

Non-User Privacy Policy

Anthropic is an AI safety and research company, building reliable, interpretable, and steerable AI systems.

We've prepared this policy ("**Policy**") to explain how our large language models are 'trained' and how personal data may be used as part of the training process. We've also included information about the ways in which personal data of individuals who are not registered users of our services may incidentally be processed as part of our services, and the privacy rights such individuals may enjoy with respect to their personal data.

To understand how we collect and use information from registered users of our services, please see our [Privacy Policy](#).

1. What Claude Is

Claude is a content-based generative AI system used by many people around the world to assist with a wide range of tasks. These include improving writing, providing coding assistance, offering strategic advice, conducting business analysis, summarizing documents, and helping with various daily activities. Claude is available to individuals, developers, and businesses around the world.

2. How Our Models are Trained

Large language models such as Claude are 'trained' on a variety of content such as text, images and multimedia so that they can learn the patterns and connections between words and/or content. This training is essential so that the model performs effectively and safely.

Training our models involves a number of separate stages, but broadly these fall into two steps: (1) pre-training and (2) post-training.

ANTHROPIC

“Pre-training” involves the analysis of content in order for the model to understand the fundamentals of language – that is, syntax, the patterns and connections between words – as well as basic facts about the world. To understand the full diversity of language, the model is trained on large amounts of content (including text) that is converted into “tokens” (short chunks of text often just a few characters long, like common words, parts of longer words, or even punctuation marks). The model processes this content and builds a complex network of relationships between tokens. At the end of pre-training, the model can predict the most likely next word in a sequence based on the words that come before it, much like a sophisticated autocomplete feature.

Models do not store text like a database, nor do they simply “mash-up” or “collage” existing content. Models identify general patterns in text in order to help people create new content, and they do not have access to or pull from the original training data once the models have been trained.

“Post-training” helps make the model more useful, effective, and safe, by integrating human feedback. The model generates multiple responses to a given prompt, which are then evaluated by human reviewers based on criteria such as quality, helpfulness, safety, and accuracy. This reinforcement learning feedback loop helps align the model’s outputs with human expectations and values, gradually improving a model’s ability to provide relevant and appropriate responses. In addition, further “fine-tuning” that provides the model with the examples of appropriate outputs (e.g., specific to a particular domain or use case) helps improve the model.

We’ve described the stages of model training in more detail in Section 4 below.

3. Collection of Personal Data

Anthropic may obtain personal data from the following sources in order to train our models:

- Publicly available information via the Internet. Our automated tool, ClaudeBot, collects content from web pages that are freely available to anyone browsing the internet. Online material is broad and varied - for example, websites, publicly shared documents, programming languages, and content such as image or video. We take steps to minimize the privacy impact on individuals as part of our data

ANTHROPIC

collection practices. For example, ClaudeBot does not access password protected pages or bypass CAPTCHA controls, and it respects ‘Do Not Crawl’ signals (robots.txt). For more information on how we collect publicly available data, see <https://anthropic.com/crawl>.

- Third Party Datasets that we purchase, are made openly available, or that we obtain through agreements with third party organizations. These may include, for example, material that third parties have gathered and made openly available, physical texts, or non-public datasets obtained from businesses.
- User inputs and outputs submitted as feedback, flagged for violating our [Usage Policy](#), or where the user has opted into the use of inputs and outputs for model training.
- Data we generate internally.

When we gather this material, we are not looking for information about particular individuals, but some of it will incidentally contain personal information; for example, names, contact information, and work or biographical details.

We only use personal data included in our training data to help our models learn about language and how to understand and respond to it. We do not use such personal data to contact people, build profiles about them, to try to sell or market anything to them, or to sell the information itself to any third party.

Anthropic may also obtain personal data from our third-party partners, such as in content returned by web search providers, when we are developing or delivering product features like web search.

4. How personal information is used at each stage of training our models

Training our models can involve a number of separate stages, described below. Any of the personal information described in Section 3 may be used for these purposes.

Gathering and filtering training material. We collect a diverse range of content such as text, images and multimedia. We describe the sources for this data in more detail in [Section 3].

When we gather this material, we are not looking for information about particular individuals,

ANTHROPIC

but some of it will incidentally contain personal information; for example, names, contact information, and work or biographical details.

Pre-training. Data is prepared for pre-training. This process can include deduplicating portions that appear repeatedly, converting the remaining characters to tokens, and randomizing certain sequences. These techniques can help reduce the likelihood of a trained model reproducing content it was trained on. We train the base model so it learns the patterns of language and basic facts about the world (Section 2 explains how this works in more detail).

Making the model helpful and safe. We then refine the trained model so that it follows instructions, gives useful answers and avoids harmful ones. We do this using carefully prepared examples, feedback and ratings from people, and by training the model to follow certain principles, for example by using [Constitutional AI](#).

Safety testing and safety tools. We build internal tools that spot harmful or policy-violating content, deliberately stress-test the model, and test each model against fixed sets of questions to measure how well and how safely it performs.

Research and publishing what we learn. A core part of Anthropic's mission is to make AI safer for everyone, not just our own products but the field as a whole. We carry out research into how AI models operate and behave - for example, what makes them more or less likely to produce harmful, biased or inaccurate responses, which safety and alignment techniques are effective, and how AI is being used by and impacting society - and we publish our findings so that others can learn from and build on them.

The models we build through these stages are general-purpose models used for tasks like conversation, writing, coding, analysis and answering questions. We use techniques such as [Constitutional AI to train our models to be aligned with human values](#). Claude's constitution includes individual privacy and freedom from undue surveillance, prevention and protection from harm, and individual wellbeing as values Claude may need to weigh up when determining how to respond to a given request.

We make the finished models available through Claude.ai, our developer tools and our business products. We do not use personal information in our training material to contact people, build profiles about them, or sell that information to anyone. The same training

ANTHROPIC

material may be used at more than one of the stages above, and we may use it again to train newer versions of our models over time.

5. Your Rights and Choices

Depending on where you live and the laws that apply in your country of residence, you may enjoy certain rights regarding your personal data, as described further below. We make all reasonable efforts to respond to such rights. However, please be aware that these rights are limited, and that the process by which we may need to action your requests regarding our training dataset are complex. We may also decline a request if we have a lawful reason for doing so. That said, we strive to prioritize the protection of personal data and comply with all applicable privacy laws.

To exercise your rights, you or an authorized agent may submit a request by [contacting us](#). After we receive your request, we may verify it by requesting information sufficient to confirm your identity. Where a third-party representative submits a request on behalf of a data subject, we require evidence of authorization to act on behalf of the data subject.

We will respond to your request within the period required by the data protection law that applies to you. For example, where the EU GDPR or UK GDPR applies, we will respond within one calendar month of receiving a verifiable request, and where your request is complex or you have made a number of requests, we may extend that period by up to a further two months. Where another law applies to you and sets a different period, we will respond within that period. Anthropic will not discriminate based on the exercising of privacy rights you may have. Set out below is a summary of the rights which you may enjoy, depending on the laws that apply in your country of residence.

- **Right to know:** the right to know what personal data Anthropic processes about you, including the categories of personal data, the categories of sources from which it is collected, the business or commercial purposes for collection, and the categories of third parties to whom we disclose it.
- **Access & data portability:** the right to request a copy of the personal data Anthropic processes about you, subject to certain exceptions and conditions. In certain cases, and subject to applicable law, you have the right to port your information.

ANTHROPIC

- **Deletion:** the right to request that we delete personal data collected from you, subject to certain exceptions.
- **Correction:** the right to request that we correct inaccurate personal data Anthropic retains about you, subject to certain exceptions. Please note that we cannot guarantee the factual accuracy of Outputs. If Outputs contain factually inaccurate personal data relating to you, you can submit a correction request and we will make a reasonable effort to correct this information—but due to the technical complexity of large language models, it may not always be possible for us to do so.
- **Objection:** the right to object to processing of your personal data, including profiling conducted on grounds of public or legitimate interest. In places where such a right applies, we will no longer process the personal data in case of such objection unless we demonstrate compelling legitimate grounds for the processing which override your interests, rights, and freedoms, or for the establishment, exercise or defense of legal claims.
- **Restriction:** the right to restrict our processing of your personal data in certain circumstances.
- **Automated decision-making:** Anthropic does not engage in decision making based solely on automated processing or profiling in a manner which produces a legal effect (i.e., impacts your legal rights) or significantly affects you in a similar way (e.g., significantly affects your financial circumstances or ability to access essential goods or services).
- **Sale & targeted Anthropic marketing** of its products and services. Anthropic does not “sell” your personal data as that term is defined by applicable laws and regulations.

For more information on submitting requests relating to data used in our training process, see [here](#).

Set out below are additional details in relation to our legal bases for processing your personal data, how we disclose your personal data, how we may transfer your personal data to different regions, how we retain your personal data, and various contact details in case you wish to contact us or a supervisory body.

ANTHROPIC

6. Uses of Personal Data and Legal Bases

We will only use your personal data in accordance with applicable laws. We use your personal data for the following purposes and rely upon the following legal bases to do so:

Purpose	Legal Basis
To improve our services and conduct research, including to develop and improve our AI models (as described in more detail in section 4)	Legitimate interests It is in our legitimate interests to train and improve our AI models. This enables us to provide more accurate and helpful services to our users over time. Our AI models also offer significant societal benefits to the wider public such as supporting scientific and medical research, facilitating education, creating content, promoting software development, improving customer service and enhancing workplace productivity. Our research also benefits the AI industry and society: it investigates the safety, inner workings, and societal impact of AI models so that artificial intelligence has a positive impact on society as it becomes increasingly advanced and capable
To prevent and investigate fraud, abuse, and violations of our Usage Policy , unlawful or criminal activity, unauthorized access to or use of personal data or Anthropic systems and networks, to protect our rights and the rights of others, and to meet legal, governmental and institutional policy obligations	Legal obligation Legitimate interests It is in our legitimate interests to protect our business, and we also have a duty to cooperate with authorities.
To provide, maintain and facilitate any products and services	Legitimate interests It is in our and our users' legitimate interests to expand our product features and deliver services that enhance platform functionality and user experience.

Use of Personal Data

7. How We Disclose Personal Data

Anthropic will disclose personal data to the following categories of third parties for the purposes explained in this Policy:

- **Affiliates.** Anthropic discloses personal data between and among its affiliates and related entities, meaning an entity that controls, is controlled by, or is under common control with Anthropic. Our affiliates will only use any personal data that we share in a manner consistent with this Policy.

ANTHROPIC

- Service providers & business partners. Anthropic may disclose personal data with service providers and business partners for a variety of business purposes, including ensuring compliance with industry standards, research, auditing, and data processing. These parties will only access, process or store personal data that we share with them in accordance with our instructions.

Anthropic may also disclose personal data in the following circumstances:

- As part of a significant corporate event. If Anthropic is involved in a merger, corporate transaction, bankruptcy, or other situation involving the transfer of business assets, Anthropic will disclose your personal data as part of these corporate transactions.
- Pursuant to regulatory or legal requirements, safety, rights of others, and to enforce our rights or our terms. We may share personal data with government authorities, law enforcement, or other third parties where, based on the information available to us, we have a good-faith belief that disclosure is reasonably necessary to (i) comply with applicable law, regulation or legal process, including for legal, tax or accounting purposes, or in response to an enforceable governmental request; (ii) prevent serious harm to any person or to property; (iii) detect, prevent, or otherwise address fraud or other illegal activity; or (iv) enforce our terms, or protect the rights, property, security, or safety of Anthropic, our users, or others.

8. Data Transfers

Anthropic is a global company, and your personal data is transferred to our servers in the US, or to other countries outside the European Economic Area (“EEA”) and the UK, for the purposes outlined in section 3. You can find more details about Anthropic’s group companies, including their locations, in our [Trust Centre](#).

Where information is transferred outside the EEA or the UK, we ensure it benefits from an adequate level of data protection by relying on:

- Adequacy decisions. These are decisions from the European Commission under Article 45 GDPR (or equivalent decisions under other laws) where they recognize that a country outside of the EEA offers an adequate level of data

ANTHROPIC

protection. We transfer your information as described in “Collection of personal data” to some countries with adequacy decisions, such as the countries with EU adequacy decisions listed [here](#) and countries with UK adequacy decisions listed [here](#); or

- Standard contractual clauses. The European Commission has approved contractual clauses under Article 46 GDPR that allows companies in the EEA to transfer data outside the EEA. These (and their approved equivalent for the UK and Switzerland) are called standard contractual clauses. We rely on standard contractual clauses to transfer information as described in “Collection of personal data” to certain affiliates and third parties in countries without an adequacy decision.

In certain situations, we rely on derogations provided for under applicable data protection law to transfer information to a third country. If you have questions about international transfers or the safeguards we apply, you can [contact us](#).

9. Data Retention and Data Lifecycle

Anthropic retains your personal data for as long as reasonably necessary for the purposes and criteria outlined in this Policy. We decide how long to keep information based on factors such as: how long it remains necessary for training, testing, maintaining and improving our models, including future versions; how long the models trained on it remain in service or supported; how long we need it to investigate, reproduce or explain a model’s behavior; any limits agreed with the organizations we obtain data from; and our other legal and regulatory obligations.

When the personal data collected is no longer required by us, we and our service providers will perform the necessary procedures for destroying, deleting, erasing, or converting it into an anonymous form as permitted or required under applicable laws.

10. Contact Information

If you live in the European Economic Area (EEA), UK or Switzerland (the “European Region”), the data controller responsible for your personal data is Anthropic Ireland, Limited. If you live

ANTHROPIC

outside the European Region, the data controller responsible for your personal data is Anthropic PBC.

If you have any questions about this Policy, or have any questions, complaints or requests regarding your personal data, you can contact us as described below:

- Anthropic PBC with a registered address at 548 Market St, PMB 90375, San Francisco, CA 94104 (United States).
- Anthropic Ireland, Limited with a registered address at 6th Floor, South Bank House, Barrow Street. Dublin 4, D04 TR29 (Ireland).

You can email us at privacy@anthropic.com and contact our Data Protection Officer at dpo@anthropic.com.

Please note that under many countries' laws, you have the right to lodge a complaint with the supervisory authority in the place in which you live or work. A full list of EU supervisory authorities' contact details is available [here](#). If you live or work in the UK, you have the right to lodge a complaint with the [UK Information Commissioner's Office](#). If you live in Brazil, you have the right to lodge a complaint with the [Brazilian Data Protection Authority \(ANPD\)](#). If you live in Australia, you have the right to lodge a complaint with the [Office of the Australian Information Commissioner](#).