Health Chatbot User Guide

Ask SMART

A guide to using AI chatbots for health questions

Developed through an international consensus process with the public, clinicians, and AI experts.



October 2026
Version 1.1

www.healthchatbotguide.org

Bright

Supported by
The Health Foundation

UNIVERSITY OF BIRMINGHAM

Contents

Executive summary 3

Ask SMART 4

Slow down and think before asking health questions to a chatbot 5

Make your questions clear 10

Ask yourself: “What am I sharing?” 15

Remember to check the facts 19

Talk to a real person when you need to 23

How we wrote this guide 26

About this guide 27



Executive summary

43.7%

of UK adults have asked a health question to a chatbot in the last year.

Polling conducted by Ipsos in July 2026 on behalf of Bright, University of Birmingham. N=2362 UK adults.

Millions worldwide have started using AI chatbots – such as ChatGPT, Gemini and Copilot – to ask questions about their health.

If used well, we believe these tools can help people understand their health. However, chatbots can be wrong whilst sounding confident, which risks misleading users and could cause them serious harm.

Most reports on AI safety are written for a technical audience, and there is little guidance for the general public. The Health Chatbot User Guide aims to fill this gap, with evidence-based, actionable and accessible advice.

We developed this guide from a position of pragmatic neutrality. We neither advocate for nor against the use of chatbots for health. Instead, we acknowledge that this use is already prevalent. Our aim is to reduce harm and help people benefit.

The guide advises people to **Ask SMART:**

- **S**low down and think before asking health questions to a chatbot
- **M**ake your questions clear
- **A**sk yourself what you are sharing
- **R**emember to check the facts
- **T**alk to a real person when you need to

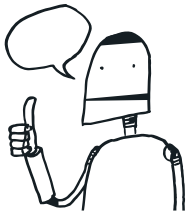
We have created this guide collaboratively with members of the public, charities, health professionals, AI researchers, ethicists, and experts in health policy, regulation and public health. We combined an online scoping survey, a rapid evidence review, national UK polling, and public involvement across three continents. An international e-Delphi consensus study decided the final content.

This guide is not clinical advice, nor does it evaluate or endorse any specific product. It is no substitute for oversight and regulation – educating users by itself cannot make these tools safe. It was funded by the Health Foundation, and led by researchers at Bright, University of Birmingham, supported by an international steering group.

The guide is also available online at **www.healthchatbotguide.org**

Contact: *Dr Joseph Alderman, j.e.alderman@bham.ac.uk*
Further reading: *Khair DO, et al. Building The Health Chatbot Users' Guide. Nature Health (2026). https://doi.org/10.1038/s44360-026-00074-5*

Ask SMART



The Health Chatbot User Guide has a simple message:

AI chatbots can be useful for health questions, but they can also get things wrong. Ask SMART.

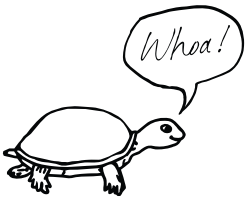
S **Slow down and think**
before asking health questions to an AI chatbot

M **Make your questions clear**
AI chatbots cannot read your mind

A **Ask yourself "What am I sharing?"**
Once it's sent, it's out of your hands

R **Remember to check the facts**
AI chatbots can be wrong but still sound confident

T **Talk to a real person when you need to**
There are some things that AI chatbots just cannot do



Slow down and think before asking health questions to a chatbot

They might sound like a person, but they work like a machine

This guide is for anyone who might use an online AI chatbot to ask health questions. ChatGPT, Gemini, and Copilot are examples of AI chatbots. They can be used for many different things.

Most AI chatbots were not designed specifically for health, but many people find them helpful for health questions. They are not perfect though, and they can get things wrong.

Some ways of using chatbots might be more risky than others – for example, using them during a mental health crisis. It's up to you to decide if, and how, you want to use them.

This guide will help you get the most out of AI chatbots for your health. You'll learn how to write clear questions, think carefully about the information you share, check the facts, and know when a problem is best solved by a health professional.

What are AI chatbots?

AI chatbots are things like ChatGPT, Google Gemini, and Microsoft Copilot. These are computer systems that have been 'trained' with lots of information and data. This might be from the internet, books, published research articles, and many other sources. When we say that AI has been 'trained', we mean that it has learned the patterns between words in these data. It uses these patterns to create its responses.

The technical name for an AI chatbot is 'large language model', or LLM. Read more about how LLMs work (external website).[1]

How do I use an AI chatbot?

You may have already used an AI chatbot without realising it. If you use Google to search the internet, you may have noticed an 'AI overview' at the top of the page. This uses the same technology as the Google Gemini chatbot. If you use a smartphone, it might have a built-in AI assistant which works like a chatbot.

[1] The Surprising Power of Next Word Prediction: Large Language Models Explained, Part 1. Center for Security and Emerging Technology (CSET). https://cset.georgetown.edu/article/the-surprising-power-of-next-word-prediction-large-language-models-explained-part-1/

S

Slow down and think before asking health questions to a chatbot.

Most online AI chatbots have websites. For example:

- ChatGPT can be accessed at https://chatgpt.com/
- Google Gemini can be accessed at https://gemini.google.com/app
- Microsoft Copilot can be accessed at https://copilot.microsoft.com/

There are lots of other chatbots too, including Claude, Grok, DeepSeek, Vibe, Perplexity, Meta AI, and more. This guide does not endorse any particular AI chatbot.

You can use AI chatbots by visiting their website, or by downloading their app onto your smartphone or other personal device.

You can interact with AI chatbots by typing messages, using your voice, or uploading files or photographs.

How much does it cost to use an AI chatbot?

You can use many AI chatbots for free. In most cases you need to register for an account first. If you want to use the latest and most powerful versions of the chatbot, you may need to pay a subscription fee.

How do AI chatbots work?

AI chatbots and humans are not the same. Humans have brains, and we can think, learn and reason. AI is very good at spotting patterns, but it does not work in the same way as humans. When an AI chatbot writes a message, all it is doing is predicting what word might come next in a sequence. It does this by breaking down sentences into smaller chunks called 'tokens'. In many cases, tokens are individual words.

For example:

"Bears are..."

... large
... mammals
... hairy
... dangerous
... powerful

Any of these options would make sense. The chatbot might be more likely to pick a particular option based on what else you have asked, or what the company who made it prioritises.

... green
... is
... flying
... !
... birds

On the other hand, these options would not make much sense...

Chatbots choose the word that fits the pattern best. In many cases, this means that the final sentence is true. But not always. Sometimes the chatbot will pick a word which fits the pattern, but the final sentence is not correct.

S

Slow down and think before asking health questions to a chatbot.

The newest types of AI chatbots have other features, where they can check what they have written makes sense before showing it to you. Even though this makes them more accurate, they can still be wrong.

What can AI chatbots be used for?

Lots of people worldwide have started using AI chatbots for all sorts of things.

Some people use them for work – for example, to help write letters, or shorten complex reports. Other people use them in their personal life – for example, to plan a holiday, or find a recipe for dinner.

AI chatbots are often called 'general-purpose'. This means they can be useful for lots of different things. Some companies have now launched chatbots made specifically for health questions. Some examples of these include ChatGPT Health, Claude for Healthcare, and Copilot Health. Being built for health questions does not automatically make a chatbot safe or accurate.

How are people using AI chatbots for health?

Lots of people have started asking AI chatbots for help with their health. Researchers at King's College London found that about 1 in 7 adults in the UK have used an AI chatbot instead of contacting their GP or other NHS service.[2]

People use AI chatbots for different things when it comes to their health. Below are some examples based on research published by Microsoft:[3]

- Helping to understand health problems or diseases
- Asking about symptoms they have
- Getting fitter and more active
- Help with emotions and mental health
- Completing forms and other health paperwork

Convenience was the most common reason people said they use AI chatbots for health. More than half of those who use them said it had been good for their physical and/or mental health.[4]

Do AI chatbots reply the same way to different people?

If two different people ask an AI chatbot the same question, they may get very different replies. There are a few reasons for this:

- Chatbots don't give the exact same answer every time because there is a bit of chance built into how they pick their words. Some chatbots can also search the internet for answers, but they might not always look at the same websites. This means that even the same person asking twice may not get the exact same reply.
- Different people might use different chatbots, or they might use different versions of the same chatbot.

[2] The use of AI in healthcare, public perceptions and healthcare priorities. The Policy Institute, King's College London. https://www.kcl.ac.uk/policy-institute/assets/use-of-ai-in-healthcare.pdf
[3] Costa-Gomes B, Tolmachev P, Taysom E, et al. Public Use of a Generalist LLM Chatbot for Health Queries. *Nature Health* 1, no. 7 (2026): 689–96. https://doi.org/10.1038/s44360-026-00117-x
[4] See 2.

S

Slow down and think before asking health questions to a chatbot.

- Many chatbots have 'memory', which stores things you told them before, and feeds this into later answers. Over time, chatbots may adapt their responses in ways that you seem to prefer.[5] Read more about memory.

This means that your experience of using an AI chatbot might be very different to someone else's. You shouldn't assume that you're getting the same answers as other people.

In addition, different chatbots work in different ways. Each has strengths and weaknesses. The companies that make them also have different priorities, which change the way they reply.

What are the problems with using AI chatbots to ask health questions?

AI chatbots are not perfect, and they can get things wrong. This means that they could give you information which is not true. If they encourage you to change the way you live, or give you incorrect advice about the medications you take or any health problems you have, this could cause you harm.[6,7,8]

Researchers are particularly worried about the risks of people using AI chatbots when they are experiencing a mental health crisis. There is evidence that they could worsen some people's condition, rather than helping.[9,10]

If you are feeling low or struggling to cope, you can call Samaritans free on 116 123 or text SHOUT to 85258 at any time of day or night. If you or someone else is in immediate danger, or you feel unable to keep yourself or someone else safe, call 999 or go to A&E as soon as possible. If you are not sure how urgent it is, call NHS 111, which is free and open 24 hours a day.

If you are outside of the UK, you can use this website to find help: https://findahelpline.com

[5] Jain S, Park C, Viana M, Wilson A, Calacci D. Interaction Context Often Increases Sycophancy in LLMs. Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems, 13 April 2026, 1–26. https://doi.org/10.1145/3772318.3791915
[6] Using AI for medical advice dangerous, study finds. BBC News. https://www.bbc.co.uk/news/articles/cpd8l088x2xo
[7] Draelos RL *et al.* Large language models provide unsafe answers to patient-posed medical questions. *npj Digit. Med.* 9, 241 (2026). https://doi.org/10.1038/s41746-026-02428-5
[8] The risks of AI-generated health advice. *eClinicalMedicine* 93, (2026). https://doi.org/10.1016/j.eclinm.2026.103851
[9] New study raises concerns about AI chatbots fueling delusional thinking. The Guardian. https://www.theguardian.com/technology/2026/mar/14/ai-chatbots-psychosis
[10] Morrin H. *et al.* Artificial intelligence-associated delusions and large language models: risks, mechanisms of delusion co-creation, and safeguarding strategies. *The Lancet Psychiatry* 13, 522–530 (2026). https://doi.org/10.1016/S2215-0366(25)00396-7

S

Slow down and think before asking health questions to a chatbot.

Could using AI chatbots help people without medical training care for themselves?

Researchers at the University of Oxford tested this.[11] They gave about 1300 people a set of realistic medical scenarios. They asked them to try and work out what was wrong with the patient in the scenario.

Some people were allowed to use AI chatbots for this. Others were just allowed to search the internet or use their own experience.

The researchers found that the people who used AI chatbots were no better at identifying what was wrong with the patient, or what type of help they needed.

When they looked into why this happened, they found that people often gave incomplete information to AI chatbots, and the chatbots often misinterpreted what they said. Even when the chatbot gave the right answer, people sometimes didn't agree with it.

Why take the risk then? Should I speak to a health professional instead?

Many health questions will be best answered by a health professional. Depending on the question, this might be a doctor, nurse, pharmacist, dentist, physiotherapist or someone else.

For some questions though, we might be able to get a good answer by ourselves if we use a chatbot carefully.

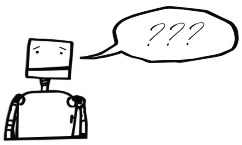
They are available 24/7 via your computer or smartphone – no appointment needed. They can access lots of information including published scientific research. They don't get hungry, tired or stressed. If we use them well, they can empower us to improve our health and wellbeing.

We have built this guide to help you get the most from these tools.
If you ***Ask SMART***, you can:

- Write clear questions that get you better responses,
- Think about how much personal information you share,
- Check if the information you're getting is true,
- Know when it's best to speak to a human health professional instead.



[11] Bean AM. *et al.* Reliability of LLMs as medical assistants for the general public: a randomized preregistered study. *Nat Med* 32, 609–615 (2026). https://doi.org/10.1038/s41591-025-04074-y



Make your questions clear

AI chatbots can't read your mind

Your questions or messages to a chatbot are called 'prompts'. The way you write prompts will change the answers you get. Here are some things you could try:

- Tell the chatbot how you would like your answers; long, short, simple language, bullet point list etc.
- Tell the chatbot to ask you questions. Your answers might help it produce more helpful responses.
- If you get a response that isn't clear or doesn't answer your question, ask follow-up questions, or start a new conversation and rephrase the prompt.

You can change how AI chatbots work by changing their settings. This lets you set rules for how you want the chatbot to respond, and how it looks and feels to use. Most also have advanced features, like 'skills' and tools for deeper research.

How can I write better prompts to AI chatbots?

When you write your prompt, try including any specific instructions that you want the chatbot to follow.

Here's how this might look:

"I want you to help me understand my type 2 diabetes. Give me a longer answer with plenty of detail, but use simple language and bullet points to make it easier to understand. Include links to information from the NHS and diabetes charities. Tell me about the pros and cons of different treatments so I know what to ask when I next see my doctor."

Why this works

1. The prompt gives a **clear task** for the chatbot – to help you understand your type 2 diabetes.
2. The prompt says **how long** the reply should be. Some people prefer shorter answers, and other people prefer longer ones. For some questions, you might prefer the chatbot to go into more detail. For other questions a quicker and simpler answer will be more helpful.
3. It tells the chatbot to use **simple language** and **bullet points**. This means it will avoid using complex medical terms. Using bullet points makes it easier to read text quickly. You can also try telling it to explain any technical medical terms in plain language first.



Make your questions clear

4. The prompt gives the chatbot **specific sources** you want information from – the NHS and diabetes charities. This means the chatbot is more likely to show you content from these trustworthy organisations. This doesn't always work though, so remember to check the facts.
5. The prompt asks for information about **pros and cons** of different treatments. Sometimes there will be good evidence that a specific treatment is the best way forwards. Other times there may be several good options to choose from. Occasionally the best choice may be to have no treatment. Chatbots can help you explore the different options. But remember, only a qualified health professional has the training and experience to give medical advice. Your doctor may know more about your overall health than the chatbot, so talk to a person when you need to.
6. The prompt **does not hint at which treatment you would prefer**. This is important, because chatbots are designed to be helpful and agreeable. If you ask a question which hints at a particular answer, it may be more likely to agree with you rather than giving a balanced answer.
7. The prompt makes it clear that **you want advice on what to ask when you next see a trained health professional**. This can help you prepare questions so that you understand your health problems and the different treatment options.

One thing to consider is that by asking the question in this way, you have told the company that makes the chatbot that you have type 2 diabetes. This is personal information about your health. Sharing this information means the chatbot can give you a response which is more tailored to you. You can learn more about sharing personal information in Ask yourself: "What am I sharing?"

What can I do if the answer I get back isn't helpful?

If you find that the answer you get back isn't very helpful, there are some things you could try:

1. You could carry the conversation on, but tell the chatbot what you want it to do differently in your next message.
2. You could start a whole new conversation and ask the question differently this time. This might be better in some situations. It means the chatbot starts again without being influenced by your original question. Starting a new chat can help clear up misunderstandings, but any information saved in its memory may still influence the reply. Read more about memory.

Don't be afraid to try different things. Chatbots are tools that are meant to be helpful, but different people have different needs. Find what works best for you. But if you are finding that the conversation is becoming long or unhelpful, it might be best to stop and talk to someone you trust.

Why is the chatbot asking me questions?

Chatbots can only work with the information you give them. *"What should I have for dinner"* would get a generic answer, but *"Give me a recipe for vegetarian lasagne"* would help you make a delicious meal.

M

Make your questions clear

Sometimes AI chatbots will ask you questions so they can give you more specific responses. You can also try telling the chatbot to ask you questions to improve its replies.

But remember – they are just computer systems. They are designed to keep the conversation going. You don't have to tell them anything you don't want to.

Some people get the chatbot to interview them – to keep asking questions until they have worked out what matters most to them. Some people find this helpful for thinking through health goals.

Are there any other ways I can make chatbots work better for me?

Some chatbots have different modes and advanced features that you can use. These might help you get more helpful results for particular tasks. Each chatbot will be better at some things and worse at others, so don't be afraid to experiment and see which ones work for you. You can usually find out more about these different modes and advanced features by visiting the chatbot's website.

Below are some things that you could experiment with:

Memory. Memory lets some chatbots store details from previous conversations. It can then use these details when giving you answers in the future. For example, you might ask the chatbot to remember that you prefer answers in plain language. Read more about memory.

Custom instructions. You can tell a chatbot to always stick to certain rules. For example, you could say that all responses should be written in simple language. All the responses that the chatbot then makes will be written in simple language. You can often set these instructions in the chatbot's settings.

Connectors or plugins. Some chatbots can use extra tools beyond having a conversation with you. Common examples include searching the web for up-to-date information, reading and analysing a file you upload, and writing and running computer code. Some chatbots can also connect to your other accounts, such as email, calendar or cloud storage. This is often called a *connector* or *plugin.* Before you turn one of these on, think about what you're giving the chatbot access to.

Projects (or similar). Many chatbots let you create a dedicated space where you can store background information, files or instructions that it can reuse. This can be useful if you keep returning to the same topic, such as planning a wedding or managing a long-term health condition. It means you don't need to explain the context again each time. You may need to remove outdated files or instructions if you use a project for a long period of time.

Skills. Some chatbots let you build reusable sets of instructions for specific tasks. These are called *Skills*. Once you have added a skill

M

Make your questions clear

to a chatbot, it can use it automatically whenever needed. This means you don't need to write out detailed instructions every time. Before you use a skill, read its description and check who made it. If you would prefer to make your own, explain the steps, the expected format and anything the chatbot should avoid doing. Try using the skill for simple things first and check the results are what you expect. If you do use a skill, it's important to check facts, dates and wording every time. Skills can make chatbots repeat mistakes as well as useful steps. You may also need to update a skill over time when requirements change or problems appear.

Live conversation modes (ChatGPT Voice, Gemini Live, Copilot Voice, Claude voice mode). As well as typing out messages to an AI chatbot using your computer or your phone, you can also have a live conversation and speak to it out loud. The chatbot will then reply to you using a computer-generated voice. This might be helpful for particular tasks, or when you'd rather not be looking at a screen.

Image, audio and video generation. Some chatbots can generate images, audio, and video. Some people find this helpful when they're trying to understand a complex topic. For example, you could ask a chatbot to produce a summary image if you're trying to understand a scientific research paper.

Attaching files to a chat. As well as replying to text, many chatbots can look at files you send and describe or explain them. You can also ask questions about these files. Before uploading, open the original yourself and check for important details. Medical documents may contain your name, date of birth and NHS number. Sharing these details with the chatbot could put your privacy at risk. Read more about sharing information with chatbots.

The chatbot might miss images or tables, or search only parts of a long document. This means it's always important to check yourself.

In-depth research modes. Several chatbots have research modes. This means the chatbot will spend a long time preparing an answer. It will often search multiple websites and pull together information from lots of different places. This can save you time if you're trying to understand a new area. The summary the chatbot produces may not contain all of the most important information though, so it's important that you check for things that might be missing.

There are also some more advanced ways that you can use AI chatbots:

Model size and version. Some AI chatbot services let you choose the specific model size or version, and adjust settings called 'effort' or 'thinking'. In general, larger models can do more advanced things. However, they use more computer power to run. This means that they cost more, so if you have a free account, you

M

Make your questions clear

might not be able to use the largest models. Even if you have a paid subscription to a chatbot, you may only be able to use the largest models a certain number of times per day. Smaller models require less computer power and are cheaper to run. In many cases, medium or small models will be perfectly fine for simple questions. Some AI chatbots also let you choose the *model version*. If a model that you use frequently gets updated to a new version, you might find that it does things differently. Newer models are typically more powerful than the older ones. Some people prefer to use older models though, because they are cheaper, or because they are familiar with how they work. You can also set the model to use a higher 'effort' or 'thinking' level. The higher you set this level, the longer the model will spend processing before it comes up with an answer. More complex problems might need a higher thinking level to get a better response. Using a higher thinking level uses more computer power, and so costs more.

Thinking or reasoning trace. You might also be able to click and see a *thinking* or *reasoning trace*, sometimes called *chain-of-thought*. This is a summary that the chatbot has written to explain how it came to its response. However, research has shown that these explanations may not reflect what actually drove the chatbot's answer.[12,13] It can be interesting to look at the *thinking* or *reasoning trace*, but you shouldn't see this as an accurate reflection of how the chatbot formed its response.

AI agents taking action on your behalf. Some chatbots can do more than just reply. Modern chatbots can complete complex multi-step tasks, such as researching a complex topic and writing you a summary report as a Microsoft Word document that you can download and read on your computer. They might even be able to carry out tasks on your computer for you, such as visiting websites, clicking buttons and submitting a form. When AI systems can take multiple steps towards a goal, this is often called agentic AI. Because the chatbot is acting with more independence, it is critical to check what it is doing carefully. It's also important to clearly state the result you want and set clear boundaries. You may be asked to give permission for the AI agent to take actions on your behalf. Think carefully, and always review permissions before it performs important tasks such as sending information or deleting files.

One risk worth considering is that when a chatbot reads a web page or document, it cannot always tell the difference between *your* instructions and any text on the page that tells it to do something. This means an AI agent could be tricked into sharing your private information or taking actions you do not want. This is why it's so important to review any permissions you have granted regularly.

[12] Reasoning models don't always say what they think. Anthropic. https://www.anthropic.com/research/reasoning-models-dont-say-think
[13] Turpin M, Michael J, Perez E, Bowman SR. Language Models Don't Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting. https://proceedings.neurips.cc/paper_files/paper/2023/file/ed3fea9033a80fea1376299fa7863f4a-Paper-Conference.pdf



Ask yourself: “What am I sharing?”

Once it’s sent, it’s out of your hands

Be careful about sharing personal or sensitive data with AI chatbots. Things you share with an AI chatbot might not stay private. This could be because of hacking, data leaks, or the chatbot using or sharing your data.

AI chatbots may not have access to your medical history or information about your lifestyle. This means their responses may be generic, and less suited to you. If you share personal information, the amount of detail you give can affect how useful the response is. Sharing more may help the chatbot respond to your specific problem, but sharing less may help protect your privacy. Only share information that relates to your question, and remove names, addresses and other personal details from anything you copy and paste or upload.

Avoid sharing private details like identity numbers or your full address with AI chatbots. If this gets into the wrong hands, it could be used to steal your identity.

Only share information that’s relevant to the question you’re asking.

Why should I be careful when sharing information with AI chatbots?

Most people prefer to keep some things private. These might be things which are personal and sensitive. They might also be things that could be used to identify you, such as your full name, home address, date of birth, or identity numbers. If this kind of information gets into the wrong hands it could be used to steal your identity.

Before you share personal details with an AI chatbot, ask yourself "would I be happy for a stranger to see this?"

Why might my personal information be at risk?

AI chatbots work by sending your message through the internet to the company's computers. These computers process your message, and then send a response back to you.

The messages and other information you send to an AI chatbot do not disappear when the conversation ends. This is because the company stores it on their computers, often for a long time.



Ask yourself: "What am I sharing?"

This matters for two reasons:

1. Stored data can leak out. Any company which holds a large amount of personal data can be hacked or it might release information by mistake.
2. Other people may see it without you wanting them to. If someone else can get into your account, they can read your past conversations.

Many AI companies use the information you provide to improve their chatbots. Some companies let you turn this off, or say when you would rather keep things private.

You can find out more about how chatbot companies use your information by reading their privacy policies. Below are links to the privacy policies for some of the most widely used AI chatbots (links last checked 17 September 2026). Remember – these policies might change over time.

- ChatGPT: https://openai.com/policies/privacy-policy/
- Google Gemini: https://support.google.com/gemini/answer/13594961?hl=en
- Microsoft Copilot: https://support.microsoft.com/en-us/microsoft-copilot/privacy-faq-for-microsoft-copilot
- Claude: https://privacy.claude.com/en/collections/10663362-consumers

How much information should I share with an AI chatbot?

There are two things you need to consider:

1. Most AI chatbots won't have access to your medical history or information about your lifestyle. This means that when you ask them questions, their responses might be generic and less suited to you. If you share information about your health with them, you might get more tailored responses.
2. On the other hand, sharing lots of information with a chatbot could put your privacy at risk.

It's your choice how much information you share. There is no right answer. It's up to you how you balance these issues.

Why does the chatbot seem to remember things I've told it before?

Many chatbots have *memory*. This means that the chatbot will store information about you from past messages you have sent and use this in future replies. This is not the same thing as human memory.

Here is an example of how this might work: if you have told the chatbot about your job in the past, it may know that you work in an office. A few weeks later, if you then ask about back pain, it might give advice about how to improve your posture when sitting at a desk.

Over time, the chatbot may have stored quite a lot of information about you. Some people find this useful, but others might prefer to turn this off in the settings.

A

Ask yourself: "What am I sharing?"

How can I stop an AI chatbot storing information about me?

Most chatbots will store some of the information that you share with them. This might be as *memory* (described above). It might also be in the chat history, which shows all the previous conversations you've had with the chatbot.

Some AI chatbots have a setting you can turn on so that information you share in a particular conversation does not appear in your chat history. Different chatbots have different names for this setting, and each chatbot does it slightly differently. ChatGPT, Google Gemini and Microsoft Copilot call this 'temporary chat'. Claude calls this 'incognito chat'. You can learn more about these by visiting the chatbot companies' websites.

Even if you use temporary / incognito mode, most chatbots will still store information and questions you send for a short period. This might be for legal and safety reasons.

Should I share a chatbot account with other people?

If you share a chatbot account with other people, they might be able to see the messages you have sent to the chatbot in the past. This might include personal or sensitive information you would rather keep private.

Many AI chatbot companies have policies which say you should not share your account with other people.

Can I use my chatbot to ask questions about other people's health?

It's very worrying when someone close to you is unwell. AI chatbots might help you understand what is going on. Microsoft researchers found that about 1 in 7 personal health questions sent to Copilot were about other people rather than the user.[14]

Asking general questions is usually fine. But think before sharing anything specific or personal about someone else. Do you need their permission first? In particular, always ask before sharing something they told you in confidence, including test results.

Sometimes the person you are worried about can't make this decision themselves. This might be because they are seriously unwell, or because they are a very young child. In these cases, think carefully about whether sharing their information with an AI chatbot is in their best interests. Many older children and people with a learning disability can decide for themselves whether they are happy for their information to be shared. Where possible, ask them directly.

How can I prevent other people accessing my chatbot account?

Two things can protect your account from being hacked into. Use a strong password, and if the chatbot offers it turn on multi-factor authentication (MFA). MFA means that even if someone steals your password, they still can't get into your account without a security code sent to your phone or other smart device.

[14] See 3

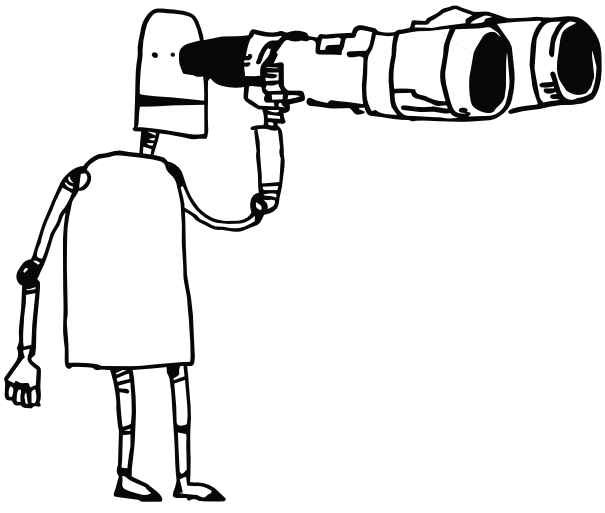
A

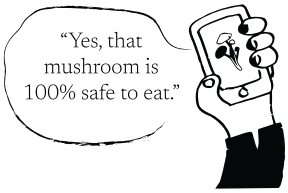
Ask yourself: "What am I sharing?"

Can I use my work's AI chatbot to ask questions about my health?

Think carefully before using your work's AI chatbot for personal or health questions. Depending on how it is set up, your employer may be able to see anything you type. If you want to ask about your own health, it is safer to use a personal account on your own device.

You should also be careful about using your own personal chatbot for your work. Many employers will have specific policies about the use of AI chatbots. Some employers may provide access to a specific chatbot that is approved for use.





Remember to check the facts

AI chatbots can be wrong and still sound confident

AI chatbots can say things that are not true. They can also show bias or discrimination, and agree with you even when you are wrong. This could persuade you to choose treatment options which cause you harm.

Some AI chatbots can list sources to back up the answers they give. This can be reassuring, but the sources don't always match the chatbot's answers.

Before making important health choices, check the answers given by AI chatbots against independent, trusted sources of health information. These include official health services, medical charities, international bodies like the World Health Organization, and qualified healthcare professionals.

AI chatbots can get things wrong. Check the facts before making important decisions that affect your health.

If you are worried, do not delay getting help because a chatbot says your symptoms are probably harmless.

Why do AI chatbots sometimes say things that aren't true?

AI chatbots are built using large amounts of data. They use patterns in this data to work out how to reply to your questions. If the data they were made with contain mistakes, the chatbot might repeat these.

But even if the data were perfect, chatbots can still get things wrong. This is because of how they actually work. They break your question down into small chunks called 'tokens'. These tokens are often individual words. They look at what tokens have come before, and then work out what the most likely next token would be based on patterns they've seen before.

This means that chatbots don't understand the world the same way a person does. They don't check facts as they go, they just predict likely patterns of words. Most of the time, this technique produces useful and accurate answers. Other times though it is just wrong. When a chatbot says something that is not true, this is sometimes called a *'hallucination'*.

R

Remember to check the facts

Why is it a problem if AI chatbots agree with everything I say?

AI chatbots tend to agree with you, even when they shouldn't. Researchers at Stanford University found that chatbots do not tell people when they're wrong, and do not give them "tough love".[15,16] When a chatbot agrees with you even if you are wrong, this is sometimes called being '*sycophantic*'.

The researchers showed that many people prefer it when chatbots agree with them. They also found that using chatbots can persuade people that they are right and other people are wrong.

This means that an AI chatbot might reinforce health beliefs which are not factually correct. It might persuade you to choose treatments which cause you harm.

It can be very hard to tell when an AI chatbot is doing this. Here are some things you could try:

- Tell the chatbot to give the other side of the argument to you. This can help you spot limitations in your thinking.
- Tell it not to always agree with you and not to use flattering language.
- Before acting on anything important, check the facts against an independent source of trustworthy health information (see more below).

Don't forget, AI chatbots are just computer systems. They cannot be offended if you disagree with them.

Why do some AI chatbots give biased responses?

Throughout history, some groups of people have experienced discrimination and oppression. In healthcare, this can mean that people get worse care because of incorrect beliefs that we now know are not true.

A 2016 study led by researchers at the University of Virginia, USA, found that many members of the public and doctors incorrectly believed that black people have thicker skin and a higher pain threshold than white people. This is not true.[17]

These types of incorrect statements can be present in the data which is used to train AI chatbots. A 2023 study led by researchers at Stanford University, USA, found that several AI chatbots produced incorrect statements based on outdated and disproven ideas about racial differences, including skin thickness, lung function, and kidney health.[18]

The companies which make AI chatbots say they are working hard to reduce these biases, and specific errors like these seem to be corrected once identified. But new research shows bias is still present, and is often

[15] AI overly affirms users asking for personal advice. Stanford Report. https://news.stanford.edu/stories/2026/03/ai-advice-sycophantic-models-research
[16] Cheng M *et al.* Sycophantic AI decreases prosocial intentions and promotes dependence. *Science* 391, eaec8352 (2026). https://doi.org/10.1126/science.aec8352
[17] Hoffman KM, Trawalter S, Axt JR, Oliver MN. Racial bias in pain assessment and treatment recommendations, and false beliefs about biological differences between blacks and whites. *Proceedings of the National Academy of Sciences* 113, 4296–4301 (2016). https://doi.org/10.1073/pnas.1516047113
[18] Omiye JA, Lester JC, Spichak S, Rotemberg V, Daneshjou R. Large language models propagate race-based medicine. *npj Digit. Med.* 6, 195 (2023). https://doi.org/10.1038/s41746-023-00939-z

R

Remember to check the facts

subtle.[19,20,21,22] The issue of bias in algorithms is an active problem rather than one that has been totally solved.

Should I trust the sources that chatbots give in their answers?

Some chatbots can provide sources to back up the answers they give. This might appear as a footnote or a reference, like in a scientific study. Other times it might appear like a link to a website.

Seeing these references might be reassuring. It suggests that the chatbot has evidence to support its response. However, research has shown that these can be inaccurate.[23,24,25,26] This means that the chatbot's reply might not be supported by the evidence it provides.

You should check these sources yourself before making important decisions about your health.

It would also be a good idea to look for independent sources of trustworthy health information.

How do I find trustworthy sources of health information on the internet?

Official health services like the NHS, medical charities, and international bodies like the World Health Organization are trustworthy sources of online health information.

You can easily find trustworthy health information by searching the PIF TICK website: https://piftick.org.uk. The PIF TICK is a trust mark run by the Patient Information Forum, a UK charity. Organisations that earn it are called Trusted Information Creators. To get the mark they must meet 10 criteria, and be reassessed every year.

You can also speak to a qualified healthcare professional. They will be able to give you evidence-based information from trustworthy sources.

How can I spot health misinformation online?

Misinformation is false information which people share because they think it is true.

Disinformation is false information which is created and shared on purpose to deceive people.

---

[19] Bouguettaya A, Stuart EM, Aboujaoude E. Racial bias in AI-mediated psychiatric diagnosis and treatment: a qualitative comparison of four large language models. *npj Digit. Med.* 8, 332 (2025). https://doi.org/10.1038/s41746-025-01746-4

[20] Docking JJ *et al.* Evaluating the Potential of Reasoning Large Language Models to Perpetuate Racial and Gender Disease Stereotypes in Health Care. *Journal of Medical Internet Research* 28, e82256 (2026). https://doi.org/10.2196/82256

[21] Liao J *et al.* Bias and Equity in LLM Applications for Healthcare: A Scoping Review. SSRN Scholarly Paper (2026). https://doi.org/10.2139/ssrn.6053515

[22] He L. *et al.* Large language model biases in health care: a scoping review and call for an integrated assessment framework. *J Am Med Inform Assoc* 33, 1822–1833 (2026). https://doi.org/10.1093/jamia/ocag089

[23] Tiller, N. B. *et al.* Generative artificial intelligence-driven chatbots and medical misinformation: an accuracy, referencing and readability audit. *BMJ Open* 16, e112695 (2026). https://doi.org/10.1136/bmjopen-2025-112695

[24] Walters, W. H. & Wilder, E. I. Fabrication and errors in the bibliographic citations generated by ChatGPT. *Sci Rep* 13, 14045 (2023). https://doi.org/10.1038/s41598-023-41032-5

[25] Largest study of its kind shows AI assistants misrepresent news content 45% of the time – regardless of language or territory. BBC Media Centre. https://www.bbc.co.uk/mediacentre/2025/new-ebu-research-ai-assistants-news-content

[26] AI search has a citation problem. Columbia Journalism Review. https://www.cjr.org/tow_center/we-compared-eight-ai-search-engines-theyre-all-bad-at-citing-news.php

R

Remember to check the facts

An international team of researchers evaluating five commonly used chatbots showed that chatbots can recite health misinformation in their responses.[27]

If you are making important decisions about your health, it is important you have access to information which is true. Unfortunately, in recent years there have been increasing amounts of health misinformation and disinformation online.

The World Cancer Research Fund has created the simple TRUST Test to help you avoid being misled. Read more about the TRUST Test (external website).[28]

The test has 5 steps:

T Too good to be true? Does it promise unrealistic results or quick fixes?

R Research-backed? Does it trace back to scientific evidence rather than personal stories or opinions?

U Understood? Has the person sharing the information understood the dangers and risks, or have they downplayed the harms and promoted unproven remedies over appropriate medical care?

S Source quality? Does it come from a trusted organisation or scientific source? What do other experts say?

T Think before you share. If it doesn't pass the TRUST Test, don't pass it on.

This last step protects other people by stopping misinformation spreading online.

[27] See 23.
[28] TRUST test. World Cancer Research Fund (WCRF). https://www.wcrf.org/living-well/health-guides-cookbooks/trust-test/



Talk to a real person when you need to

There are some things that AI chatbots just cannot do

Do not use an AI chatbot in an emergency. Seek immediate medical help instead.

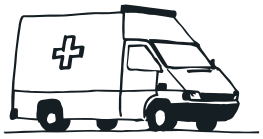
Talk to a qualified healthcare professional or someone else you trust if:

- You think a chatbot is giving you incorrect information,
- The conversation is going around in circles,
- You find its responses confusing,
- Your condition is getting worse,
- You are thinking about starting or stopping a medication,
- Or you are not sure what to do.

Using an AI chatbot is your decision. You don't need to tell a healthcare professional you've used them if you don't want to. But if its response has worried you, it may be helpful to talk to someone about it.

If the conversation with an AI chatbot starts to feel distressing, inappropriate, unusually important, or simply isn't going anywhere, stop and speak to someone you trust.

AI chatbots are tools which can help you manage your health, but they cannot replace healthcare professionals. Humans and AI chatbots work in different ways and are good at different things.



Do not use an AI chatbot in a medical emergency or mental health crisis

This is the one thing you should definitely not use a chatbot for. AI chatbots can miss things that are very risky. When researchers at the Icahn School of Medicine at Mount Sinai tested ChatGPT Health against realistic emergency

T

Talk to a real person when you need to

scenarios, it missed half of the cases that doctors agreed needed emergency care, including people at risk of suicide.[29,30]

If you think you or someone close to you is having a health emergency or is in immediate danger, get help from a trained health professional straight away. If you are in the UK, call 999 or go to A&E as soon as possible.

If you are not sure how urgent it is, call NHS 111, which is free and open 24 hours a day. If you are feeling low or struggling to cope, you can call Samaritans free on 116 123 or text SHOUT to 85258 at any time of day or night.

In many countries outside of the UK, you can often reach the emergency services by dialling 911 or 112. You can use this website to find help: https://findahelpline.com

Given how good AI chatbots are getting, why should I still speak to a health professional?

AI chatbots are computer programs which have been trained on vast amounts of information. Many chatbots have access to the internet, including published scientific research. Some studies have shown that AI chatbots can be better than humans at working out the answers to written medical questions and exam-style problems.[31,32]

However, there is much more to being a healthcare professional than simply knowing facts. Healthcare professionals can review your past medical history, ask you questions you may not have thought were relevant, and examine you to look for evidence of certain diseases. They will look at your existing test results and may recommend that you have further tests to help you get a diagnosis.

Once you have a diagnosis, they can recommend treatments which should help. Importantly, they also take into account your values and wishes and help you make health decisions that are right for you.

Healthcare professionals are also held accountable by regulations and laws in ways that AI chatbots are not.

When should I speak to a qualified health professional rather than a chatbot?

It might be better to speak to a healthcare professional if:

- You think a chatbot is giving you incorrect information,
- The conversation is going around in circles,

[29] Ramaswamy A *et al.* ChatGPT Health performance in a structured test of triage recommendations. *Nat Med* 32, 1671–1675 (2026). https://doi.org/10.1038/s41591-026-04297-7
[30] ChatGPT Health 'under-triaged' half of medical emergencies in a new study. NBC News. https://www.nbcnews.com/health/health-news/chatgpt-health-under-triaged-half-medical-emergencies-rcna261409
[31] Brodeur PG *et al.* Performance of a large language model on the reasoning tasks of a physician. *Science* 392, 524–527 (2026). https://doi.org/10.1126/science.adz4433
[32] AI outperforms doctors in Harvard trial of emergency triage diagnoses. The Guardian. https://www.theguardian.com/technology/2026/apr/30/ai-outperforms-doctors-in-harvard-trial-of-emergency-triage-diagnoses

T

Talk to a real person when you need to

- You find its responses confusing,
- Your condition is getting worse,
- You are thinking about starting or stopping a medication,
- Or you are not sure what to do.

If you can't easily get an appointment with a healthcare professional, you could also try talking to someone else you trust. You could also look for trustworthy sources of health information online. Read more about how to find trustworthy health information.

Should I tell my doctor that I've used an AI chatbot?

It's your choice if you use an AI chatbot to ask questions about your health. You don't need to tell anyone if you don't want to, even your doctor. Getting information about your health is nothing to be embarrassed about.

But if a chatbot response has worried you, or you want more information from someone qualified to give advice, it might be helpful to talk about it.

Many doctors would prefer to know exactly what their patients are worried about. This means they can give advice on the right things. Your doctor or another healthcare professional will be able to help you understand information you've been given by a chatbot. There might be some things which are useful, but there might also be areas where the chatbot has got things wrong.





How we wrote this guide

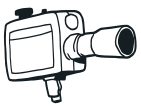
The Health Chatbot User Guide has been co-created by an international team.

We started by forming a **steering group**, which included patients and public representatives, health professionals, AI experts, people who work in regulation and health policy, ethicists, public health professionals, and academic researchers.



We then conducted a series of **scientific studies**:

- Our **online survey** helped us understand what people think a user guide for AI chatbots should contain. We were interested to hear both from AI experts, and from the wider public. The survey contained a series of quantitative and qualitative questions. Though different people had different perspectives, we used a qualitative framework analysis to map the similarities in what people were saying to us.
- We then ran a **rapid evidence review** to summarise what guidance was already published on how AI chatbots can be used safely to ask health questions. We built our search strategy and overall approach based on systematic review methodology, but made pragmatic changes to allow us to run the study more quickly. Specifically, we limited our search to January 2023 onwards, and in English only. We excluded studies that were not focused on health, or which did not explicitly contain guidance for the public when using general-purpose AI chatbots.
- We commissioned Ipsos to conduct **UK-wide polling** to understand how the public perceives chatbots, and how they use them.



We presented these findings to **patient and public groups** based in the UK and Australia. They helped us decide both what the first draft of the guide should contain, and how we should communicate the content to the public. They advised us that a website and PDF version wouldn't be enough, and that we should create **social media** and **video explainer content** too.



Finally, we used a **Delphi consensus study** to create the final guide. This involved an online survey where participants voted on whether each part of the guide should be included and/or amended. We received over 27,000 words of feedback, and used this to improve the guide. The guide was then reviewed in a **consensus meeting** with 13 members of the public who read it line-by-line, making final edits before ratifying this final version.

About this guide

Version

Version 1.0, published 19 October 2026. The field of AI chatbots for health is rapidly evolving, and this guide will become outdated quickly. We intend to revisit and update the content.

Suggested citations

We announced this programme in a Correspondence published in Nature Health in February 2026:

Khair DO, Kale AU, Agbakoba R et al. Building The Health Chatbot Users' Guide. Nat. Health 1, 555–556 (2026). https://doi.org/10.1038/s44360-026-00074-5

We are preparing a scientific paper about the research we conducted to develop the Health Chatbot User Guide, which will be submitted for publication in the autumn of 2026. In the interim, if you want to cite this PDF version, we suggest the following citation:

Laws E, Khair D, Kale A, Ong A, Balyot D, Alderman JE and the Health Chatbot User Guide author group. The Health Chatbot User Guide: a guide to using AI chatbots for health questions. 19 October 2026. DOI: 10.5281/zenodo.23503850

Authorship and acknowledgements

This project was led by researchers from the Bright group at the University of Birmingham. Dr Elinor Laws and Dr Duaa Khair led the research which underpinned the guide and its content, supported by Dr Aditya Kale and Dr Ariel Ong who contributed to methodology, study design and rapid review screening. Danica Balyot was the project manager, supported also by Nathalie Madoc, Charlotte Radovanovic and Sarah Bowkley. Dr Joseph Alderman developed the initial concept and was principal investigator, in collaboration with Dr Charlotte Blease.

The author group for the Health Chatbot User Guide includes the steering group, patient and public advisory group, consensus group, and scientific advisors (the scientific report will include affiliations as well as names):

Agbakoba R, Aston D, Auchterlonie M, Ayre J, Blease C, Cohen D, D'Souza H, Downing A, Fielding K, Freeborn C, Gilbert D, Goh E, Khay A, Kholwadia AK, Leung C, Lim E, Liu Y, Machirori M, Mackintosh M, Mateen B, McAteer SJ, McCradden M, Mistry P, Nicholls R, Pearson R, Pfohl S, Robinson T, Shelmerdine S, Sims W, Smith P, Steele C, Taylor D, Thomas M, Whitaker JSM, Williams B, Yates S.

We also wish to acknowledge the contribution of those who completed our online surveys and who supported this work in other ways. **Thank you.**

Funding

We are grateful to the Health Foundation who funded this work as part of a wider grant investigating the use and regulation of large language models in healthcare.

We also acknowledge the National Institute for Health and Care Research – Drs Alderman, Laws, Kale and Khair are all supported by the NIHR with personal awards (clinical lecturer; academic clinical fellow x 3 respectively). Without the paid research time and training enabled by these awards, this work could not have been delivered.

Competing interests and editorial independence

In building this guide we have collaborated with researchers and AI experts who work for technology companies, including those which develop AI chatbots. We believe this has strengthened the work – by adopting an inclusive approach, we have drawn from a broad range of perspectives.

In order to preserve editorial independence and mitigate competing interests, steering group members' input was considered as advisory, with all decisions on content and wording taken by our Delphi panel and consensus group – the latter comprised entirely of members of the public. No members of the local University of Birmingham research team declare any competing interests, and no funding or other support was received from any steering group member or their organisation.

Trademarks and copyright

ChatGPT, Gemini, Copilot, Claude and other named products are trademarks of their respective owners. Their mention in this guide does not imply endorsement of these products by the authors, or of this guide by their owners.

The text of the Health Chatbot User Guide © The authors, 2026, and is published under a Creative Commons Attribution 4.0 International licence (CC BY 4.0): https://creativecommons.org/licenses/by/4.0/. You are welcome to reuse, adapt or amend any portion of the text in the guide, as long as you attribute us.

Illustrations by Bridge Studios.

You are welcome to refer to the Health Chatbot User Guide and Ask SMART, and to use our logo to link to www.healthchatbotguide.org. However, please don't use our name or logo on modified versions of the guide, or in ways that imply our endorsement. For other uses, contact j.e.alderman@bham.ac.uk.