
This copy is for reference purposes only and may not be the current version of Anthropic's Usage Policy in force. For the current version of Anthropic's Usage Policy, please visit <https://www.anthropic.com/legal/aup>.

Usage Policy

Effective Date: November 12, 2026

Our Usage Policy (also referred to as our “Acceptable Use Policy” or “AUP”) applies to anyone who submits inputs to Anthropic’s products or services, all of whom we refer to as “users.” This includes individuals using our apps (such as [Claude.ai](#) and Claude Code), developers and businesses using our API and developer platforms, customers accessing Claude through cloud providers and authorized resellers, and the end users of products or services integrating Claude. The Usage Policy is intended to prevent our products and services from being used to create or contribute to real-world harm, and to promote the responsible use of our products and services.

Please review this Usage Policy carefully before using our products or services.

We will update our policy as our technology and the associated risks evolve or as we learn about unanticipated risks.

Our Usage Policy is divided into three categories:

- **Universal Usage Standards:** Our Universal Usage Standards are prohibitions that apply to all users and use cases.
- **High-risk Use Case Requirements:** Our High-risk Use Case Requirements apply to specific use cases that pose an elevated risk of harm to individuals.
- **Additional Use Case Guidelines:** Our Additional Use Case Guidelines apply to certain other use cases, including consumer-facing chatbots and agents, products serving minors, agentic use, and software listed in our directories.

Our Usage Policy focuses on *how* our products and services are used. Its prohibitions and requirements do not categorically prohibit use by any particular industry or line of business, provided the use of our products and services complies with this Usage Policy.

Anthropic’s Safeguards Team implements detection and monitoring to enforce our Usage Policy. If we suspect that you may have violated our Usage Policy, we may warn you or throttle, limit, suspend, or terminate your access to our products and services.

ANTHROPIC

Our products also include real-time safeguards that may block or limit certain outputs, but encountering a block does not by itself indicate a violation of the Usage Policy. For certain dual-use or high-risk use cases covered by real-time safeguards in select domains, we may offer programs where you can apply to be granted access to our products and services with adjusted safeguards, such as our [Cyber Verification Program](#) and [Life Sciences Verification Program](#), subject to program-specific eligibility criteria.

If you believe that our model outputs are potentially inaccurate, biased, or harmful, please notify us at usersafety@anthropic.com, or report it directly in our product through the “report issues” thumbs down button or similar feedback features (where available). You can read more about our Safeguards practices and recommendations in our [Safeguards Support Center](#).

This Usage Policy is calibrated to strike an optimal balance between enabling beneficial uses and mitigating potential harms. Anthropic may enter into contracts with certain governmental customers that tailor use restrictions to that customer’s public mission and legal authorities if, in Anthropic’s judgment, the contractual use restrictions and applicable safeguards are adequate to mitigate the potential harms addressed by this Usage Policy.

Universal Usage Standards

Do Not Violate Applicable Laws or Engage in Illegal Activity

This includes using our products or services to:

- Illegally produce, acquire, sell, or distribute controlled or regulated substances
- Engage in or facilitate human trafficking or prostitution
- Infringe, misappropriate, or violate the intellectual property rights of a third party
- Violate any other applicable laws or regulations in your jurisdiction

Do Not Compromise Critical Infrastructure

This includes using our products or services to:

- Facilitate the destruction or disruption of critical infrastructure such as power grids, water treatment facilities, medical devices, telecommunication networks, emergency services, or transportation systems
- Obtain unauthorized access to critical systems such as election, healthcare, or financial systems
- Interfere with the operation of military bases and related infrastructure

Do Not Compromise Computer or Network Systems Without Authorization

ANTHROPIC

This includes using our products or services to:

- Discover or exploit vulnerabilities in systems, networks, or applications without authorization of the system owner
- Gain unauthorized access to, or escalate unauthorized privileges within, systems, networks, applications, or devices through technical attacks or social engineering
- Create or install tools designed to intercept communications or monitor devices without authorization of the system owner
- Develop persistent unauthorized remote access tools designed to operate below normal system security levels, including firmware modifications or hardware implants
- Create automated tools designed to compromise multiple systems at scale without authorization
- Bypass security controls such as authentication systems, endpoint protection, or monitoring tools without authorization
- Create or distribute malware, ransomware, or other types of malicious code
- Conduct denial-of-service attacks, or develop tools for denial-of-service attacks or managing botnets

This section does not prohibit security research, testing, or tool development conducted on systems you own or operate, with the authorization of the system owner, or within the scope of an authorized bug bounty program or vulnerability disclosure policy, provided the activity complies with applicable law. For work covered by our real-time cyber safeguards, users may apply for adjusted access through our Cyber Verification Program.

Do Not Develop or Design Weapons

This includes using our products or services to:

- Develop, produce, modify, design, or test weapons, explosives, dangerous materials, weapons functionalities of weapons systems, or other systems, including retrofitted hardware, designed to deliver lethal or destructive force
- Develop or operate software or components for the testing or operation of weapons or weapons functionalities of weapons systems, including targeting, fire control, or engagement
- Weaponize or integrate weapons onto drones, vehicles, or any unmanned or autonomous platforms, or develop weapons delivery systems
- Illegally acquire, sell, or transfer weapons, or circumvent regulatory controls (including licensing or export controls) for their acquisition or transfer
- Design, synthesize, produce, acquire, operate, test, deploy, or plan the deployment of biological, chemical, radiological, or nuclear weapons, explosive or improvised explosive devices, or their precursors or delivery mechanisms, including developing

ANTHROPIC

weapons-critical technology (such as fissile material production, enrichment, or device design)

- Modify, design, or test biological or chemical agents to adjust or optimize their lethality, transmissibility, virulence, environmental persistence, or resistance to detection or medical countermeasures

Do Not Incite Violence or Hateful Behavior

This includes using our products or services to:

- Incite, facilitate, or promote violent extremism, terrorism, or hateful behavior
- Provide material support for organizations or individuals associated with violent extremism, terrorism, or hateful behavior
- Facilitate or promote any act of violence or intimidation targeting individuals, groups, animals, or property
- Engage in or promote discriminatory practices or behaviors against individuals or groups on the basis of one or more protected attributes, including race, ethnicity, religion, national origin, gender, sexual orientation, caste, or disability

Do Not Compromise Privacy or Identity Rights

This includes using our products or services to:

- Violate privacy rights as defined by applicable privacy laws, such as sharing personal information without consent or accessing private data unlawfully
- Misuse, collect, solicit, or gain access without permission to private information such as non-public contact details, health data, biometric or neural data (including facial recognition), or confidential or proprietary data
- Compile, aggregate, or enrich personal data from multiple sources to build profiles of individuals for sale, licensing, or disclosure to third parties without those individuals' awareness and agreement, in furtherance of use prohibited under this Usage Policy

Do Not Compromise Children's Safety

This includes using our products or services to:

- Create, distribute, promote, or facilitate the creation of child sexual abuse material ("CSAM"), including AI-generated CSAM, or other sexually exploitative content involving minors
- Facilitate the trafficking, sextortion, or any other form of exploitation of a minor
- Facilitate minor grooming, including generating content designed to impersonate a minor
- Facilitate child abuse of any form, including instructions for how to conceal abuse
- Promote or facilitate pedophilic relationships, including via roleplay with the model

ANTHROPIC

- Fetishize or sexualize minors, including in fictional settings or via roleplay with the model

When we detect CSAM (including AI-generated CSAM), activity facilitating the production or dissemination of CSAM, or scenarios where a minor may be at imminent risk of sexual abuse or exploitation, we will report our findings to relevant authorities. We define a minor or child to be any individual under the age of 18 years old, regardless of jurisdiction.

Any specific age eligibility requirements for use of our products will be set out in our Terms of Service.

Do Not Engage in Cruel, Abusive, or Psychologically Harmful Conduct

This includes using our products or services to:

- Promote, facilitate, or glamorize suicide or self-harm, such as soliciting methods, instructions for engaging in these behaviors, or instructions for hiding them from others
- Promote or facilitate eating disorders, including generating pro-eating disorder content or requesting methods for restriction, purging, or compensatory exercise
- Create or share content that promotes harmful idealization of body weight, shape, or type, including shaming, ranking, or critiquing people's bodies based on their size or shape, or presenting weight or shape as a measure of a person's inherent worth or discipline
- Create, distribute, threaten to share, promote, or facilitate the creation of non-consensual intimate imagery ("NCII"), including AI-generated material or tools designed to produce it
- Shame, humiliate, intimidate, bully, harass, or celebrate the suffering of individuals
- Coordinate the harassment or intimidation of an individual or group
- Promote, trivialize, or glorify graphic violence, gore, sexual violence, or animal cruelty or abuse
- Develop a new product or service, or support an existing product or service, that employs or facilitates deceptive techniques with the intent of causing emotional harm
- Engage in sustained and needless abusive or cruel behavior toward our models

Do Not Create or Spread Disinformation

This includes using our products or services to:

- Create or disseminate deceptive or misleading information about a group, entity, or person, or about events, laws, regulations, procedures, or the standards of an institution or governing body

ANTHROPIC

- Create, disseminate, or amplify conspiratorial narratives meant to target a specific group, individual, or entity
- Impersonate real individuals or entities, or create fake personas or accounts, to falsely attribute content or mislead audiences about its origin, or to convince people they are interacting with a human when they are not. This does not prohibit clearly disclosed parody, satire, or fiction.
- Create or disseminate false or misleading information related to medical, health, or science issues

Do Not Engage in Deceptive Campaigns or Artificial Activity

This includes using our products or services to:

- Create or operate fake personas, accounts, media outlets, or organizations, or impersonate real people, organizations, or media outlets, in order to mislead people about their origin or attribution, how widely a view is held, or whether they are interacting with a human (e.g., sockpuppets, astroturfing, fake reviews, or bots that claim to be human)
- Misrepresent or intentionally conceal the sponsorship of content or messaging intended to influence public opinion, officials, or other decision-makers
- Distribute content through networks of fake or ostensibly independent outlets, websites, or accounts to hide that it comes from a common source
- Manipulate the sources from which search engines or AI systems draw answers by seeding them with content that misrepresents its origin, authorship, or independence (e.g., networks of sites posing as unaffiliated sources corroborating the same claims)
- Build, improve, maintain, or contribute to tools, systems, or infrastructure designed to facilitate or enable deceptive campaigns, including astroturfing, sockpuppet account creation, coordinated inauthentic behavior, or scaled impersonation
- Market, sell, or solicit funding for any of the above as a product or service

Do Not Undermine Democratic Processes

This includes using our products or services to:

- Generate automated calls or messages to voters or public officials at scale that conceal their artificial origin
- Create content, or operate accounts or systems, that falsely appear to come from or on behalf of a real candidate, official, election authority, campaign, or news outlet, including fabricated statements, endorsements, or synthetic media not disclosed as synthetic

ANTHROPIC

- Generate or disseminate false or misleading information in political and electoral contexts, including about candidates, parties, policies, voting procedures, or election security
- Encourage, glorify or facilitate the disruption of electoral or civic processes, including interference with voting systems, vote counting, certification processes, or any other infrastructure used for voting
- Create or disseminate content or engage in behavior designed to suppress voter turnout or discourage legitimate political participation through deception or intimidation

Do Not Use for Surveillance, Censorship, or Prohibited Law Enforcement or Criminal Justice Purposes

This includes using our products or services to:

- Make or suggest decisions in law enforcement or criminal justice processes, including decisions to investigate, charge, arrest or detain, prosecute, sentence, or grant or revoke early or conditional release
- Target or track a person or their property without that person's consent by monitoring their physical location, biometrics, movements, online activities, communications, or associations, whether in real time, on an ongoing basis, or through analysis of previously collected data. This includes using our products for facial recognition, battlefield management, predictive policing, or monitoring in law enforcement investigations.
- Build, improve, or enhance the surveillance capabilities of tools, systems, or infrastructure designed for the surveillance uses prohibited in this section
- Identify, unmask, or locate people from anonymous or pseudonymous activity, or compile dossiers or aggregate personal information on individuals, in order to harass, intimidate, surveil, or cause harm to the target (e.g., doxxing)
- Utilize models to assign scores or ratings to individuals based on an assessment of their trustworthiness or social behavior without notification or their consent
- Build or support emotional recognition systems or techniques that are used to infer emotions of a natural person, except for medical or safety reasons
- Analyze or identify specific content to censor on behalf of a government organization
- Utilize models as part of any biometric categorization system for categorizing people based on their biometric data to infer their race, political opinions, trade union membership, religious or philosophical beliefs, sex life or sexual orientation
- Utilize models as part of any law enforcement application that violates or impairs the liberty, civil liberties, or human rights of natural persons

ANTHROPIC

This section does not prohibit analysis or tracking that individuals have agreed to, such as for fraud detection or anti-money laundering screening on a financial account, or consumer apps where a user has opted in. It also does not prohibit aggregated or anonymized analysis, content moderation, authorized security research, legal research and analysis by law enforcement agencies or courts, or journalism, provided these activities are not used for the purposes prohibited above.

Do Not Engage in Fraudulent, Abusive, or Predatory Practices

This includes using our products or services to:

- Facilitate the production, acquisition, or distribution of counterfeit or illicitly acquired goods
- Promote or facilitate the generation or distribution of spam
- Undertake, facilitate, or generate content for fraudulent activities, schemes, scams, extortion, phishing (including websites or domains that mimic legitimate webpages), social engineering, or malware
- Create, buy, or sell falsified documents including fake IDs, licenses, currency, government documents, or financial or employment records, or misrepresent professional credentials
- Develop, promote, or otherwise facilitate the sale or distribution of fraudulent or deceptive products
- Engage in deceptive or abusive practices that exploit individuals, including based on age, disability, or a specific social or economic situation, such as pyramid schemes, predatory lending, or abusive debt collection
- Deploy, develop, or support a product or service that uses subliminal, manipulative, or deceptive techniques to distort behavior by impairing decision-making
- Engage in actions or behaviors that circumvent the guardrails or terms of other platforms or services
- Plagiarize or submit AI-assisted work without proper permission or attribution

Do Not Abuse Our Platform

This includes using our products or services to:

- Coordinate malicious activity across multiple accounts, or create multiple accounts, to avoid detection, circumvent product guardrails, or generate identical or similar inputs that otherwise violate our Usage Policy
- Utilize automation in account creation or to engage in spammy behavior
- Circumvent a ban through the use of a different account, such as the creation of a new account, use of an existing account, or providing access to a person or entity that was previously banned
- Access, facilitate, or provide others (including end users) with access to our products or services in violation of our [Supported Regions Policy](#)

ANTHROPIC

- Intentionally bypass or interfere with capabilities, restrictions, or guardrails established within our products, including by circumvention of the account sign-up or log-in process or by instructing the model to produce harmful outputs (e.g., jailbreaking or prompt injection), without prior authorization from Anthropic
- Use stolen, fraudulent, or unauthorized payment methods, or otherwise evade payment for our services
- Use outputs to train an AI model (e.g., “model scraping” or “model distillation”) without prior authorization from Anthropic
- Resell, proxy, or otherwise provide access to Claude through unauthorized means, including services that route requests through consumer subscriptions or misrepresent the product or client being used

Do Not Generate Sexually Explicit Content

This includes using our products or services to:

- Depict or request sexual intercourse or sex acts
- Generate content related to sexual fetishes or fantasies
- Facilitate, promote, or depict incest or bestiality
- Engage in erotic chats

High-risk Use Case Requirements

Certain uses of our products and services can impact an individual's health, legal rights, finances, livelihoods, or access to housing, education, or essential services. When our products are used for these high-risk use cases, we require that a qualified person remain accountable for what is delivered or decided, and that the impacted individual knows AI was involved in the advice or decision. These requirements apply to anyone who uses our products to make a High-risk AI Recommendation or control equipment capable of High-risk Physical Actions.

These requirements for high-risk use cases are meant to protect individuals. However, they do not replace local legal requirements and restrictions. Users should always default to following local laws if those laws conflict with or add additional requirements to the requirements listed below. Our products and services may not be suitable for all use cases.

Requirements

For High-risk AI Recommendations or Physical Actions, users must implement the following controls:

ANTHROPIC

- **Qualified human-in-the-loop for advice and decisions** — A qualified person must meaningfully review a High-risk AI Recommendation, with the authority to change or modify it, before it is delivered as advice to an individual who might rely on it or before it is used to implement a decision. A qualified person is someone with the training or experience to evaluate the output in that field, and who holds a license where applicable law requires one for the underlying activity. The qualified person remains responsible for the accuracy and appropriateness of what is delivered or decided.
- **Disclosure for advice and decisions** — An individual who receives advice or who is the subject of a decision based on a High-risk AI Recommendation must be clearly told that AI was used to produce it. Users do not need to identify Anthropic, Claude, or the model used.
- **Human control and independent safety limits for physical actions** — For High-risk Physical Actions, there must be a qualified individual who is able to observe the equipment's operation and stop it at any time, and the equipment must stop or hold a safe state when that individual intervenes or when connection to our services is lost. Any operating limits to keep the equipment within safe bounds, such as speed, force, reach, temperature, pressure, voltage, energy output, dose, or operating area, must be enforced by the equipment or a controller independent of model output.

These requirements apply alongside applicable law. For physical devices, they do not replace vehicle, aviation, medical device, machinery, or workplace safety law or certification.

High-risk AI Recommendations

The requirements apply when users use our products or services to provide a recommendation in any of the following, which are “High-risk Areas”:

1. **Legal** — advising an individual on their legal rights, obligations, or liability, or on what to do in a specific legal matter, or preparing legal documents for them to sign or file.
2. **Medical** — diagnosing a condition; recommending, starting, stopping, or adjusting a treatment or medication, including dose; interpreting an individual's symptoms or test results to triage, provide a diagnosis, advise on treatment risks or side effects, or create a treatment plan. This includes assessing, diagnosing, or treating a mental health condition.
3. **Finance** — a personalized recommendation to buy, sell, hold, or allocate among specific investment products, such as securities, funds, digital assets, or annuities;

ANTHROPIC

advising an individual on their own taxes, such as what to report, deduct, or claim; or preparing a tax return for them to sign or file.

4. **Credit** — approving, denying, or setting the amount, rate, or limit of a loan, mortgage, credit line, or similar product, or reducing or closing an existing one, or producing a credit score or creditworthiness assessment used for such decisions.
5. **Insurance** — underwriting or pricing a policy; determining the outcome of a claim; making a coverage, prior-authorization, or medical-necessity determination; cancelling or declining to renew a policy.
6. **Housing** — approving, denying, or ranking a rental or purchase application; terminating a tenancy or initiating eviction or foreclosure.
7. **Employment** — screening, ranking, advancing, or rejecting candidates; hiring; promoting or demoting; setting or changing shifts, hours, or pay; evaluating a worker's performance or behavior, including through disciplining, suspending, or terminating a worker.
8. **Education and credentials** — admitting, rejecting, or ranking applicants; grading an exam or assessment that determines admission, graduation, licensure, or certification (outside of multiple-choice testing); granting, denying, or revoking a professional license or certification; making a determination of academic misconduct or suspending or expelling a student.
9. **Healthcare access** — setting triage or waitlist priority; determining eligibility for, denying, or discontinuing a treatment, service, or care program.
10. **Public benefits and services** — determining eligibility for, or the amount of, government benefits, assistance, or essential services, including suspension and termination.
11. **Legal status and adjudication** — deciding or recommending the outcome of a civil, administrative, or arbitral matter; determining immigration, asylum, or citizenship status.

Prohibition Reminder: Use of our products or services to make or suggest decisions in law enforcement or criminal justice processes is prohibited under our Universal Usage Standards and is not a High-risk AI Recommendation.

High-risk Physical Actions

When connecting our models to hardware that can take physical actions, users should assess and understand 1) the range of possible outcomes if something goes wrong, 2) the nature and extent of potential impact to individuals, groups or populations, and 3) the degree of control the user maintains relative to the autonomy the model has, including possible risks such as prompt injection or manipulation. The requirements apply when

ANTHROPIC

outputs of our products and services are acted upon by hardware without human approval where the hardware or linked equipment can:

1. **Move through shared space** — steer, accelerate, brake, or navigate a vehicle, vessel, aircraft, drone, or mobile robot in a place where people may be present, or on roads, waterways, or airspace.
2. **Apply force that could cause injury** — actuate arms, presses, lifts, hoists, doors, gates, conveyors, or other machinery capable of striking, crushing, pinning, or dropping a load on a person nearby.
3. **Control hazardous energy or materials** — operate heating elements or open flame, pressure systems, high-voltage or high-current circuits, lasers or radiation sources above eye- and skin-safe levels, or control the release, mixing, or dispensing of hazardous chemical or biological materials, at levels capable of causing injury.
4. **Act on the human body** — administer or adjust a drug, fluid, gas, electrical stimulation, or radiation or other energy dose to a person, or cut, penetrate, or move instruments within a person's body or tissue, including through a surgical, interventional, worn, attached, or implanted device.
5. **Control safety systems** — arm, disarm, trigger, or override fire suppression, emergency shutoffs, alarms, ventilation, or locks and barriers that could confine a person.
6. **Run industrial processes** — control plant or process equipment where a fault could injure workers or a member of the public (including diffuse or indirect harms, such as food poisoning or defective products).

Prohibition Reminder: Our Universal Usage Standards prohibit actions related to weapons, critical infrastructure, violence, and surveillance.

Exclusions

The requirements do not apply to:

AI Recommendations

Using our products to assist in workflows that are related to, but do not actually provide, High-risk AI Recommendations. Some non-exhaustive examples are:

- General information or educational content, including explaining what a law says, how a condition or treatment works, or how investing works, without applying it to the individual's circumstances.
- Wellness content, such as education on sleep, stress, nutrition, or exercise.

ANTHROPIC

- Explaining or helping an individual understand advice, a diagnosis, or a decision that a qualified person has already given them (as long as it is not an additional High-risk AI Recommendation).
- Internal drafting, research, summarization, or analysis that is not the ultimate High-risk AI Recommendation delivered.
- Carrying out a decision a person has already made, or applying a fixed rule, formula, or algorithm where the model exercises no judgment about the individual.
- Business operations that do not advise or decide about a specific individual.

Physical Actions

- Consumer and household devices operating within the manufacturer's built-in safety limits, where the worst possible outcome is discomfort or inconvenience.
- Generating plans, code, toolpaths, waypoints, or commands that a qualified person reviews before they are executed on equipment.
- Monitoring, sensing, or reporting on equipment without controlling it.
- Laboratory automation handling non-hazardous materials inside enclosed instruments.

Favorable Decisions

Where permitted by law, human-in-the-loop review is not required before acting on a High-risk AI Recommendation wholly favorable to the individual in the following circumstances:

- Paying an insurance claim.
- Approving coverage or prior authorization for care a clinician has ordered.
- Granting eligibility for a healthcare service or program.
- Granting or continuing a public benefit or service.

Additional legal obligations such as disclosure may still apply. A partial approval, reduced amount, or approval with conditions is not considered “wholly favorable.”

Additional Use Case Guidelines

The below use cases — regardless of whether they involve High-risk AI Recommendations or Physical Actions — must comply with the additional guidelines provided.

ANTHROPIC

- **All consumer-facing chatbots or AI agents that communicate directly with external users** must **disclose** to users that they are interacting with AI rather than a human. This disclosure must be provided at a minimum at the beginning of each chat session or in the product interface. The disclosure does not need to identify Anthropic, Claude, or the model used.
- **Products serving minors**, including organizations providing minors with the ability to directly interact with products that incorporate our API(s), must comply with the **additional guidelines** outlined in our Support [article](#) and all applicable laws.
- **Agentic use cases** must comply with the Usage Policy. Users are responsible for ensuring that agents they build or deploy, including actions those agents take through tools, browsers, or connected systems, comply with this Usage Policy.
- **Software** listed in our Directories must comply with our [Software Directory Policy](#).

We also provide additional informational resources in our [Safeguards Support Center](#), including:

- Our [Child Safety Guidance for Developers](#) for detailed guidance on identifying and mitigating child safety risks.
- Additional information about our [mental health safeguards](#) on Claude.ai.