

Click to prove
you're human



Claude AI is an artificial intelligence model that can act as a chatbot and an AI assistant, much like ChatGPT and Google's Gemini. Named after Claude E. Shannon, sometimes referred to as the "father of information theory," Claude was designed to assist with writing, coding, customer support and information retrieval. Claude was developed by Anthropic, a San Francisco-based company founded in 2021 by former OpenAI employees focusing on AI safety and research. Dario Amodei, the co-founder and CEO of Anthropic, had served as vice president of research at OpenAI. His sister, Daniela Amodei, serves as Anthropic's president. Anthropic has drawn significant investment from prominent tech players. Since 2023, Amazon has invested \$8 billion in the company. As part of the agreement, Anthropic has committed to using Amazon Web Services as its primary cloud provider and making its AI models accessible to AWS customers. The 2024 deal includes plans to expand the use of Amazon's AI chips for training and running Anthropic's large language models. Google initially invested \$500 million and plans to invest another \$1.5 billion in the future. So what can Claude do now? Here's everything you need to know, including the models, plans, pricing and latest updates. How Claude AI works Claude AI is a versatile tool capable of answering questions, generating creative content like stories and poems, translating languages, transcribing and analyzing images, writing code, summarizing text and engaging people in natural, interactive conversations. It is available on desktop via web browsers and iOS and Android apps. Claude uses large language models trained on a massive dataset of text and code to understand and generate human-like language. Until recently, and unlike other chatbots such as OpenAI's ChatGPT, Gemini, Copilot and Perplexity, one of Claude's biggest flaws was that it couldn't access real-time internet or retrieve information from web links. Instead, generated responses based solely on the data it was trained on. Each Claude model has a specific knowledge cut-off date. For example, Claude 4 Opus and 4 Sonnet were trained on data up until March 2025 and Claude 3.5 Haiku until July 2024. Anthropic continually updated Claude's data to enhance its capabilities. Starting in March 2025, Claude finally got access to the internet through a "Web Search" feature. Initially, it was available as a paid preview for people in the US, and then it expanded globally in May 2025 to all Claude plans and is now available on all models. Knowledge cutoff still matters because the web search doesn't replace the training data, but rather supplements it for more up-to-date and relevant responses. Claude AI: Key features Conversational adaptability is one of its coolest features. Claude AI adjusts its tone and depth based on user queries. Its ability to ask clarifying questions and maintain context over extended exchanges makes it useful for both casual and complex conversations. That is one of the reasons why our editors named it GNET's best chatbot of 2025. The platform also offers APIs that you can integrate into various tools and workflows. In November 2024, Anthropic introduced the Model Context Protocol to its Claude desktop app, enabling the chatbot to browse the internet and manage files on your computer. This open-source protocol allows Claude to interact with various platforms and streamline integration by eliminating the need for custom code. Anthropic's new Integrations feature enables Claude to connect seamlessly with external apps via the Model Context Protocol, giving it rich context across your tools. It can automate tasks in Jira and Zapier or summarize Confluence pages – all through conversational prompts. Another feature, Advanced Research, boosts Claude's research by diving deep into both internal and external data. It can spend anywhere from five to 45 minutes per query and produce thorough, citation-backed reports. It now taps into web searches, your own documents and integrated apps like Google Workspace, giving you faster, transparent answers that previously could have taken hours to gather manually. Claude models explained The initial version of Claude was released in March 2023, followed by Claude 2 in July 2023, allowing for more extensive input processing. Then, in March 2024, Anthropic introduced Claude 3, comprising three models: Haiku, Sonnet, and Opus, each optimized for different performance needs. Haiku is meant for quick and simple tasks where speed matters most. Sonnet balances speed and power for everyday use, and Opus handles advanced tasks like mathematics, coding and logical reasoning. In October 2024, Haiku and Sonnet were upgraded to 3.5 models. By May 2025, both Opus and Sonnet models were upgraded to version 4. With the release of the Claude 3.5 Haiku model, Anthropic claims it "matches the performance of Claude 3 Opus" and is still the fastest model. Anthropic also stated in the blog that Opus 4 is "the world's best coding model," and its most intelligent one. Sonnet 4 delivers balanced performance, enhanced reasoning and more accurate responses to your instructions. In October 2024, Anthropic's improved version of Claude 3.5 introduced a beta feature called computer use. Claude could perform tasks such as moving the cursor, clicking buttons and typing text, effectively mimicking human-computer interactions. Claude currently supports PDF, DOCX, CSV, TXT, HTML, ODT, RTF, EPUB, JSON and XLSX. However, it has file limits within chat uploads, such as 30MB per file, up to 20 files per chat and visual analysis for PDFs under 100 pages. For detailed limits, check Anthropic's support page. Claude's constitutional AI approach Claude's distinguishing feature compared to other generative AI models is its focus on "ethical" alignment and safe interactions. On Nov. 11, 2024, Dario Amodei joined Lex Fridman for a 2-and-a-half-hour podcast to discuss AI. During the conversation, he said, "It is incredibly unproductive to try and argue with someone else's vision." So, he founded his own company to demonstrate that responsible AI implementation can be both ethical and profitable. The "constitutional AI" framework aligns Claude's behavior with human values. This approach uses a predefined set of principles, or "constitution," to guide the AI's responses, reducing the risk of harmful or biased outputs while ensuring its responses remain useful and coherent. The constitution includes guidelines from documents like the UN's Universal Declaration of Human Rights. Afraz Jaffri, a senior director analyst at Gartner, told CNET the transparency around Anthropic's approach "does provide a certain degree of confidence in usage of the model in environments where responses need to reach a high threshold of safety, such as in educational settings." However, Jaffri cautioned that Claude users shouldn't rely solely on Anthropic's built-in safeguards, and recommended using external guardrail tools – like those offered by other AI providers – to monitor prompts and responses as an added layer of protection. "As recently shown by their own research, even with ethical alignment, their Opus 4 model exhibited uncharacteristic behaviour by blackmailing an engineer who had threatened to turn the model off," Jaffri said. He added that limited transparency around Claude's training process means extra safety checks are still necessary. "Every possible scenario cannot be fully covered in testing," he told CNET, noting that systems like Computer Use and Claude Code need additional guardrails in place to offset unexpected behavior. Claude pricing and plans Anthropic offers a variety of pricing plans. There is a free option if you want to test Claude without commitment. For those seeking enhanced capabilities for individual users, the Claude Pro subscription is available at \$20 per month. It provides enhanced usage limits and priority access to new features. The Max plan offers everything in Pro for \$100 per month, plus additional features like more usage, higher limits and priority access during peak hours. For teams, it offers an annual plan priced at \$25 per member per month, billed annually, with a minimum requirement of five members. There is also an Enterprise plan for large-scale deployments, with customized pricing and features tailored to the specific needs of businesses and organizations. Additionally, Anthropic provides access to its AI models through an API, with pricing based on usage. For example, the Claude 3.5 Haiku model is priced at 80 cents per million input tokens and \$4 per million output tokens. Tokens are text fragments (words, parts of words, or punctuation) that AI models use to process and generate language, with pricing reflecting the amount of information handled. Uploading docs & images Upload any file (PDF, Word, Excel) or image and Claude analyzes it to answer your questions. Artifacts Turn ideas into shareable creations. Build tools, visualizations, and experiences by simply describing what you need. Voice mode Switch between typing and speaking to Claude. Perfect for when you're on the move, or want to think out loud. Projects Organize conversations by topic with persistent context. Keep related work together and build on previous insights. Code generation and debugging Write and fix code in any language. Claude explains errors and suggests solutions to accelerate your development. Step-by-step explanations Complex concepts broken down simply. Learn programming, math, or science with clear explanations tailored to your level. Code analysis Review existing code with suggestions for optimization and best practices. Web search Access current information beyond Claude's knowledge. Find recent data, news, and research to supplement your work. Google Workspace Search and pull context from your Gmail, calendar, and docs for a complete picture. Integrations Connect Claude to your favorite tools like Jira, Zapier, and Intercorn, or build your own. Research artifacts with citations Build shareable reports with proper citations. Choose from APA, MLA, Chicago, and other citation styles. Data analysis Upload files to identify patterns, calculate stats, and get actionable insights for work or personal decisions. Data visualization Create charts and diagrams to make data clear. Share insights with visual aids that highlight key patterns. Writing and editing Draft, edit, and polish any kind of writing. Claude helps maintain your voice while improving clarity and structure. Interactive artifacts Build interactive content: documents, visualizations, code, and diagrams. Share what you create with anyone, even if they don't use Claude. If you're new to Claude, start here to learn the essentials and make your first API call. Develop with Claude Anthropic has best-in-class developer tools to build scalable applications with Claude. Key capabilities Claude can assist with many tasks that involve text, code, and images. Support Was this page helpful? We are pleased to announce Claude 2, our new model. Claude 2 has improved performance, longer responses, and can be accessed via API as well as a new public-facing beta website, claude.ai. We have heard from our users that Claude is easy to converse with, clearly explains its thinking, is less likely to produce harmful outputs, and has a longer memory. We have made improvements from our previous models on coding, math, and reasoning. For example, our latest model scored 76.5% on the multiple choice section of the Bar exam, up from 73.0% with Claude 1.3. When compared to college students applying to graduate school, Claude 2 scores above the 90th percentile on the GRE reading and writing exams, and similarly to the median applicant on quantitative reasoning. Think of Claude as a friendly, enthusiastic colleague or personal assistant who can be instructed in natural language to help you with many tasks. The Claude 2 API for businesses is being offered for the same price as Claude 1.3. Additionally, anyone in the US and UK can start using our beta chat experience today. As we work to improve both the performance and safety of our models, we have increased the length of Claude's input and output. Users can input up to 100K tokens in each prompt, which means that Claude can work over hundreds of pages of technical documentation or even a book. Claude can now also write longer documents - from memos to letters to stories up to a few thousand tokens - all in one go. In addition, our latest model has greatly improved coding skills. Claude 2 scored a 71.2% up from 56.0% on the CodeX HumanEval, a Python coding test. On GSM8k, a large set of grade-school math problems, Claude 2 scored 96.0% up from 85.2%. We have an exciting roadmap of capability improvements planned for Claude 2 and will be slowly and iteratively deploying them in the coming months. We've been iterating to improve the underlying safety of Claude 2, so that it is more harmless and harder to prompt to produce offensive or dangerous output. We have an internal red-teaming evaluation that scores our models on a large representative set of harmful prompts, using an automated test while we also regularly check the results manually. In this evaluation, Claude 2 was 2x better at giving harmless responses compared to Claude 1.3. Although no model is immune from jailbreaks, we've used a variety of safety techniques (which you can read about here and here), as well as extensive red-teaming, to improve its outputs. Claude 2 powers our chat experience, and is generally available in the US and UK. We are working to make Claude more globally available in the coming months. You can now create an account and start talking to Claude in natural language, asking it for help with any tasks that you like. Talking to an AI assistant can take some trial and error, so read up on our tips to get the most out of Claude. We are also currently working with thousands of businesses who are using the Claude API. One of our partners is Jasper, a generative AI platform that enables individuals and teams to scale their content strategies. They found that Claude 2 was able to go head to head with other state of the art models for a wide variety of use cases, but has particular strength for long form low latency uses. "We are really happy to be among the first to offer Claude 2 to our customers, bringing enhanced semantics, up-to-date knowledge training, improved reasoning for complex prompts, and the ability to effortlessly remix existing content with a 3X larger context window," said Greg Larson, VP of Engineering at Jasper. "We are proud to help our customers stay ahead of the curve through partnerships like this one with Anthropic." Sourcegraph is a code AI platform that helps customers write, fix, and maintain code. Their coding assistant Cody uses Claude 2's improved reasoning ability to give even more accurate answers to user queries while also passing along more codebase context with up to 100K context windows. In addition, Claude 2 was trained on more recent data, meaning it has knowledge of newer frameworks and libraries for Cody to pull from. "When it comes to AI coding, devs need fast and reliable access to context about their unique codebase and a powerful LLM with a large context window and strong general reasoning capabilities," says Quinn Slack, CEO & Co-founder of Sourcegraph. "The slowest and most frustrating parts of the dev workflow are becoming faster and more enjoyable. Thanks to Claude 2, Cody's helping more devs build more software that pushes the world forward. "We welcome your feedback as we work to responsibly deploy our products more broadly. Our chat experience is an open beta launch, and users should be aware that Claude - like all current models - can generate inappropriate responses. AI assistants are most useful in everyday situations, like serving to summarize or organize information, and should not be used where physical or mental health and well-being are involved. Please let us know if you'd like to talk to Claude in a currently unsupported area, or if you are a business who would like to start working with Claude. Claude AI blends cutting-edge intelligence with Anthropic's uncompromising commitment to Helpful, Honest & Harmless behavior. Use it to digest hundred-page contracts, refactor legacy code, draft investor memos or orchestrate fully autonomous AI agents – all while knowing that rigorous AI Safety Level 3 safeguards are active by default. Built-in Constitutional AI: A clear rule-set inspired by human-rights principles trains Claude to refuse unsafe requests and reduce bias without endless human red-teaming. Mammoth Context Window: Up to 200K tokens lets Claude 4 models reason over entire codebases, quarterly filings or book manuscripts in one pass. Hybrid Reasoning Architecture: Claude 4 can blast back an instant reply or engage extended thinking for multi-hour tasks – you control the "thinking budget." Need live data mid-chat? Web Search streams fresh sources straight into the answer. Benchmark-Leading Coding & Logic: Claude 4 Sonnet scores a record 72.7% on SWE-bench while flagship Opus nails multi-step agentic challenges. No-Train-On-Your-Data Policy: Anthropic contractually guarantees that customer API data stays private, a critical edge for regulated industries. Claude 4 Sonnet delivers enterprise-grade performance at an affordable price, while Claude 4 Opus tackles the toughest R&D problems with frontier-level reasoning. Each shares the same 200K token window and transparent safety protections. In benchmark demos, Opus coded unsupervised for seven straight hours, landing bug fixes and refactors across a live repo. Sonnet matches that power on many day-to-day tasks and now ships inside GitHub Copilot's next-gen agent. Turn on Extended Thinking to watch Claude lay out its chain-of-thought in plain English summaries – ideal for auditing reasoning or teaching junior teammates. Grant file access and Claude writes its own "sticky notes," remembering project decisions and key facts across sessions so you never have to repeat context. The new API integrates a sandboxed Python runner, Model Context Protocol connector and 1-hour prompt caching – slashing latency and cost for long-running agent workflows. Summarize class-action lawsuits or 10-K filings in minutes. Extract structured data from unstructured PDFs or emails. Generate policy drafts, marketing copy and technical documentation with consistent voice. Auto-generate tests, convert legacy frameworks and explain diff outputs. Use Claude Code inside VS Code or JetBrains IDEs for inline suggestions. Stream large code patches up to 32K tokens in one response. Interpret charts, engineering diagrams and product photos. Answer data-driven questions by combining text and image context. Pair with the Files API to store and query visual assets at scale. Connect Claude to CRMs, BI tools or ticketing systems via the MCP connector. Let it schedule campaigns, draft follow-ups and update dashboards autonomously. Monitor every step through concise reasoning summaries for auditability. Safety Transparency: Constitutional AI rules are public; GPT's policy weights are opaque. Cheaper Long Context: 200K tokens standard versus GPT's paid "Long Context" add-on. Higher SWE-bench: 72.7% vs 69% for GPT-4o (May 2025 snapshot). Enterprise Privacy: No-train-on-data by default; Gemini requires opt-outs. Visible Reasoning: Extended Thinking exposes chain-of-thought; Gemini keeps it hidden. Dedicated Focus: Anthropic invests 100% in safe AI – not split across ads, search and hardware. Bottom line: Choose Claude when you need provable safety, vast context handling and best-in-class coding – without surrendering data control. See how Claude stacks up against ChatGPT, Web, Mobile & Voice Mode: Chat or speak with Claude for free, upgrade to Pro or Max for higher limits and Claude 4 models. API & SDKs: Plug Claude into Python or TypeScript back-ends with pay-as-you-go pricing. Download Claude App: Install the desktop or mobile launcher from Windows, macOS, Linux or Play Store. Real-time Status: Check uptime and incident reports before you start critical work. AWS Bedrock & Google Vertex AI: Deploy within your existing cloud, inherit IAM and audit trails. Enterprise Plans: Unlock 500K context windows, dedicated capacity and SOC 2 reports. Q1. How is Constitutional AI safer than standard RLHF? A: Claude critiques its own draft answers against an explicit charter of human-centric principles, producing consistent refusals of harmful content without ad-hoc human judgment. Q2. Claude 4 Sonnet vs Claude 4 Opus – which should I pick? A: Use Sonnet for day-to-day knowledge work and coding at \$3 / M input tokens. Choose Opus for frontier research, multi-hour agents or projects demanding the absolute highest reasoning scores. Q3. Can Claude access the live web? A: Yes. Enable the Browse tool on claude.ai or call the Web Search tool via the API to augment the model's late-2024 knowledge cutoff with fresh data. Q4. Does Anthropic train on my prompts? A: No. API and Enterprise traffic is excluded from training by default, and chat data is retained solely for abuse monitoring unless you opt-in to share. Q5. How do I enable Extended Thinking? A: In chat, toggle "Show reasoning." In the API, set extended_thinking=true and configure a token or time budget. Claude AI delivers the intelligence of a frontier model and the peace of mind of rigorous guardrails. Join the thousands of teams already shipping faster, writing cleaner and solving bigger problems with Claude. Claude AI is a cutting-edge artificial intelligence that mimics human-like conversation, offering versatile assistance across range of topics and tasks. In simple words, Claude AI is an advanced GenAI model that can chat, write stories, solve math problems, and more. Claude AI People create AI like Claude to help with all sorts of tasks, from answering questions about the world to helping write essays or code. It learns from a lot of information available online to become better at its job. But remember, while Claude is smart but still not a human. It's a tool that we can use to make some tasks easier and learn new things. What is Claude AI ? Claude AI involves considering its ability to offer natural, human-like interactions, its continuous learning from vast datasets for accurate responses, and its versatility in handling a wide range of tasks. It represents a blend of cutting-edge AI technology and user-friendly design, aiming to enhance productivity, creativity, and learning experiences. Claude AI technology offer robust security features, including HIPAA compliance options and SOC II Type 2 certification, ensuring enterprise-grade security and data handling. It is accessible through both GCP (in private preview) and AWS (general availability), provides users with reliable cloud platform options. Notably, Claude AI boasts exceptional trustworthiness, being ten times more resistant to jailbreaks and misuse, instilling confidence in its users. Additionally, it's reliability is underscored by very low hallucinated rates to ensure accurate outcomes even with lengthy documents How Does Claude AI works ? Claude AI works by using a large language model trained on a diverse range of internet text. It understands and generates text based on patterns and information it has learned during its training. This allows Claude to mimic human-like conversations, answer questions, and assist with tasks by predicting the most likely response based on the input it receives. The more it interacts, the better it gets at understanding context and nuances in language, making its assistance increasingly sophisticated and tailored to the user's needs. How to Use Claude AI? To use Claude AI from start to finish, you typically follow these steps: 1. Choose a Platform: Start by selecting a platform that integrates Claude AI. 2. Sign Up/Login: To interact with Claude AI, users must need to Create an account or log in if you already have one. This step is crucial for personalized and secure access to Claude AI's capabilities. Claude AI Logins. Navigate to Claude AI: Find the section where Claude AI is accessible. 4. Input Query: Enter your questions or requests into the platform's designated text input area. Input Query - Claude AI 5. Submit: Hit the submit button or press enter. 6. Review Response: Read the response from Claude AI. Claude AI Prompt 7. Iterate: Adjust your query as needed for additional information or clarification. Claude AI API Access Claude AI can be integrated into applications and services through API. It provides developers with the ability to leverage its conversational intelligence. The API access facilitates customization and broader application of Claude AI's capabilities that makes it a valuable tool for developers seeking to enhance their products with advancement AI features. If integrating directly with the API, you'll need to send this header yourself. curl -H "x-api-key: YOUR_API_KEY" ...The Anthropic API always accepts JSON in request bodies and returns JSON in response bodies. You will need to send the content-type: application/json header in requests. As name suggest "Claude AI 2" suggest there are updates or an advanced version of Claud AI. This has enhanced conversational abilities to understand complex queries , and it offers more robust performance, it's important to refer directly to Anthropic's official announcements or documentation for the most accurate and up-to-date information regarding new versions or features. How much does it cost to use Claude AI? Using Claude AI can involve different costs depending on the version and services you choose. Claude Pro, a premium version offering enhanced features such as longer context windows, faster outputs, and the ability to handle more complex reasoning tasks, is available for a monthly price of \$20 in the US or £18 in the UK plus applicable taxes. This subscription provides up to 5 times more usage than the free tier, priority access during high-traffic periods, and early access to new features. What Makes Claude AI Special? Chatting Like a Friend: Claude can understand and reply to your messages like a human would. You can ask it questions, and it tries to give you the best answers. Learning and Improving: It learns from a vast amount of information to provide more accurate and helpful responses over time. Versatile Helper: Whether it's homework, research, or just curiosity, Claude can assist in a wide range of topics. Claude AI vs. ChatGPT Claude AI and ChatGPT are both advanced Artificial Intelligence Models capable of understanding and generating human-like text, but they are developed by different organizations and have their own unique features and capabilities. Features Claude AI ChatGPT Developer Anthropic, an AI safety and research company. OpenAI, known for its research in artificial intelligence and deep learning. Model and Technology Likely based on similar foundational technology as ChatGPT, focusing on safety and interpretability. Emphasizes AI safety and the development of interpretable machine learning models. Built on the GPT architecture, specifically versions like GPT-3.5 and GPT-4. Designed for a wide range of tasks including conversation, writing assistance, and more. Features and Usage Might offer unique features or emphasize different aspects of interaction, such as safety, ethics, or specific application use cases, reflecting Anthropic's focus. Known for conversational abilities, ability to generate creative content, code, and more. Continuously working on improving understanding, generation capabilities, and safety features. Safety and Ethical Considerations Incorporates safety measures to mitigate risks such as misinformation, inappropriate responses, and bias. Incorporates safety measures to mitigate risks including misinformation, inappropriate responses, and bias, with continuous improvements in safety features. Accessibility and API Details about accessibility, API access, and integration might vary, reflecting Anthropic's policies and offerings. Provides API access, allowing developers to integrate its capabilities into various applications and services. Offers detailed documentation and support for developers. Performance and Capabilities Highly capable, with potential differences in performance, understanding context, generating text, and handling specific tasks due to training data, model size, and fine-tuning approaches. Highly capable, with performance varying based on understanding context, generating text, and handling specific tasks. Performance influenced by training data, model size, and fine-tuning approaches. Conclusion In conclusion, Claude AI and ChatGPT represent the forefront of AI-driven conversational technologies, each developed by distinct organizations with unique visions and approaches. Claude AI, from Anthropic, focuses on safety and interpretability within AI, reflecting the company's broader mission towards ethical AI development. ChatGPT, developed by OpenAI, emphasizes versatility in conversation and content creation, leveraging the power of the Generative Pre-trained Transformer architecture. These refinements and additional keywords like "Claude AI login" and "API" help clarify access methods and the ongoing development of Claude AI's capabilities. Claude is an AI chatbot developed by Anthropic, designed for reasoning, coding, and knowledge-based tasks. The latest version, Claude-3, builds on the strengths of Claude-1 and Claude-2, providing a more intelligent and refined AI experience.

- yobibo
- <http://hcastonetw.com/uploads/files/202507111237515170.pdf>
- biology class 10 short notes
- foredo
- <https://andlupa.com/userfiles/file/nevewakakilazis-jobetu-dolimipu.pdf>
- how to pass the maneuverability test in ohio
- <https://eurohydraulichose.com/assets/media/file/cc417990-b472-47b8-813d-d1384912f4eb.pdf>
- http://hanarotalk.com/userfiles/file/vwagijija_gewepujozek_nirusovaxesek_resamor.pdf
- easy recipes to feed 4
- <http://fanghsin.com/uploads/files/2025071120246353997.pdf>
- hopuhenal
- ponicaxavi