying on “legitimate interests”—not user consent—as a legal basis to train Meta AI on the public posts and comments of users above the age of 18. The ICO mandated certain measures from Meta, including making it simpler for users to object to the data processing and providing them with a longer window to do so, as well as providing more transparency about how the data is being used, and highlighted that it did not provide final regulatory approval and that Meta is bound to demonstrate “ongoing compliance” [112]. Meta confirmed it will be sending in-app notifications to inform users about the feature and maintains that private messages are not within the scope of training data [69].

Regulator inquiries and guidelines from EU member states have typically operated on a case-by-case basis, taking into consideration each AI system’s business model and operational specifications. The European Data Protection Board (EDPB), which oversees the consistent application of data protection rules across the EU, began harmonizing the GDPR’s interpretation regarding AI training this year. Following the launch of ChatGPT and the Italian regulator’s enforcement action against OpenAI, the EDPB launched a dedicated task force to investigate privacy issues in the ChatGPT service [29]. The main findings for AI training were based on OpenAI’s practice of “web scraping” (the automatic extraction of data from third-party platforms). This practice is different from obtaining data posted by users on proprietary platforms, as is currently the case in MetaAI and Grok’s processing operations. Nevertheless, some of the EDPB’s recommendations may be more broadly applicable, such as the requirement that users be clearly and demonstrably informed that any prompts, including text queries and file uploads, may be used for training purposes.

Aiming to trigger more comprehensive discussions at the supranational level, the DPC requested the EDPB for a supervisory opinion on relevant issues when processing personal data to develop and train AI models [44]. The EDPB’s much-awaited Opinion 48/2024 was published in December 2024, focusing specifically on (1) when and how an AI model can be considered as “anonymous”; (2) how organizations can demonstrate the appropriateness of legitimate interest as a legal basis in the development and (3) deployment phases; and (4) consequences of the unlawful processing of personal data in the development phase of an AI model on subsequent processing operations [28]. With regard to balancing organizations’ “legitimate interests” and user rights, the EDPB highlighted the importance of transparency and meeting users’ reasonable expectations considering the complexity of the technologies and the variety of potential uses of personal data. The EDPB did not address user consent; the EDPB noted that the opinion does not cover all EU data protection issues in the context of AI development and deployment.

Importantly, while the EU is openly debating the lawfulness of AI training under different conditions, none of the regulator inquiries or opinions to date address the issue of using data obtained from messaging or end-to-end encrypted services. While data posted on social media platforms under public settings may generally be considered public, messaging data is private by nature.⁴⁸ Legal regimes typically recognize that private communications must be treated to a higher standard of privacy protection because of their

⁴⁷Following Brexit, the UK retained a copy of the GDPR in domestic law, which is known as the UK GDPR.

⁴⁸There are also nuances regarding the public nature of publicly posted data: “(...) the mere fact that personal data is publicly accessible does not imply that ‘the data subject has manifestly made such data public’. In order to rely on the exception laid down in Article 9(2)(e) GDPR, it is important to ascertain whether the data subject had intended, explicitly and by a clear affirmative action, to make the personal data in question accessible to the general public.” See [29].

more intimate and vulnerable qualities, as embodied by the EU’s e-Privacy directive.⁴⁹ Thus, the use of private messages shifts the paradigm and necessarily calls for more comprehensive debates around user consent. Moreover, training on what is advertised or commonly thought to be E2EE data may frustrate users’ legitimate expectations as to the privacy and security of their data.

Although the regulation of AI training in the EU is still in flux and conclusions are uncertain, the recent inquiries in the EU and the UK have indicated that simply updating the company’s privacy policy is insufficient to comply with GDPR standards. Besides providing users and regulators with proactive transparency and disclosure, the substance of how companies are collecting personal data, the types of data harnessed and how they are being used also matter. If private messages are used as training data, questions related to the indispensability of user consent, adequate interfaces for collecting consent and options to exercise control over data become substantially more pressing.

6.4 Other Jurisdictions

The EU GDPR is seen by some as the gold standard for online privacy protection [133]. Multiple jurisdictions have enacted new regulations or updated existing privacy laws, many of which emulate the GDPR’s provisions, indicating a global trend to increase data regulation. Due to the GDPR’s influential role, European regulators’ stance on AI training issues will likely trigger ripple effects across the globe [33].

For instance, following the Irish regulator’s reaction to Meta AI, Brazil’s data privacy regulator (National Agency for Data Protection (ANPD)) similarly suspended Meta’s privacy policy update allowing the training of Meta AI on the data of Brazilian users [55]. Brazil’s first comprehensive data protection regulation (the “LGPD”) was enacted in 2020 and is heavily inspired by the GDPR. Brazil is reported to be WhatsApp’s second biggest market worldwide, with more than 124 million active users [126]. The ANPD raised concerns as to Meta’s legal basis for processing sensitive and non-sensitive data, transparency in how the changes in the privacy policy were communicated, and accessibility to exercise the right to opt out.

The regulator’s technical report also noted discrepancies in Meta’s standards for EU and Brazilian users, despite similarities between the GDPR and the LGPD. Unlike EU users, Brazilians were required to go through eight steps to access an opt-out form hosted in a complex interface [18]. Additionally, while EU users were notified about the privacy policy update via email and app notifications, Brazilian users were not directly informed and were only able to access the update via Meta’s Privacy Policy Center. After several rounds of negotiations and subsequent fixes carried out by Meta, the ANPD lifted the suspension [56]. Meta AI was launched in Brazil in October 2024, although the feature is still undergoing full regulatory inspection.

Meta has recently faced severe enforcement actions for privacy violations in other jurisdictions that have updated their privacy laws to be more robust. In November 2024, South Korea’s privacy regulator fined Meta 21.6 billion won (roughly 14 million euros) for illegally collecting sensitive personal information from Facebook users, including data about their political views and sexual orientation, and sharing it with thousands of advertisers without users’ express consent [66]. In July 2023, Meta subsidiaries were fined by the Australian Competition and Consumer Commission a total of \$20 million Australian dollars for misleading consumers about the privacy protections in Meta-owned VPN app Onavo Protect [49]. Although the app was marketed as a privacy tool in app stores, it gathered extensive user data for Meta’s commercial purposes.

These enforcement actions show a trend of increased regulatory oversight over digital corporations’ data practices stemming from legal privacy obligations. As corporations are faced with multiple cumulative regulatory standards, they may be encouraged to conform to the combination of the most stringent regulations provided by different jurisdictions [33]. However, different jurisdictions will exercise varying degrees of leverage over tech companies based on the company’s market share and number of users by country or region, as well as the local law’s penalty provisions and the value of the currency.

6.5 Privacy, Antitrust and Consumer Data

Big tech corporations are increasingly at the center of discussions about both antitrust and data protection regulation. Access to vast amounts of consumer data is a strategic asset that may reinforce online platforms’ dominant positions by creating barriers to entry for competitors, reducing consumer choice, or unfairly

⁴⁹In the US, the “Federal Wiretap Act” (18 U.S.C. §§ 2510-2522), also known as “Title III”, prohibits the interception and disclosure of wire, oral and electronic communications.

manipulating markets [42]. Mergers and acquisitions between tech corporations may involve the merging of databases and data sharing practices in violation of privacy laws [111], at times blurring the lines between privacy and antitrust concerns. We briefly provide two illustrative examples.

At the time of Meta’s acquisition of WhatsApp in 2014, the European Commission approved the acquisition [47] with the understanding that Facebook would be unable to do automated matching between Facebook users’ accounts and WhatsApp users’ accounts [48]. However, in August 2016 WhatsApp announced updates to its terms of service and privacy policy including the possibility of linking user data across WhatsApp and Facebook. The aftermath resulted in a €110 million fine on Meta (formerly Facebook) for providing misleading information under EU merger regulation [48]. Had a separation between WhatsApp and Facebook data been enforced beyond levying a fine, this antitrust measure could have had far-reaching implications for the integration of AI assistants across WhatsApp and Facebook.

In November 2024, India’s antitrust regulator targeted Meta for anti-competitive behavior related to data sharing [106]. The Competition Commission of India (CCI) imposed a penalty of roughly 25.4 million dollars on Meta after WhatsApp users were forced to share their data with other Meta-owned companies by means of a 2021 privacy policy update with no opt-out mechanism [124]. The regulator also ordered WhatsApp to refrain from sharing user data for advertising purposes with other Meta-owned applications for five years, highlighting that the data sharing resulted in denial of market access for competitors.

These examples illustrate how data privacy decisions can lead to broader antitrust concerns beyond the scope of privacy and vice versa. While companies often have monetary incentives to access larger pools of data and integrate features throughout a broad range of products, this may come at the expense of higher scrutiny by regulators. This is particularly timely observation as the past four years have seen a wave of antitrust lawsuits against Big Tech initiated by the US and the EU [101].

7 Recommendations

Next, we present our recommendations for integrating AI assistants into E2EE applications. Our recommendations address both compatibility with E2EE from a security perspective (Recommendations 1 and 2, in Section 7.1), and disclosure and consent from a legal and normative perspective (Recommendations 3 and 4, in Section 7.2). As a recap, the summary of our recommendations is as follows:⁵⁰

1. **Training.** Using end-to-end encrypted content to train *shared AI models* is not compatible with E2EE.
2. **Processing.** Processing *E2EE content* for AI features (such as inference or training) may be compatible with end-to-end encryption only if the following recommendations are upheld:
 - (a) Prioritize *endpoint-local* processing whenever possible.
 - (b) If processing E2EE content for *non-endpoint-local* models,
 - (i) No *third party* can see or use any *E2EE content*⁵¹ without breaking encryption,⁵² and
 - (ii) A user’s *E2EE content* is exclusively used to fulfill that user’s requests.
3. **Disclosure.** Messaging providers should not make unqualified representations that they provide E2EE if the default for any conversation is that *E2EE content* is used (e.g., for AI inference or training) by any *third party*.
4. **Opt-in consent.** AI assistant features, if offered in E2EE systems, should generally be off by default and only activated via opt-in consent. Obtaining meaningful consent is complex, and requires careful consideration including but not limited to: scope and granularity of opt-in/out, ease and clarity of opt-in/out, group consent, and management of consent over time.

⁵⁰This bullet-point summary of our recommendations is identical to that in Section 1.

⁵¹Or any derivative of E2EE content (i.e., any non-trivial function of content). See Section 3.1 for full definitions.

⁵²This may sound paradoxical; see § 4.2 for detailed discussion of technical configurations that may achieve this condition.

7.1 Compatibility with E2EE Security

Next, we describe our recommendations to preserve compatibility with E2EE from a security perspective, informed by our analysis of from Section 4 and guided by our E2EE feature evaluation framework described in the Appendix A.

Recommendation 1 (Training) *Using end-to-end encrypted content to train shared AI models is not compatible with E2EE.*

As discussed in Section 4.2, training a non-endpoint-local *shared* model with E2EE data is incompatible with the security guarantees of E2EE encryption; regardless of how the data reaches the model or how the model is stored, the fact that other users can query a model that is a function of other users’s E2EE data is not compatible with E2EE confidentiality. In practice, this can lead to leakage of sensitive data or malicious extraction via adversarial attacks.

It may be helpful to consider the following simpler model, based on an analogy. We can think of AI models as a “bundle” of the training data that went into them, which is somehow encoded in the values of its parameters—not necessarily perfectly encoded,⁵³ but captured well enough that the model’s outputs are highly dependent on this bundle of data. Then, an AI assistant trained on E2EE messaging data can be thought of as a *bundle of private messages*. The model receives queries from users, performs a computation on this bundle of messages, and outputs a response. Through this lens, we can see that entities with access to the model itself—for example, the service provider, or if the model is stored directly in end-user devices—now have access to (a bundle of) messages not intended for them—a bundle of E2EE content—violating E2EE confidentiality. Further, even if no entity has direct access to this bundle (for example, if the model is somehow encrypted), any user can still *perform computations* on this bundle of messages, and receive outputs that depend on it—which are E2EE content, being derivatives of E2EE plaintext messages—with just query access to the model. In other words, model outputs can be thought of as being “tainted” by any E2EE content in the underlying training data, and thus models trained on E2EE content produce outputs that are necessarily dependent on these training data and thus are themselves E2EE content. So, training shared AI assistants with E2EE data *definitionally* undermines the confidentiality guarantees of E2EE technologies.

Recommendation 2 (Inference) *Processing E2EE content for AI features (such as inference or training) may be compatible with end-to-end encryption only if the following recommendations are upheld:*

1. *Prioritize endpoint-local processing whenever possible.*
2. *If processing E2EE content for non-endpoint-local models,*
 - (i) *No third party can see or use any E2EE content⁵⁴ without breaking encryption,⁵⁵ and*
 - (ii) *A user’s E2EE content is exclusively used to fulfill that user’s requests.*

Application developers should try to prioritize local inference whenever possible, as this option is fully compatible with E2EE.

If performing non-endpoint-local inference, in order to maintain E2EE confidentiality, it is critical that no third party (including the model provider) handles any E2EE content unencrypted. The one cryptographic tool that might permit this, FHE, is not currently practical for this use (see Section 4.2).

Another approach to protecting the confidentiality of user content when processing it for AI features would be to use TEEs, which are much more practical than FHE. As discussed in Section 2.4.2, TEE security is different from encryption security: a system that relies on TEE security for message confidentiality does not achieve E2EE security and cannot accurately be described as offering E2EE. If exploring TEE-based approaches, we recommend that service providers do not describe these as E2EE (see also Recommendation 3), and exercise great care around user expectations, consent, and potential unanticipated harms if transitioning a system from one type of security guarantee to another (see also Recommendation 4).

⁵³E.g., the bundle may be disorganized, and it may be difficult to extract every single piece of training data on demand.

⁵⁴Or any derivative of E2EE content (i.e., any non-trivial function of content). See Section 3.1 for full definitions.

⁵⁵This may sound paradoxical; see § 4.2 for detailed discussion of technical configurations that may achieve this condition.

7.2 Towards Meaningful Disclosure and Consent

In practice, legal consent and disclosure are often analyzed together. For the purposes of this section, we distinguish disclosure as the set of texts in advertising and user interfaces that the user is likely to read, despite formally not being part of a service’s privacy policy or ToS. This allows us to recommend changes to the set of communications that may actually change some individuals’ behavior. This kind of disclosure can also be described as one-way communication from the service to the user.

Consent, then, encompasses assent to the formal terms of service and privacy policies that bind the user and service in a contractual agreement. Opt-in consent mechanisms are those which allow users to take an affirmative step to agree to legal terms, such as pop-up banners for clicking a button or checking a box. These mechanisms may display brief disclosures informing users of the most important terms and may link to the pages where the complete terms and policies are displayed. Conversely, opt-out mechanisms are interfaces that provide consumers with a streamlined process to opt out of certain data processing activities after consent has been given, or as an alternative to the initial opting-in. Opt-out can therefore be considered withdrawal of consent. While users may not read all the fine print (see Section 5.1), the content contained in terms of service meaningfully affects users’ rights and are subject to legally defined floors.

Using these definitions, the line between disclosure and consent can be blurry in some situations. We include terms of service and privacy policies are certainly included under the scope of consent and not disclosure. Elements of a service’s chat interface are included in disclosure, but not necessarily consent. We treat advertising, which is generally not considered binding on a service, as a disclosure. Cookie banners, ubiquitous since the passage of GDPR, overlap both disclosure and consent. Although easily dismissed by many consumers, they are prominently displayed on websites and contain non-legal language designed to be legible, a common characteristic of disclosures. Consumers then select their cookie preferences, making a two-way communication characteristic of consent.

Next, we present our Recommendations 3 and 4, which are respectively focused on disclosure and consent.

Recommendation 3 (Disclosure) *Messaging providers should not make unqualified representations that they provide E2EE if the default for any conversation is that plaintext-dependent versions of E2EE content is used (e.g., for AI inference or training) by any third party.*

“Unqualified representation” is a term often used in consumer protection law. It generally refers to a statement made without caveats, and, in this paper, excludes caveats made in fine print, such as footnotes or terms of service. So, a splash page for an E2EE messaging service which boldly states “E2EE Used for All Messages*” would be making an unqualified representation, even if the asterisk points to a footnote that says “Unless opted out by user.” or “Except when AI features are used.” Recommendation 3 says that messaging services should not make unqualified claims that they are using E2EE when their implementation does not achieve E2EE confidentiality.

Moreover, as discussed in Section 6, to avoid regulator scrutiny for deceptive practices, messaging services hoping to make AI available to or train AI in E2EE contexts should err on the side of over-disclosure.

The requirements are clear for many possible implementations. A service with a configuration which falls into Category 1 (as defined in Section 3.2) would be able to advertise itself as using E2EE with no qualifications or caveats. On the other hand, services with configurations which fall into Categories 3–5 would need to make *qualified* claims regarding their use of E2EE. Most Category 2 services will require additional, specific disclosures based on their features.

For example, in WhatsApp, many points of the advertising and app interface indicate that “text and voice messages” are E2EE. Many users may see messages invoking Meta AI, which are not encrypted, as a “text or voice message,” and thus expect that their requests to Meta AI *are* E2EE. Meta reinforces this false, but natural, misunderstanding. When entering a Meta AI chat for the first time, WhatsApp users are shown a screen which states “Personal messages stay private. Meta may use your AI messages to improve AI quality. Personal messages are never sent to Meta, can’t be read, and remain end-to-end encrypted.” The scope of “personal messages” is critical, and ambiguous, in this disclosure. A sufficient disclosure would explicitly state that the messages in the following conversation are not encrypted and are sent to Meta, and as such should not include sensitive information.

Besides clarity and accuracy, another important component of disclosure is giving users notice of said disclosures. In other words, services should ensure that their disclosures are actually reaching users. Since

E2EE services often do not require an email address to sign up, relevant disclosure updates should be sent to all users via the service itself.

Recommendation 4 (Consent, opt-in and opt-out) *AI assistant features, if offered in E2EE systems, should generally be off by default and only activated via opt-in consent. Obtaining meaningful consent is complex, and requires careful consideration including but not limited to: scope and granularity of opt-in/out, ease and clarity of opt-in/out, group consent, and management of consent over time.*

Despite normative critiques highlighting the insufficiency of consent, consent is still the main legal tool currently available for giving users greater control over their privacy rights, and implementing consent interfaces well can make an important difference.

We recommend that opt-in express consent be the standard mechanism for allowing messaging services to train AI on user data. Consent for using messaging services must not be conflated with consent for additional features such as AI assistants. Consent for separate features should be collected in separate instances, each with a specifically tailored opt-in banner or other form of disclosure. Although the banner should refer to a more detailed privacy policy and/or terms of service, there should be enough information on the banner to meet previously mentioned disclosure requirements.

Messaging services should recognize the inherently private nature of communications, even (or especially) in environments that mix E2EE and non-E2EE data. Opt-in consent banners should serve the dual purpose of (i) clearly informing users about the exact scope of data that will be used for training purposes, including whether the data is protected by E2EE or not (or by some other security mechanism such as a TEE); (ii) providing users with a real choice, making sure to disclose how their choices may affect E2EE protection. It should be made clear to users which data is kept under E2EE (or TEE) protection, including any differences between individual and group settings. Additionally, terms of service and privacy or security policies should accurately describe all technical security implementations and associated risks.

Digital services should also provide clear opt-out mechanisms upon signing up, including a clear disclosure of what features still work if a user opts out. Each option should carry its own opt-out mechanism. A clear opt-out mechanism is one that is not concealed by dark patterns or unnecessary obstructions.⁵⁶ As a rule of thumb, opting out or withdrawing consent for services to which users have previously accepted must be made as easily accessible as the initial conditions of consent for using the feature. When possible, opting out from allowing AI training on one’s data should not be a precondition to the use of an entire app or to the use of AI features in the app. Relatedly, services should consider adding separating the opt-in/opt-out for use of an AI feature and for training based on such use. Especially regarding AI training, it must be clearly disclosed that the opt-out is forward-looking: data which the company has already used for training will remain in their models, regardless of any opt-out exercised.⁵⁷ Finally, consent should be subject to periodical renewals and opt-in banners should be re-validated after reasonable periods of time.

7.3 Applying Our Recommendations to Apple Intelligence

Having established our recommendations, we return to discuss the example of Apple (introduced in Section 2.5). Analyzing the strengths and shortcomings of Apple’s system, though based on the limited information available, provides a helpful illustration of our recommendations in practice.

Many parts of Apple’s system are compliant with our recommendations, and provide a positive illustration of what can be done to preserve E2EE security when deploying AI assistants. Some parts of Apple’s system raise significant concerns related to our recommendations, and other parts are not well enough documented to support a detailed analysis of compliance with our recommendations.

Apple’s system has three modes of processing, as described in Section 2.5. We proceed case by case.

Case I (Local Models). In Case I, no user data is used for training [13], which is compliant with Recommendation 1. Further, all processing of queries—including, importantly, queries involving E2EE content—occurs locally on device [12]. Therefore, Case I is also fully compliant with E2EE security according to Recommendation 2, and features relying on Case I may be publicized as E2EE without qualifications according to Recommendation 3. Notably, Apple Intelligence prioritizes Case I whenever possible, as specified

⁵⁶For instance, requiring users to navigate an eight-step process to access the opt-out form is not a clear mechanism [18].

⁵⁷

in Recommendation 2(a). Lastly, at the time of writing, Apple Intelligence is off by default, and has to be manually activated by users in their device settings [53], which is in line with Recommendation 4.

Case II (Server Model). Just like Case I, no user data is used to train Apple’s PCC models [13], which is compliant with Recommendation 1. However, notwithstanding the security-conscious design choices within Apple’s PCC, processing inference queries inside trusted hardware provides a different kind of security from E2EE security, and thus is not compatible with E2EE under Recommendation 2. Therefore, in line with Recommendation 3, Apple should not advertise Case II as E2EE without qualifications; rather, Apple should clearly articulate the limitations introduced by the TEE processing, and explicitly state the contexts in which E2EE applications (both native and third-party) no longer meet E2EE guarantees. Such notices should be prominently displayed. Importantly, even users who have Apple Intelligence off should be aware of these limitations, due to the systemic E2EE properties mentioned in Section 2.1.2. Finally, Case II appears to adhere to Recommendation 4, as Apple Intelligence as a whole is off by default and, to our limited knowledge, seems to have reasonable opt-in consent mechanisms.

Case III (Third-Party Model). On Recommendation 1, if users do not sign in to their ChatGPT accounts, then OpenAI does not use requests to train its models [16]—thus, Case III processing is compliant with Recommendation 1 for users who are not logged in. However, if a user *is* logged into their OpenAI account and has not configured their settings otherwise⁵⁸, queries outsourced to ChatGPT may be used for training OpenAI models [115], directly opposing Recommendation 1. Thus, overall, Case III does not uphold Recommendation 1.

On Recommendation 2, if Siri is integrated within an application that is E2EE, then processing queries with ChatGPT via Siri would be in violation of Recommendation 2(b)(i), as E2EE content would be used by a third party. This is also true if ChatGPT can respond to requests on third-party E2EE or ADP-E2EE applications. Furthermore, if a user has logged into their OpenAI account and has not configured their settings otherwise, ChatGPT can log requests and use query data to train its models. Using ChatGPT to process queries on E2EE content would therefore be in violation of Recommendation 2(b)(ii), since the user’s E2EE content would not solely be used for the purpose of fulfilling that user’s requests.

Hence, according to Recommendation 3, Apple should not advertise Case III as E2EE without qualifications; it should clearly disclose the contexts in which E2EE data could be processed by third parties and how those third parties could use it.

Finally, since the ChatGPT extension in Siri is off by default and must be explicitly enabled in settings [16], Case III abides by Recommendation 4. Furthermore, it seems users are always made aware of when ChatGPT is being used to respond to a query [15]. Even if a user chooses not to be asked for permission before each query is sent to OpenAI, a small banner remains at the bottom of the response stating that it has been generated via ChatGPT. An additional disclosure that content sent to ChatGPT is no longer E2EE could further clarify security guarantees to users.

8 Discussion

8.1 Crypto Wars 3.0

Law enforcement agencies and the global intelligence community have sought over decades to both limit the use of strong encryption and exploit the naturally occurring weaknesses of encrypted systems [23].⁵⁹ These efforts are in response to the technical reality that end-to-end encryption provides confidentiality to anyone who uses it. The so-called “Crypto Wars” can be described as the confrontation across regulatory, corporate, and research settings between state power and encryption proponents to determine whether there exists a technical compromise that would facilitate “lawful access” to confidential data [125].

The Crypto Wars have evolved as encryption has evolved and technical proposals have evolved from weak cryptographic standards to various forms of key escrow [2]. The latest battles in the protracted war between lawful access and secure encryption are characterized by proposed technology solutions framed as sophisticated alternatives to backdoors. Rather than mandating the implementation of weak encryption protocols that allow a third-party to exfiltrate data, design features like client-side scanning would preserve

⁵⁸ChatGPT accounts have model training on by default [114].

⁵⁹Very recent developments suggest an interesting possible shift, though it is too soon to be clear [54, 88]. See also [19].

strong encryption algorithms while accessing E2EE content in other ways, namely on the device (the “client”) and in the cloud using cryptographic computational techniques not requiring decryption.

What do backdoors, client-side scanning and homomorphic encryption have to do with AI assistants? Like data-driven algorithms, lawful access is incompatible with the privacy and confidentiality of end-to-end encryption, especially when users are uninformed of the risks. Furthermore, they are extensions or features of private messaging applications, not meant to enhance privacy and confidentiality; rather, they are designed to undermine it by accessing data that would not otherwise be accessible. A recurring concern in the Crypto Wars has been the risk of large-scale, hard-to-detect abuse and compromise: Even if we all agreed that AI assistants are beneficial, if E2EE data can be accessed by AI assistants for good, that method of access could also be used to track and target users, undermining the core of the security guarantee that E2EE provides.

8.2 Business Incentives to Integrate AI (and Fundamental Rights)

Companies with data-driven business models may see messaging data as a valuable resource for refining AI algorithms. These companies might be more inclined to collect and analyze user interactions to enhance personalization, improve accuracy, or monetize insights, than they are offer E2EE, which places strong limits on what the platform can do. The temptation to use this data for training AI models will likely grow as AI assistants become increasingly embedded in everyday devices. Adding to these incentives in the private messaging context is the fact that messaging data is a potentially valuable source of untapped human-authored data in an environment where concerns are arising about the scarcity of such data.

E2EE has become established as a key feature of private messaging and a key part of major messaging providers’ offerings and marketing. Interestingly, in this case, the business incentives to protect proprietary and user data appear aligned with human rights protections. However, within the last couple of years, we have seen a shift towards data-heavy AI features, following the public release of ChatGPT by OpenAI. In what appears to be the new business model for Apple—which for years built its brand around privacy—and its competitors, services are seeking out and leveraging large sources of user data. The speed with which companies are turning to AI tools, including incorporating them into E2EE environments, is happening at a pace that does not permit thorough community deliberation over associated risks and externalities; and the trend appears to be corporate-driven even in the absence of clear user demand, and in the face of user concerns [38, 87]. Yet informed debate and discussion to consider the compatibility of privacy tools like end-to-end encryption and AI, data-driven features requires human rights and legal advocates, as well as technologists, to carefully consider impacts and harms in a more deliberate process.

8.3 But Can’t I Just Paste That Into ChatGPT?

Of course, a user can always take their E2EE messages and paste them into an AI assistant. That doesn’t violate E2EE confidentiality. So why are we making such a big deal about the idea of E2EE content being sent from a user’s device, where the user has them decrypted anyway, to an integrated AI assistant?

At a glance, integrating AI features in an E2EE messaging app might seem like the same thing as pasting one’s own messages into an AI assistant—just streamlined, which might be a convenient benefit for the user. In fact, we believe it is different in many ways, though a full exposition is beyond this paper’s scope; interested readers may refer to the theories of privacy briefly discussed in Section 5.2.3 for additional context.

In a nutshell, though, the two key differences are in the *service provider’s promise* about an E2EE application—that it does not access or process E2EE content, that it designs its systems to uphold E2EE confidentiality—and in the *systemic nature of the breach of confidentiality*. If a user pastes their E2EE messages into an AI assistant, or indeed if they post them publicly online, that is a intentional act undertaken by an individual that may be a breach of social norms (depending on the context) but neither breaches the E2EE *service provider’s promise* or creates a *systemic breach* of confidentiality by which similar processes can happen as a matter of course, as part of a regular app feature that is easily accessible and whose frequent use is encouraged by the very service provider who offers the E2EE.

9 Conclusion

AI models are being developed at a rapid pace, and being incorporated into ubiquitous applications, including E2EE systems. In this work, we performed a detailed examination of the technical and legal tensions between AI models and E2EE applications, identifying serious concerns.

Deployment of AI features in E2EE environments should be accompanied by informed discussions on potential risks and harms. Our work represents a first step towards this; we hope this conversation continues involving free speech, privacy, and human rights advocates, civil society representatives, and others outside academia, to develop frameworks and standards for the responsible deployment of AI and E2EE. It is critical for these debates to occur promptly, before unilateral industry rollout.

Author contributions

MK had the initial idea and co-led the overall structuring of the paper and recommendations with SP. AF, BLH, and DY contributed equally to designing and writing technical background, implications, and recommendations. DF and JL contributed equally to designing and writing legal background, implications, and recommendations. SA led the Apple case study and provided research support throughout other aspects of the paper. SP was faculty lead and KC provided consultation.

Acknowledgments

We are grateful to Peter Hall, Daji Landis, Julia Len, Megan Richards, Nathalie Smuha, and Michael A. Spector for helpful discussion.

References

- [1] In re Blackbaud, Inc., Customer Data Breach Litigation, 567 F. Supp. 3d 667 (D.S.C. 2021).
- [2] Harold Abelson et al. “Keys under doormats: mandating insecurity by requiring government access to all data and communications ‡”. In: *Journal of Cybersecurity* (2015). URL: <https://doi.org/10.1093/cybsec/tyv009>.
- [3] A. Acquisti and J. Grossklags. “Privacy and rationality in individual decision making”. In: *IEEE Security Privacy* 3.1 (2005), pp. 26–33. DOI: 10.1109/MSP.2005.22.
- [4] Alessandro Acquisti and Jen Grossklags. “What Can Behavioral Economics Teach Us About Privacy?” In: *Digital Privacy: Theory, Technologies and Practices* (2007). PERMALINK: <https://perma.cc/56BA-YXXQ>. URL: <https://www.heinz.cmu.edu/~acquisti/papers/Acquisti-Grossklags-Chapter-Etrics.pdf>.
- [5] *AI (and other) Companies: Quietly Changing Your Terms of Service Could Be Unfair or Deceptive*. PERMALINK: <https://perma.cc/A2ZP-DWV7>. 2024. URL: <https://www.ftc.gov/policy/advocacy-research/tech-at-ftc/2024/02/ai-other-companies-quietly-changing-your-terms-service-could-be-unfair-or-deceptive>.
- [6] Meta et. al. “Europe needs regulatory certainty on AI”. In: *Open Letter* (2024). URL: <https://euneedsai.com/>.
- [7] Apple. *Apple Intelligence now features Image Playground, Genmoji, Writing Tools enhancements, seamless support for ChatGPT, and visual intelligence*. 2024. URL: <https://www.apple.com/newsroom/2024/12/apple-intelligence-now-features-image-playground-genmoji-and-more/>.
- [8] Apple. *Apple Platform Security*. Tech. rep. 2024. URL: https://help.apple.com/pdf/security/en_US/apple-platform-security-guide.pdf.
- [9] Apple. *Hardware Integrity*. 2024. URL: <https://security.apple.com/documentation/private-cloud-compute/hardwareintegrity>.

- [10] Apple. *iCloud data security overview*. PERMALINK: <https://perma.cc/EP89-TPED>. 2024. URL: <https://support.apple.com/en-us/102651>.
- [11] Apple. *Integrating your app with Siri and Apple Intelligence*. PERMALINK: <https://perma.cc/W7D2-ZPBM>. 2024. URL: <https://developer.apple.com/documentation/appintents/integrating-your-app-with-siri-and-apple-intelligence>.
- [12] Apple. *Introducing Apple Intelligence, the personal intelligence system that puts powerful generative models at the core of iPhone, iPad, and Mac*. PERMALINK: <https://perma.cc/YRT8-QPPA>. 2024. URL: <https://www.apple.com/newsroom/2024/06/introducing-apple-intelligence-for-iphone-ipad-and-mac>.
- [13] Apple. *Introducing Apple’s On-Device and Server Foundation Models*. PERMALINK: <https://perma.cc/UB4P-BLRJ>. 2024. URL: <https://machinelearning.apple.com/research/introducing-apple-foundation-models>.
- [14] Apple. *Legal — Ask Siri, Dictation & Privacy*. PERMALINK: <https://perma.cc/83H3-ZFA7>. 2024. URL: <https://www.apple.com/legal/privacy/data/en/ask-siri-dictation>.
- [15] Apple. *Use Apple Intelligence with Siri on iPhone*. 2024. URL: <https://support.apple.com/en-gb/guide/iphone/iph17bafef6/ios>.
- [16] Apple. *Use ChatGPT with Apple Intelligence on iPhone*. 2024. URL: <https://support.apple.com/en-mo/guide/iphone/iph00fd3c8c2/ios#:~:text=The%20ChatGPT%20Extension%20and%20privacy,tied%20to%20your%20Apple%20Account..>
- [17] Brooke Auxier et al. *Americans and Privacy: Concerned, Confused and Feeling Lack of Control Over Their Personal Information*. Pew Research Center, 2019. URL: <https://www.pewresearch.org/internet/2019/11/15/americans-and-privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal-information/>.
- [18] Maria Badillo. “Processing of Personal Data for AI Training in Brazil: Takeaways from ANPD’s Preliminary Decisions in the Meta Case”. In: *Future of Privacy Forum* (2024). PERMALINK: <https://perma.cc/24GE-ABNN>. URL: <https://fpf.org/blog/processing-of-personal-data-for-ai-training-in-brazil-takeaways-from-anpds-preliminary-decisions-in-the-meta-case/>.
- [19] Jim Baker. *Rethinking Encryption*. Lawfare. Oct. 2019. URL: <https://www.lawfaremedia.org/article/rethinking-encryption>.
- [20] Jack M. Balkin. “Information Fiduciaries and the First Amendment”. In: *UC Davis Law Review* (2016).
- [21] Ero Balsa, Helen Nissenbaum, and Sunoo Park. “Cryptography, Trust and Privacy: It’s Complicated”. In: *Proceedings of the 2022 Symposium on Computer Science and Law, CSLAW 2022, Washington DC, USA, November 1-2, 2022*. Ed. by Daniel J. Weitzner, Joan Feigenbaum, and Christopher S. Yoo. ACM, 2022, pp. 167–179. DOI: 10.1145/3511265.3550443. URL: <https://doi.org/10.1145/3511265.3550443>.
- [22] Oren Bar-Gill, Omri Ben-Shahar, and Florencia Marotta-Wurgler. “Searching for the Common Law: The Quantitative Approach of the Restatement of Consumer Contracts”. In: *University of Chicago Law Review: Vol. 84: Iss. 1, Article 2* (2017). PERMALINK: <https://perma.cc/5MF2-E7WP>. URL: <https://chicagounbound.uchicago.edu/uclrev/vol84/iss1/2/>.
- [23] Steven M. Bellovin et al. “Going Bright: Wiretapping without Weakening Communications Infrastructure”. In: *IEEE Security & Privacy* (2013). URL: <https://ieeexplore.ieee.org/document/6357177>.
- [24] Steven M. Bellovin et al. “It’s Too Complicated: How the Internet Upends Katz, Smith, and Electronic Surveillance Law”. In: (Nov. 2016). URL: <https://jolt.law.harvard.edu/assets/articlePDFs/v30/30HarvJLTech1.pdf>.
- [25] Katie Benner and Paul Mozur. “Apple Sees Value in Its Stand to Protect Security”. In: *The New York Times* (2016). URL: <https://www.nytimes.com/2016/02/21/technology/apple-sees-value-in-privacy-vow.html>.

- [26] Katie Benner and Nicole Perlroth. *How Tim Cook, in iPhone Battle, Became a Bulwark for Digital Privacy*. The New York Times. URL: <https://www.nytimes.com/2016/02/19/technology/how-tim-cook-became-a-bulwark-for-digital-privacy.html>.
- [27] European Data Protection Board. “Guidelines 03/2022 on Deceptive design patterns in social media platform interfaces: how to recognise and avoid them”. In: (2022). PERMALINK: <https://perma.cc/RLY5-FN2X>. URL: https://www.edpb.europa.eu/system/files/2023-02/edpb_03-2022_guidelines_on_deceptive_design_patterns_in_social_media_platform_interfaces_v2_en_0.pdf.
- [28] European Data Protection Board. “Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models”. In: (2024). PERMALINK: <https://perma.cc/MB7M-KMHF>. URL: https://www.edpb.europa.eu/system/files/2024-12/edpb_opinion_202428_ai-models_en.pdf.
- [29] European Data Protection Board. “Report of the work undertaken by the ChatGPT Taskforce”. In: (2024). PERMALINK: <https://perma.cc/J6UQ-ZM82>. URL: https://www.edpb.europa.eu/our-work-tools/our-documents/other/report-work-undertaken-chatgpt-taskforce_en.
- [30] European Data Protection Board. “Statement 03/2021 on the ePrivacy Regulation”. In: (2021). PERMALINK: <https://perma.cc/92MS-2HXJ>. URL: https://www.edpb.europa.eu/system/files/2021-03/edpb_statement_032021_eprivacy_regulation_en_0.pdf.
- [31] August Bourgeois and Laurens Vandercruysse. “Reinvigorating Consent: Exploring New Paradigms for Privacy and Data Sharing”. In: *Proceedings of 22nd European Conference on Computer-Supported Cooperative Work*. European Society for Socially Embedded Technologies (EUSSET), 2024. DOI: 10.48340/ecscw2024_ep08.
- [32] Andrea J. Boyack. “The Shape of Consumer Contracts”. In: *Denver Law Review* (2023).
- [33] Anu Bradford. “The Brussels Effect: How the European Union Rules the World”. In: *Oxford University Press* (2020). PERMALINK: <https://perma.cc/4AND-6QJB>. URL: <https://doi.org/10.1093/oso/9780190088583.001.0001>.
- [34] Ferdinand Brasser et al. “Software grand exposure: {SGX} cache attacks are practical”. In: *WOOT*. 2017.
- [35] H Brignull et al. *Deceptive patterns – user interfaces designed to trick you*. Apr. 2023. URL: <https://www.deceptive.design/>.
- [36] Hannah Brown et al. “What does it mean for a language model to preserve privacy?” In: *FACCT*. 2022.
- [37] Nicholas Carlini et al. “Extracting Training Data from Large Language Models”. In: *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021, pp. 2633–2650. ISBN: 978-1-939133-24-3. URL: <https://www.usenix.org/conference/usenixsecurity21/presentation/carlini-extracting>.
- [38] Mesut Cicek, Dogan Gursoy, and Lu Lu. “Adverse impacts of revealing the presence of “Artificial Intelligence (AI)” technology in product and service descriptions on purchase intentions: the mediating role of emotional trust and the moderating role of perceived risk”. In: *Journal of Hospitality Marketing & Management* (2024). DOI: 10.1080/19368623.2024.2368040. URL: <https://doi.org/10.1080/19368623.2024.2368040>.
- [39] Sam Clark and Pieter Haeck. “Europe’s privacy patrol is spoiling Big Tech’s AI party”. In: *Politico* (2024). PERMALINK: <https://perma.cc/A8YF-MM3U>. URL: <https://www.politico.eu/article/europe-privacy-patrol-vengeance-block-ai-artificial-intelligence/>.
- [40] Jessica Rich Claudia Bourne Farrell. “Mortgage Company Settles Information Security Charges”. In: (2005). URL: <https://www.ftc.gov/news-events/news/press-releases/2005/09/mortgage-company-settles-information-security-charges>.
- [41] Julie E. Cohen. “Examined Lives: Informational Privacy and the Subject as Object”. In: *52 Stan. L. Rev.* 1373-1438 (2000). PERMALINK: <https://perma.cc/VR6L-87AS>. URL: <https://scholarship.law.georgetown.edu/facpub/810>.

- [42] Julie E. Cohen. “Law for the Platform Economy”. In: *51 U.C. Davis L. Rev. 133-2014* (2017). PERMALINK: <https://perma.cc/59AP-9F32>. URL: <https://scholarship.law.georgetown.edu/facpub/2015>.
- [43] Julie E. Cohen. “What Privacy Is For”. In: *Georgetown Law Faculty Publications and Other Works* (2013). URL: <https://scholarship.law.georgetown.edu/facpub/2526/>.
- [44] Data Protection Commission. “Data Protection Commission welcomes conclusion of proceedings relating to X’s AI tool ‘Grok’”. In: *Press Release* (2024). PERMALINK: <https://perma.cc/L6F3-LGMP>. URL: <https://www.dataprotection.ie/en/news-media/press-releases/data-protection-commission-welcomes-conclusion-proceedings-relating-xs-ai-tool-grok>.
- [45] Data Protection Commission. “The DPC’s Engagement with Meta on AI”. In: *News and Media* (2024). PERMALINK: <https://perma.cc/Y65M-U72C>. URL: <https://www.dataprotection.ie/en/news-media/latest-news/dpcs-engagement-meta-ai>.
- [46] European Commission. “Consumer protection: manipulative online practices found on 148 out of 399 online shops screened”. In: *Press Release* (2023). PERMALINK: <https://perma.cc/J24T-LCEH>. URL: https://ec.europa.eu/commission/presscorner/detail/en/ip_23_418.
- [47] European Commission. “Mergers: Commission approves acquisition of WhatsApp by Facebook”. In: *Press Release* (2014). PERMALINK: <https://perma.cc/5EU4-KNWE>. URL: https://ec.europa.eu/commission/presscorner/detail/en/ip_14_1088.
- [48] European Commission. “Mergers: Commission fines Facebook €110 million for providing misleading information about WhatsApp takeover”. In: *Press Release* (2016). PERMALINK: <https://perma.cc/V9MC-YHGQ>. URL: https://ec.europa.eu/commission/presscorner/detail/en/ip_17_1369.
- [49] Australia Competition and Consumer Commission. “20m penalty for Metacompanies for conduct liable to mislead consumers”. In: *Media Release* (2023). PERMALINK: <https://perma.cc/2WPA-LQ5A>. URL: <https://www.accc.gov.au/media-release/20m-penalty-for-meta-companies-for-conduct-liable-to-mislead-consumers-about-use-of-their-data>.
- [50] A Feder Cooper and James Grimmelman. “The Files are in the Computer: Copyright, Memorization, and Generative AI”. In: *arXiv preprint arXiv:2404.12590* (2024).
- [51] Victor Costan and Srinivas Devadas. “Intel SGX explained”. In: *IACR Cryptol, EPrint Arch* (2016).
- [52] Loredana Crisan. “Launching Default End-to-End Encryption on Messenger”. In: (2023). URL: <https://about.fb.com/news/2023/12/default-end-to-end-encryption-on-messenger/>.
- [53] Andrew Cunningham. *Report: New “Apple Intelligence” AI features will be opt-in by default*. Ars Technica. PERMALINK: URL: <https://arstechnica.com/gadgets/2024/06/report-new-apple-intelligence-ai-features-will-be-opt-in-by-default/>.
- [54] Cybersecurity and Infrastructure Security Agency (CISA) et al. *Enhanced Visibility and Hardening Guidance for Communications Infrastructure*. Dec. 2024. URL: <https://www.cisa.gov/resources-tools/resources/enhanced-visibility-and-hardening-guidance-communications-infrastructure>.
- [55] Autoridade Nacional de Proteção de Dados. “ANPD determina suspensão cautelar do tratamento de dados pessoais para treinamento da IA da Meta”. In: (2024). PERMALINK: <https://perma.cc/WWA4-NQ9K>. URL: <https://www.gov.br/anpd/pt-br/assuntos/noticias/anpd-determina-suspensao-cautelar-do-tratamento-de-dados-pessoais-para-treinamento-da-ia-da-meta>.
- [56] Autoridade Nacional de Proteção de Dados. *Meta cumpre exigências da ANPD e poderá retomar, com restrições, o uso de dados pessoais para treinamento de inteligência artificial*. PERMALINK: <https://perma.cc/3M5Q-ZCJN>. 2024. URL: <https://www.gov.br/anpd/pt-br/assuntos/noticias/meta-cumpre-exigencias-da-anpd-e-podera-retomar-com-restricoes-o-uso-de-dados-pessoais-para-treinamento-de-inteligencia-artificial>.
- [57] Gregory M. Dickinson. “Privately Policing Dark Patterns”. In: *Georgia Law Review* (2023).
- [58] Artem Dinaburg. *Understanding Apple’s On-Device and Server Foundation Models release*. Accessed: 2024-10-21. 2024. URL: <https://blog.trailofbits.com/2024/06/14/understanding-apples-on-device-and-server-foundations-model-release/>.

- [59] Roger Dingledine and Nick Mathewson. “Anonymity Loves Company: Usability and the Network Effect”. In: *5th Annual Workshop on the Economics of Information Security, WEIS 2006, Robinson College, University of Cambridge, England, UK, June 26-28, 2006*. 2006. URL: <http://weis2006.econinfosec.org/docs/41.pdf>.
- [60] Discord. “Discord’s audio/video end-to-end encryption (DAVE) protocol”. In: (2024). URL: <https://github.com/discord/dave-protocol/blob/main/protocol.md>.
- [61] Charles Duhigg. “How Companies Learn Your Secrets”. In: *The New York Times Magazine* (2012). URL: <https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html>.
- [62] Cynthia Dwork. “Differential Privacy: A Survey of Results”. In: *Theory and Applications of Models of Computation, 5th International Conference, TAMC 2008, Xi’an, China, April 25-29, 2008. Proceedings*. Ed. by Manindra Agrawal et al. Vol. 4978. Lecture Notes in Computer Science. Springer, 2008, pp. 1–19. DOI: 10.1007/978-3-540-79228-4_1. URL: https://doi.org/10.1007/978-3-540-79228-4_1.
- [63] Benj Edwards. *So far, AI hasn’t been profitable for Big Tech*. Ars Technica. PERMALINK: <https://perma.cc/B3R3-GA2Y>. URL: <https://arstechnica.com/information-technology/2023/10/so-far-ai-hasnt-been-profitable-for-big-tech>.
- [64] *End-to-End Encryption FAQ*. URL: <https://tsf.telegram.org/manuals/e2ee-simple> (visited on 12/13/2024).
- [65] Apple Security Engineering et al. *Private Cloud Compute: A new frontier for AI privacy in the cloud*. 2024. URL: <https://security.apple.com/blog/private-cloud-compute/>.
- [66] Euronews. “South Korea fines Meta €14 million for illegally collecting information on Facebook users”. In: *Euronews Tech News* (2024). PERMALINK: <https://perma.cc/8MY9-X4MP>. URL: <https://www.euronews.com/next/2024/11/05/south-korea-fines-meta-14-million-for-illegally-collecting-information-on-facebook-users>.
- [67] Shufan Fei et al. “Security vulnerabilities of SGX and countermeasures: A survey”. In: *ACM Computing Surveys (CSUR)* (2021).
- [68] Geoffrey A. Fowler. “I tried to read all my app privacy policies. It was 1 million words.” In: *The Washington Post* (2022). URL: <https://www.washingtonpost.com/technology/2022/05/31/abolish-privacy-policies/>.
- [69] Stefano Fratta. “Building AI Technology for Europeans in a Transparent and Responsible Way”. In: *Meta - Newsroom for Europe, Middle East and Africa* (2024). PERMALINK: <https://perma.cc/6377-A64J>. URL: <https://about.fb.com/news/2024/06/building-ai-technology-for-europeans-in-a-transparent-and-responsible-way/>.
- [70] Jennifer Stisa Granick and Daniel Kahn Gillmor. *The Vital Role of End-to-End Encryption*. ACLU: News and Commentary. PERMALINK: <https://perma.cc/RN4P-K94X>. URL: <https://www.aclu.org/news/privacy-technology/the-vital-role-of-end-to-end-encryption>.
- [71] Nico Grant. *Google’s A.I. Search Errors Cause a Furor Online*. The New York Times. URL: <https://www.nytimes.com/2024/05/24/technology/google-ai-overview-search.html>.
- [72] Matthew Green. *Looking back at the Snowden revelations*. A Few Thoughts on Cryptographic Engineering. PERMALINK: <https://perma.cc/F6PV-SW9H>. URL: <https://blog.cryptographyengineering.com/2019/09/24/looking-back-at-the-snowden-revelations>.
- [73] Andy Greenberg. *Signal Is More Than Encrypted Messaging. Under Meredith Whittaker, It’s Out to Prove Surveillance Capitalism Wrong*. 2024. URL: <https://www.wired.com/story/meredith-whittaker-signal/>.
- [74] Britta Hale and Chelsea Komlo. “On end-to-end encryption”. In: *Cryptology ePrint Archive* (2022). URL: <https://ia.cr/2022/449>.
- [75] Padraic Halpin. “Top EU privacy regulator opens probe into Google’s AI compliance”. In: *Reuters* (2024). PERMALINK: <https://perma.cc/2LUC-U7V8>. URL: <http://archive.today/2024.09.12-004248/https://www.reuters.com/technology/artificial-intelligence/top-eu-privacy-regulator-opens-probe-into-googles-ai-compliance-2024-09-11/>.

- [76] Melissa Heikkilä. “Nobody knows how AI works”. In: *MIT Technology Review* (2024). URL: <https://www.technologyreview.com/2024/03/05/1089449/nobody-knows-how-ai-works/>.
- [77] Kashmir Hill. “How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did”. In: *Forbes* (2022). PERMALINK: <https://perma.cc/D944-EEFW>. URL: <https://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/>.
- [78] Rand Hindi. *Making ChatGPT Encrypted End-to-end*. Accessed: 2024-10-21. 2023. URL: <https://www.zama.ai/post/chatgpt-privacy-with-homomorphic-encryption>.
- [79] Chris Jay Hoofnagle and Jennifer M. Urban. “Alan Westin’s Privacy Homo Economicus”. In: *49 Wake Forest Law Review 261 (2014)*, *UC Berkeley Public Law Research Paper No. 2434800* (2014). PERMALINK: <https://perma.cc/2EWE-7DU9>. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2434800.
- [80] *How to back up your chat history | WhatsApp Help Center*. URL: https://faq.whatsapp.com/481135090640375/?cms_platform=android (visited on 12/13/2024).
- [81] Hongsheng Hu et al. “Membership inference attacks on machine learning: A survey”. In: *ACM Computing Surveys (CSUR)* 54.11s (2022), pp. 1–37.
- [82] *I link therefore I am*. 2019. URL: <https://signal.org/blog/i-link-therefore-i-am/>.
- [83] Ella Jakubowska. “A beginner’s guide to EU rules on scanning private communications: Part 1”. In: *European Digital Rights - EDRi* (2021). PERMALINK: <https://perma.cc/63ZP-2D44>. URL: <https://edri.org/our-work/a-beginners-guide-to-eu-rules-on-scanning-private-communications-part-1/>.
- [84] Kristin Madigan Jay Mayfield Jessica Lyon. “Dental Practice Software Provider Settles FTC Charges It Misled Customers About Encryption of Patient Data”. In: (2016). URL: <https://www.ftc.gov/news-events/news/press-releases/2016/01/dental-practice-software-provider-settles-ftc-charges-it-misled-customers-about-encryption-patient>.
- [85] David Kaplan, Jeremy Powell, and Tom Woller. “AMD memory encryption”. In: *White paper* (2016).
- [86] David Kaye. “Report on encryption, anonymity, and the human rights framework”. In: *Human Rights Council* (2015). URL: <https://www.undocs.org/A/HRC/29/32>.
- [87] Alec Tyson and Emma Kikuchi. *Growing public concern about the role of artificial intelligence in daily life*. 2023. URL: <https://www.pewresearch.org/short-reads/2023/08/28/growing-public-concern-about-the-role-of-artificial-intelligence-in-daily-life/>.
- [88] Mallory Knodel. *The FBI Wants You to Know it Has Not Changed its Position on Encryption. And That’s a Problem*. Tech Policy Press. Dec. 2024. URL: <https://www.techpolicy.press/the-fbi-wants-you-to-know-it-has-not-changed-its-position-on-encryption-and-thats-a-problem>.
- [89] Mallory Knodel et al. *Definition of End-to-end Encryption*. Accessed: 2024-11-26. 2023. URL: <https://datatracker.ietf.org/doc/html/draft-knodel-e2ee-definition-11>.
- [90] Sangho Lee et al. “Inferring fine-grained control flow inside {SGX} enclaves with branch shadowing”. In: *USENIX Security*. 2017.
- [91] Maeve Duggan Lee Rainie. *Privacy and Information Sharing*. Pew Research Center, 2016. URL: <https://www.pewresearch.org/internet/2016/01/14/privacy-and-information-sharing/>.
- [92] Lawrence Lessig. *Code and Other Laws of Cyberspace*. Basic Books, 1999. ISBN: 0-465-03912-X.
- [93] Tom Lewis. *How Generative AI Is Making Customer Experience Worse*. Forbes. URL: <https://www.forbes.com/councils/forbestechcouncil/2024/03/26/how-generative-ai-is-making-customer-experience-worse/>.
- [94] Caroline Schmitz Linda Holleran Kopp Ryan Mehm. “FTC Requires Zoom to Enhance its Security Practices as Part of Settlement”. In: (2020). URL: <https://www.ftc.gov/news-events/news/press-releases/2020/11/ftc-requires-zoom-enhance-its-security-practices-part-settlement>.

- [95] Bart van der Sloot Linnet Taylor Luciano Floridi. *Group Privacy, New Challenges of Data Technologies*. Springer Cham, 2017. ISBN: 978-3-319-46606-4.
- [96] Kevin Litman-Navarro. *We Read 150 Privacy Policies. They Were an Incomprehensible Disaster*. 2019. URL: <https://www.nytimes.com/interactive/2019/06/12/opinion/facebook-google-privacy-policies.html>.
- [97] Kim Mackrael. “Meta to European Union: Your Tech Rules Threaten to Squelch the AI Boom”. In: *The Washington Post* (2024). URL: <https://www.wsj.com/tech/ai/meta-to-european-union-your-tech-rules-threaten-to-squelch-ai-boom-35297c03>.
- [98] Moxie Marlinspike. *Technology preview: Private contact discovery for Signal*. PERMALINK: <https://perma.cc/DR7X-3G2E>. 2017. URL: [\url{https://signal.org/blog/private-contact-discovery/}](https://signal.org/blog/private-contact-discovery/).
- [99] Florencia Marotta-Wurgler. “Will Increased Disclosure Help? Evaluating the Recommendations of the ALI’s Principles of the Law of Software Contracts”. In: *University of Chicago Law Review: Vol. 78: Iss. 1, Article 9* (2011). PERMALINK: <https://perma.cc/BP4K-VPTQ>. URL: <https://chicagounbound.uchicago.edu/uclrev/vol78/iss1/9/>.
- [100] Kim Martineau. *What is retrieval-augmented generation?* PERMALINK: <https://perma.cc/48JR-WJ2U>. Aug. 2023. URL: <https://research.ibm.com/blog/retrieval-augmented-generation-RAG> (visited on 12/11/2024).
- [101] Elaine McArdle. “(Anti)Trust Issues”. In: *Harvard Law Bulletin* (2024). PERMALINK: <https://perma.cc/HVX2-742Z>. URL: <https://hls.harvard.edu/today/antitrust-issues/>.
- [102] Alecia M McDonald and Lorrie Faith Cranor. “The cost of reading privacy policies”. In: *I/S: A Journal of Law Policy for the Information Society* 4 (2009), pp. 543–897.
- [103] Meta. “Building AI Technology for the UK in a Responsible and Transparent Way”. In: *Newsroom for Europe, Middle East and Africa* (2024). PERMALINK: <https://perma.cc/VGZ6-QGHK>. URL: <https://about.fb.com/news/2024/09/building-ai-technology-for-the-uk-in-a-responsible-and-transparent-way/>.
- [104] James C. Miller III. *FTC Policy Statement on Deception*. Letter. URL: https://www.ftc.gov/system/files/documents/public_statements/410531/831014deceptionstmt.pdf.
- [105] Dan Milmo. “Elon Musk’s X under pressure from regulators over data harvesting for Grok AI”. In: *The Guardian* (2024). PERMALINK: <https://perma.cc/SPP5-QD5S>. URL: <https://www.theguardian.com/technology/article/2024/jul/26/elon-musks-x-under-pressure-from-regulators-over-data-harvesting-for-grok-ai>.
- [106] Dipak Mondal. “Anti-trust regulator levy Rs 213 crore penalty on Meta for ‘forcing’ WhatsApp users to share data”. In: *The New Indian Express* (2024). PERMALINK: <https://perma.cc/863A-8AVC>. URL: <https://www.newindianexpress.com/business/2024/Nov/18/anti-trust-regulator-levy-rs-213-crore-penalty-on-meta-for-forcing-whatsapp-users-to-share-data>.
- [107] Helen Nissenbaum. *Privacy in Context: Technology, Policy and the Integrity of Social Life*. Stanford Law Books, 2010. ISBN: 0804752370.
- [108] Midas Nouwens et al. “Dark Patterns after the GDPR: Scraping Consent Pop-ups and Demonstrating their Influence”. In: *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*. CHI ’20. ACM, Apr. 2020. DOI: 10.1145/3313831.3376321. URL: <http://dx.doi.org/10.1145/3313831.3376321>.
- [109] NOYB. “noyb urges 11 DPAs to immediately stop Meta’s abuse of personal data for AI”. In: *Forced Consent Consent Bypass* (2024). PERMALINK: <https://perma.cc/92R9-XX5D>. URL: <https://noyb.eu/en/noyb-urges-11-dpas-immediately-stop-metas-abuse-personal-data-ai>.
- [110] Oasis Labs. *Smart Privacy for Web3 & AI*. PERMALINK: <https://perma.cc/9CXH-8964>. URL: [\url{https://oasisprotocol.org/}](https://oasisprotocol.org/).
- [111] Information Commissioner’s Office. “Due diligence when sharing data following mergers and acquisitions”. In: *For Organizations* (). PERMALINK: <https://perma.cc/WMJ8-W8AJ>. URL: <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/data-sharing/data-sharing-a-code-of-practice/due-diligence/\#mergers>.

- [112] Information Commissioner’s Office. “ICO statement in response to Meta’s announcement on user data to train AI”. In: *News and blogs* (2024). PERMALINK: <https://perma.cc/JWT8-6RC4>. URL: <https://ico.org.uk/about-the-ico/media-centre/news-and-blogs/2024/09/ico-statement-in-response-to-metas-announcement-on-user-data-to-train-ai/>.
- [113] Information Commissioner’s Office. “When can we rely on legitimate interests?” In: *For Organizations* (). PERMALINK: <https://perma.cc/PQZ4-CM76>. URL: <https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/lawful-basis/legitimate-interests/when-can-we-rely-on-legitimate-interests/>.
- [114] OpenAI. *data ControlsFAQ*. 2024. URL: <https://help.openai.com/en/articles/7730893-data-controls-faq>.
- [115] OpenAI. *How your data is handled when you use ChatGPT through Apple’s integrations*. 2024. URL: <https://help.openai.com/en/articles/9737562-how-your-data-is-handled-when-you-use-chatgpt-through-apple-s-integrations>.
- [116] OpenAI. *Setting up ChatGPT with Apple Intelligence*. 2024. URL: <https://help.openai.com/en/articles/10269382-setting-up-chatgpt-with-apple-intelligence>.
- [117] Saurabh Panjwani et al. “Understanding the privacy-personalization dilemma for web search: a user perspective”. In: *2013 ACM SIGCHI Conference on Human Factors in Computing Systems, CHI ’13, Paris, France, April 27 - May 2, 2013*. Ed. by Wendy E. Mackay, Stephen A. Brewster, and Susanne Bødker. ACM, 2013, pp. 3427–3430. DOI: 10.1145/2470654.2466470. URL: <https://doi.org/10.1145/2470654.2466470>.
- [118] Nicolas Papernot and Abhradeep Guha Thakurta. *How to deploy machine learning with differential privacy*. Accessed: 2024-10-21. 2021. URL: <https://www.nist.gov/blogs/cybersecurity-insights/how-deploy-machine-learning-differential-privacy>.
- [119] European Parliament. “Proposal for a regulation on privacy and electronic communications”. In: *Legislative Train Schedule* (2024). PERMALINK: <https://perma.cc/796G-R3J7>. URL: <https://www.europarl.europa.eu/legislative-train/theme-connected-digital-single-market/file-jd-e-privacy-reform>.
- [120] Rafael Pass, Elaine Shi, and Florian Tramer. “Formal abstractions for attested execution secure processors”. In: *Eurocrypt*. 2017.
- [121] Trevor Perin and Moxie Marlinspike. “The Double Ratchet Algorithm”. In: (2016). URL: <https://signal.org/docs/specifications/doubleratchet/doubleratchet.pdf>.
- [122] Fabio Petroni et al. *Language Models as Knowledge Bases?* 2019. arXiv: 1909.01066 [cs.CL]. URL: <https://arxiv.org/abs/1909.01066>.
- [123] Natalia Ponomareva et al. “How to DP-fy ML: A Practical Guide to Machine Learning with Differential Privacy”. In: *Journal of Artificial Intelligence Research* 77 (July 2023), 1113–1201. ISSN: 1076-9757. DOI: 10.1613/jair.1.14649. URL: <http://dx.doi.org/10.1613/jair.1.14649>.
- [124] Reuters. “India restricts WhatsApp sharing data with other Meta entities, imposes 25.4mln fine”. In: *Reuters - Boards, Policy Regulation* (2024). URL: <https://www.reuters.com/technology/indias-competition-regulator-imposes-254-mln-fine-meta-whatsapps-2021-privacy-2024-11-18/>.
- [125] Engineering National Academies of Sciences and Medicine. *Decrypting the Encryption Debate: A Framework for Decision Makers*. The National Academies Press, 2018. URL: <https://nap.nationalacademies.org/catalog/25010/decrypting-the-encryption-debate-a-framework-for-decision-makers>.
- [126] Gadjoo Sevilla. “WhatsApp outage results in a scramble for messaging alternatives”. In: *eMarketer* (2022). PERMALINK: <https://perma.cc/LN6D-8DQ4>. URL: <https://www.emarketer.com/content/whatsapp-outage-results-scramble-messaging-alternatives>.
- [127] Justin Sherman. “Data Brokers Are a Threat to Democracy”. In: *Wired* (2021). PERMALINK: <https://perma.cc/AH8Z-JAR7>. URL: <https://www.wired.com/story/opinion-data-brokers-are-a-threat-to-democracy/>.

- [128] Reza Shokri et al. “Membership inference attacks against machine learning models”. In: *2017 IEEE symposium on security and privacy (SP)*. IEEE. 2017, pp. 3–18.
- [129] Internet Society. *Traceability in End-to-End Encrypted Environments*. en-US. PERMALINK: <https://perma.cc/4N4V-DD3U>. URL: <https://www.internetsociety.org/resources/doc/2024/traceability-in-end-to-end-encrypted-environments/> (visited on 11/19/2024).
- [130] Daniel J. Solove. “Murky Consent: An Approach to the Fictions of Consent in Privacy Law”. In: *Boston University Law Review* (2024). URL: <https://ssrn.com/abstract=4333743>.
- [131] Daniel J. Solove. “Privacy Self-Management and the Consent Dilemma”. In: *126 Harvard Law Review 1880 (2013)*, *GWU Legal Studies Research Paper No. 2012-141*, *GWU Law School Public Law Research Paper No. 2012-141* (2013). PERMALINK: <https://perma.cc/4THR-HS4F>. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2171018.
- [132] Jasmine E. McNealy Stephanie T. Nguyen. *I, Obscura — Illuminating deceptive design patterns in the wild*. 2021. URL: <https://pacscenter.stanford.edu/wp-content/uploads/2021/07/I-Obscura-Zine.pdf>.
- [133] European Data Protection Supervisor. “Data Protection”. In: *Data Protection* (). PERMALINK: <https://perma.cc/H8KV-SQXN>. URL: https://www.edps.europa.eu/data-protection/data-protection_en.
- [134] Eli Tan. “When the Terms of Service Change to Make Way for A.I. Training”. In: *The New York Times* (2024). URL: <https://www.nytimes.com/2024/06/26/technology/terms-service-ai-training.html>.
- [135] The Economist. *The world’s most valuable resource is no longer oil, but data*. URL: <https://www.economist.com/leaders/2017/05/06/the-worlds-most-valuable-resource-is-no-longer-oil-but-data>.
- [136] Jo Van Bulck et al. “LVI: Hijacking transient execution through microarchitectural load value injection”. In: *IEEE S&P*. 2020.
- [137] Max Van Kleek et al. “Better the Devil You Know: Exposing the Data Sharing Practices of Smartphone Apps”. In: *Proceedings of the 2017 CHI Conference on Human Factors in Computing Systems*. CHI ’17. Denver, Colorado, USA: Association for Computing Machinery, 2017, 5208–5220. ISBN: 9781450346559. DOI: 10.1145/3025453.3025556. URL: <https://doi.org/10.1145/3025453.3025556>.
- [138] Stephan Van Schaik et al. “Sok: Sgx. fail: How stuff gets exposed”. In: *IEEE S&P*. 2024.
- [139] Pablo Villalobos et al. “Position: Will we run out of data? Limits of LLM scaling based on human-generated data”. In: *Forty-first International Conference on Machine Learning*.
- [140] Adrian Weckler. “Meta ‘pauses’ AI data collection in EU following Irish DPC request”. In: *The Irish Independent* (2024). PERMALINK: <https://perma.cc/XDG6-ZH5H>. URL: <https://www.independent.ie/business/technology/meta-pauses-ai-data-collection-in-eu-following-irish-dpc-request/a2000443736.html>.
- [141] Wex. “Consent”. In: *Legal Information Institute* (2022). PERMALINK: <https://perma.cc/RB7X-UL4A>. URL: <https://www.law.cornell.edu/wex/consent>.
- [142] WhatsApp. “WhatsApp Encryption Overview, v.8”. In: (2024). URL: <https://whatsapp.com/security/WhatsApp-Security-Whitepaper.pdf>.
- [143] Lukas Wutschitz et al. “Rethinking privacy in machine learning pipelines from an information flow control perspective”. In: *arXiv preprint arXiv:2311.15792* (2023).
- [144] Yuanzhong Xu, Weidong Cui, and Marcus Peinado. “Controlled-channel attacks: Deterministic side channels for untrusted operating systems”. In: *IEEE S&P*. 2015.
- [145] Zoom. “Security Guide: Zoom”. In: (2021). URL: <https://explore.zoom.us/docs/doc/Zoom-Security-White-Paper.pdf>.

A E2EE Features Evaluation Framework

In the course of evaluating the many ways in which AI features might be introduced in E2EE applications, we created a framework that describes necessary considerations for determining whether a specific feature of an E2EE messaging application can be implemented such that it is still “E2EE,” which was helpful in guiding our analysis. While we applied the framework largely to the context AI assistants in E2EE messaging applications, we believe there are many other features already widely deployed in mainstream E2EE applications whose implementations would benefit from use of this framework. AI has been a catalyst for our thinking about the role of data collection and use in today’s communications apps. While simple, we hope the framework facilitates a principled discussion of the implications of features in E2EE environments.

Our framework distinguishes between three core axes that must be considered in the context of application features: (1) the *technical implementation* of the feature, and its confidentiality guarantees (or lack thereof); (2) the *data inputs* to the feature, i.e., what type of user data will interface with it; and (3) the level of *user control* over the feature. We discuss these components in more detail below.

Layer #1: technical implementation. Does the underlying implementation of the AI assistant preserve the confidentiality guarantees of E2EE?

We do so by presenting an abstract taxonomy of possible implementations, and labeling their compatibility with E2EE confidentiality. This captures the intuition that there are better and worse implementations of the same feature from the perspective of confidentiality, and that certain implementations do not admit any configuration that is confidentiality-respecting. For our purposes, there are two main confidentiality considerations for the implementation of an AI feature:

- *Does the training of the AI system preserve the confidentiality of user data?* Training on personal data is fundamentally incompatible with E2EE, and thus no implementation of training on message data can guarantee confidentiality.
- *Is inference performed in a way that preserves the confidentiality of user data?* Inference is in principle compatible with E2EE confidentiality, whether done locally (on-device) or through central servers.

Layer #2: Data inputs. The next layer of our framework considers the data that is used as input to the AI features (for both training and inference). This layer is agnostic to a particular technical implementation, and is concerned exclusively with the data being used. There are four main considerations regarding the data that is fed to an AI features:

- *What type of data is being used?* Types of data may include, for example, messages, links, images, user profile information, metadata, etc. Different AI features may be considered more or less sensitive due to the type of data it interacts with.
- *Who is the intended recipient of the data?* The intended endpoint of communication may be either other users of the application, or the service provider itself. If it is the latter, then the data is not sensitive from the service provider’s point of view.
- *What is the “semantic” value of the data?* The semantic value of data relates to the information content of the message. Data with higher semantic value is more sensitive.
- *Is the data referential?* Relevant to the prior point, some data may have “intrinsic” meaning (e.g., messages), while other data may have “referential” meaning (e.g., URLs). Intrinsic data is generally more sensitive than referential data.

Layer #3: User control. The last layer of our framework considers the level of control that the user has over an AI feature in regards to data privacy. There are four main considerations related to user control:

- *Can the user opt out?* If the feature is on by default, does the user have the ability to turn the feature off? Being able to turn a feature off is more compatible with E2EE.
- *What is the default configuration for the AI feature?* The AI feature can be either on or off by default within a user’s application, with the latter being more compatible with E2EE than the former.
- *What is the scope of the authorization?* Some features have *full scope* of authorization, i.e., turning on the feature authorizes its usage across the entire application state. Conversely, features may have

contextual scope of authorization, in which case the feature must be turned on for individual application contexts. Contextual scope gives users more fine-grained control for how the AI assistant interacts with their data. For example, in a messaging application, it may be the case that the AI assistant must be individually authorized for each chat. Contextual scope is more compatible with E2EE.

- *What is the barrier to toggling off?* Disabling an AI assistant can involve a range of friction and effort for the user, and require different levels of technical fluency. For example, an easy-to-find toggle button in the user settings requires less effort from the user and presents less friction than one that is buried in the settings, or that requires sending a request to the service provider. Low barriers to toggling off are more aligned with the principles of E2EE.

These layers are roughly ordered from most technical and objective to the most context-dependent and normative. For example, the underlying implementation of a feature (Layer #1) may reveal a user's data in plaintext, and therefore immediately preclude that feature from being compatible with E2EE just on the basis of the desired cryptographic guarantees. Conversely, arguing the acceptability of certain user control setups (Layer #3) necessitates reference to normative principles.