



Embase AI

Frequently Asked Questions



ELSEVIER

Advancing human progress together

Contents

| | |
|---|---|
| 1. Functionality and usage | 4 |
| 1.1. What is Embase AI?..... | 4 |
| 1.2. What does the filter funnel provide?..... | 4 |
| 1.3. I have access to Embase via my institution, but I don't see the Embase AI tab. What should I do? | 4 |
| 1.4. Can results pages be opened in a separate tab? | 4 |
| 1.5. What about opening the record details page of a reference? Can it be opened in a separate tab? | 5 |
| 1.6. Is it possible to use languages other than English for the query? | 5 |
| 1.7. Is the LLM able to identify tradenames? If yes, where does that information come from? | 5 |
| 1.8. Can Embase AI be used to fill out search forms, e.g. PICO?..... | 5 |
| 1.9. Can Embase AI be used to create summaries of a predefined set of publications, i.e. of Embase records found in a regular search? | 5 |
| 1.10. Is there any limit in the number of questions I can run? | 5 |
| 1.11. Is Embase AI available via third-party vendors (OVID, Proquest, or STN)? | 5 |
| 1.12. Is it possible to integrate Embase AI with other software? | 5 |
| 1.13. Do you have any tips/tricks to start working with Embase AI? | 6 |
| 2. Technology and search methodology..... | 7 |
| 2.1. What is the underlying search query for the AI response?..... | 7 |
| 2.2. How are the top thirty references for response generation selected? | 7 |
| 2.3. What is the level of creativity/hallucinations in Embase AI results? And how do you reduce the risk of hallucinations?..... | 7 |
| 2.4. What are you doing to tackle different form of bias?..... | 7 |
| 2.5. Is there reproducibility in results?..... | 7 |
| 3. Content and coverage..... | 8 |
| 3.1. Will clinical trials.gov data be included in the Embase AI search? | 8 |
| 3.2. I'm subscribed to Embase Classic, would it be used to generate a response in Embase AI? | 8 |
| 3.3. I see a "full text" link in the reference section. What does it do? | 8 |
| 4. Other questions | 9 |
| 4.1. What's the long-term roadmap for this product?..... | 9 |

| | |
|--|---|
| 4.2. Is Embase AI counter compliant? If not, how can I request Embase AI usage statistics? | 9 |
| 4.3. How does Embase AI ensure data privacy and security? | 9 |
| 4.4. What happens when I click on the thumbs up/thumbs down button? | 9 |
| 4.5. Is the chat history automatically saved? | 9 |

1. Functionality and usage

1.1. What is Embase AI?

Embase AI is a conversational search feature generative AI that quickly delivers responses to scientific questions while providing a fully traceable route to the reference data sources.

Embase AI:

- Provides concise responses to scientific queries.
- Converts a natural language question into a working Embase query.
- Responses are backed by citations to the source of the information, coming from Embase.
- Synthesizes information from Embase into a tailored response.
- Allows multiple follow-up questions.

1.2. What does the filter funnel provide?

The filter funnel provides the user the ability to filter by publication year, publication type, and study types.

By default, for the response generation Embase AI doesn't include preprints, conference abstracts, clinicaltrials.gov articles and article in press or in process. This is to ensure the whole value of Embase is used as we are leveraging the power of manual indexing to ensure the results to generate the response are the highest quality possible. If you want to expand your search using the beforementioned records, you can always toggle it on to get records with the latest information on a topic.

1.3. I have access to Embase via my institution, but I don't see the Embase AI tab. What should I do?

Embase AI is only accessible for registered users. If you have access via your institution but you don't see the Embase AI tab you will need to register via the "sign in" button on the top right corner.

1.4. Can results pages be opened in a separate tab?

In this version you can only open the results page in the same tab. You can either use the "back" button in your browser or navigate through the Embase menus to get back to Embase AI page. We are working in ways to solve this and let users open results page in different tabs.

The user can open the results page via the link at the end of the Embase Query.

1.5. What about opening the record details page of a reference? Can it be opened in a separate tab?

Yes! You can open the record details page of a reference by right clicking on the record and select "Open link in a new tab".

1.6. Is it possible to use languages other than English for the query?

You can use any language you want on Embase AI, but the responses will be in English by default in most of the languages. If you want a translation of the summary you can add "*Translate the response into [language]*" in your preferred language. **Note:** Language translation has not been fully tested, so exercise caution when using languages other than English.

1.7. Is the LLM able to identify tradenames? If yes, where does that information come from?

Not yet. Embase AI only generates summaries of the records it retrieves after translating the natural language query into an Embase query.

1.8. Can Embase AI be used to fill out search forms, e.g. PICO?

Embase AI cannot be used to fill out search forms. It works separately creating Embase queries from natural language questions. However, it does interpret framework questions and will separate each of the parts of the research questions using the Boolean operator AND between them.

1.9. Can Embase AI be used to create summaries of a predefined set of publications, i.e. of Embase records found in a regular search?

Yes! Embase AI can generate a summary when asking to summarize up to 30 records added as PUIs. In the future you will be able to send the documents directly to Embase AI from the results page.

1.10. Is there any limit in the number of questions I can run?

We have set up a limit of tokens for each user on Embase AI. According to our tests you can ask between 350 to 500 questions in a month depending on the complexity of your query. Once you reached the limit you won't be able to ask another question with your username until the first of the following month. Even though we believe the limit is high enough to never hit it, if you encounter it please [contact support](#).

1.11. Is Embase AI available via third-party vendors (OVID, Proquest, or STN)?

No. Embase AI will be only available through Embase provided by Elsevier

1.12. Is it possible to integrate Embase AI with other software?

Currently, Embase AI is a standalone product. However, integration with third-party software is something we will continue to explore based on user feedback.

1.13. Do you have any tips/tricks to start working with Embase AI?

Best way to start working with Embase AI is experimenting. Some suggestions we could give to you are:

- Start with broad questions and use follow up questions to further refine the query.
- If you encounter an error when performing your search:
 - Rerun the query again.
 - Try to make it broader to exclude some terms from the Embase query.
- You don't need to ask a proper question into Embase AI. If you have a research framework question you can always use that format. For example, you can ask as a question: *population = smokers intervention = hypnosis comparison = nicotine patch outcome = quit smoking.*
- You can always specify publication dates for your response. To do so, add a text like "from the 90s" or "using documents from the 1980s".
- If you only want a specific publication type, you can always append it to your question. You can use: *Using only reviews as publication type, what is the best treatment for tuberculosis?*
- In case you get a response with a text "The information you're looking for can be found in the provided references." You can ask for the summary using "Summarize them" as a follow up question.
- Always Export your session after you're finished.
- Advanced searching tips
 - You can ask for a summarization of up to 30 documents. Just use the PUIs of the records and ask it so summarize. For example you can ask *find the following records L2033499183 L2033736252 L2038883182 L2033066919 L2034623970 and summarize them*
 - If you want to add a specific term into your query, specifically mention it during the prompt generation. For example, you can add "Term should be a key term" at the end of your question.
 - You can also provide a full Embase query directly to Embase AI as a question. After retrieving the top-30 documents, you can ask "summarize them" as a follow-up question. **Note:** Embase AI accepts up to 2000 characters.
 - For further information, users can access [Embase AI use cases and prompt guide | Embase Support Center](#)

2. Technology and search methodology

2.1. What is the underlying search query for the AI response?

An LLM transforms the natural language query into an Embase query taking into account medical descriptors based on the latest version on Emtree, free text with the corresponding field codes and different subheadings.

2.2. How are the top thirty references for response generation selected?

Embase AI uses a combination of lexical (BM25) and semantic (cross-encoder) methods. The BM25 plus cross-encoder approach combines traditional and neural methods for effective document retrieval. First, BM25 — a fast, term-based ranking algorithm — quickly retrieves a broad set of relevant candidate documents based on keyword matching. Then, a cross-encoder, a neural model that jointly encodes the query and each candidate document, re-ranks these candidates by deeply assessing their semantic relevance. This two-stage process balances efficiency and accuracy, enabling precise retrieval of the most relevant documents.

2.3. What is the level of creativity/hallucinations in Embase AI results? And how do you reduce the risk of hallucinations?

The prompt engineering that guides our large language models (LLMs) has been designed to be strict, with clear instructions and scope. For example, the responses generated by Embase AI must match the intent of your query. If the AI does not find relevant information, the user should be informed. Additionally, every claim or assertion made by Embase AI is always backed by a reference to the source text, ensuring transparency and traceability. The responses are also regularly tested against rigorous evaluation frameworks to further reduce the risk of hallucinations.

2.4. What are you doing to tackle different form of bias?

We take bias very seriously. Embase AI draws exclusively on the content in Embase, enabling us to point directly to the abstracts behind any claims or assumptions it makes. Our search tools identify the abstracts that most closely match your query — this ensures that content is selected based on its ability to respond to your question, not the number of citations it has received, or the journal it was published in.

Our prompt engineering also plays an important role, instructing the LLM to filter out ‘unsafe’ responses; these are typical responses that exacerbate prejudice, harm or stereotypes against specific individuals or groups. We also have easy feedback mechanisms for users to report harmful or biased responses they receive. These reports are manually reviewed by our team.

2.5. Is there reproducibility in results?

LLM's do not guarantee reproducibility. However, the LLMs used for Embase are maximized for focus and minimized for creativity. This ensures key concepts of the response are the same for more consistent results.

In addition, thousands of new records are added to Embase.com per day, making new information available that can change the responses in the summarization. However, users can save and reuse

the Embase query generated by the LLM to reproduce searches on Embase.com. It is also possible to save the chat history for future reference.

Embase follows Elsevier's Responsible AI Principles. Please see the details here: <https://www.elsevier.com/about/policies-and-standards/responsible-ai-principles>

3. Content and coverage

3.1. Will clinical trials.gov data be included in the Embase AI search?

Yes! All content on Embase will be used in Embase AI.

3.2. I'm subscribed to Embase Classic, would it be used to generate a response in Embase AI?

Yes! All your Embase subscriptions (Embase Classic and Embase French Local Collection) will be part of the search on Embase AI

3.3. I see a "full text" link in the reference section. What does it do?

The "full text" link will take you to the publisher site where you will be able to download the article should you have access to it via your institution. This feature works the same way the feature found in the results page in Embase.

4. Other questions

4.1. What's the long-term roadmap for this product?

There is a roadmap for further development, this is not currently available but will be shortly.

4.2. Is Embase AI counter compliant? If not, how can I request Embase AI usage statistics?

The searches done via Embase AI are not COUNTER compliant for the time being, but we expect it to be COUNTER compliant in the upcoming months. You can always ask for Embase AI usage stats by contacting [customer support](#).

4.3. How does Embase AI ensure data privacy and security?

As we embed GenAI features in Embase and other products, we do so in line with Elsevier's [Responsible AI Principles](#) and [Privacy Principles](#).

For Embase AI, we use OpenAI's large language model (LLM) ChatGPT hosted on Microsoft Azure and have an agreement in place that information passed to this service will not be stored or used for training purposes. Our use of OpenAI's LLM is private, meaning there is no data exchange or use of our data to train OpenAI's public model.

Embase AI minimizes hallucinations by using only high-quality, curated Embase content. This grounds Embase AI when generating responses. Unlike many other natural language processing tools out there, Embase AI shows its workings with clear references to the journals and documents it uses to generate a response. In addition, Embase AI adheres to GDPR to guarantee user privacy. Elsevier does not store personal user information or chat history on our systems, unless done so in a compliant way that improves the product (like analytics or personalization). Elsevier does not share this data.

You can also rest easy knowing that the journals that users can have confidence in Embase based upon the rigorously vetted peer-reviewed journals selected for inclusion in Embase by independent experts on the Embase Content Selection team

4.4. What happens when I click on the thumbs up/thumbs down button?

Elsevier tracks the number of thumbs up or thumbs down at an organization level and for usage stats only. Feedback is not linked to the corresponding response and, in accordance with Elsevier's Responsible AI principles, Elsevier/Embase does not use them to fine tune any AI model.

4.5. Is the chat history automatically saved?

No, but we are working on this feature so that users will be able to save their conversation history and revisit those conversations later.