Public sources in TrevantAI

All source connections: what they offer and how they are used. The same texts are in the Help centre (ℹ️ next to the name of each source).

62 sources in 12 categories · Status: 2026-11-20

Contents

1. Adding public sources
2. For the administrator
3. Overview of the sources
4. Law and regulations
   - Rechtspraak (Dutch case law)
   - Laws and regulations
   - Official publications
   - Local regulations (CVDR)
   - EUR-Lex (EU law)
   - legislation.gov.uk (United Kingdom legislation)
5. Government and politics
   - Council information (OpenBesluitvorming)
   - House of Representatives (parliamentary papers)
   - European Parliament (documents)
6. Spatial planning and environment
   - Spatial plans
   - BAG (Dutch addresses and buildings)
   - Environmental documents (Environment and Planning Act, DSO)
7. Environment and climate
   - KNMI (weather data per station)
   - Luchtmeetnet (air quality per monitoring station)
   - Rijkswaterstaat (water level, temperature, waves and discharge)
8. Statistics
   - CBS StatLine (Statistics Netherlands)
   - RIVM and other third parties (via CBS)
   - Eurostat (European statistics)
9. Companies and finance
   - KVK Business Register (Dutch companies)
   - GLEIF (LEI register of legal entities)
   - Annual reports of listed companies (ESEF)
   - OpenCorporates (companies worldwide)
   - Companies House (UK companies)
   - SEC EDGAR (US listed companies)
   - TED (European public tenders)
   - Source connection Sanctions lists (EU, UK, US)
   - Source connection UN sanctions list (Security Council)
   - Source connection OpenSanctions (sanctions and PEPs)
10. Science
    - OpenAlex (scholarly publications)
    - Crossref (publications with a DOI)
    - arXiv (preprints)
    - OpenAIRE (European open research)
    - Unpaywall (open versions of articles)
11. Health and medicine
    - ClinicalTrials.gov (clinical studies)
    - Europe PMC (biomedical literature)
    - PubMed (medical literature)
    - WHO Global Health Observatory
    - EMA (European Medicines Agency)
    - openFDA (U.S. FDA)
    - DailyMed (U.S. product labeling)
12. Heritage and knowledge
    - Europeana (European cultural heritage)
    - Wikipedia (encyclopaedia)
    - Wikisource (historical texts)
    - Dutch National Monuments Register (RCE)
    - Rijksmuseum (collection)
    - RKD (artists and works of art)
    - Nationaal Archief (archival inventories)
    - Sound & Vision (Open Images)
    - Wayback Machine (web archive)
13. News and media
    - Source connection Feeds (RSS, Atom, JSON Feed)
    - GDELT (worldwide news)
    - Video sites (Vimeo, Dailymotion, TED)
14. Social media
    - Social media and marketplaces as sources
    - Marketplace listings as a source
    - Bluesky (social media)
    - Mastodon (social media)
    - X (formerly Twitter)
    - TikTok (video by link)

15. Religion and theology
    - Theological sources: overview
    - Verse numbering: why a verse sometimes has another number
    - Asking questions about theological sources
    - Bibles (public domain)
    - Marginal notes to the Statenvertaling
    - Original text and lexicon (Hebrew and Greek)
    - Cross-references, places and people
    - Creeds and catechisms
    - Sefaria (Jewish texts)

Adding public sources

On the Sources page, under **Explore & add sources**, you will find the **Public sources** tab. Use it to search public sources and add items as sources. The tab is only there if the administrator has switched on at least one source (Administration > Source connections).

Searching and adding

1. Under **Source**, choose where you want to search. If there is an active project, only the sources of that project are listed (Sources of a project).
2. Fill in the search fields and choose **Search**. The fields differ per source.
3. Tick the items you want (at most 20 at a time) and choose **Add to sources**. Retrieval runs in the background; the progress is shown in Task & Process Monitor. If a project is active, the items are linked to it.

If an item is already in the sources, this is indicated, with **View**.

The sources

- **Rechtspraak**: judgments and opinions of Dutch courts. You can search by ECLI number, court, area of law and period, but not for words in the text.
- **Laws and regulations**: the current text from wetten.overheid.nl. Search by title or BWB number; after adding, the text is searchable per article.
- **Official publications**: Staatsblad, Staatscourant, municipal and other gazettes, parliamentary papers and proceedings.
- **Council information**: documents, meetings and motions of Dutch municipal councils, provincial councils and water authorities (OpenBesluitvorming).
- **Spatial plans**: see *Adding spatial plans*.
- **CBS StatLine**: statistical tables. A table becomes a snapshot; very large tables are shortened.
- **BAG (Dutch addresses and buildings)**: address, dwelling unit and building from the Kadaster's registration: year of construction, floor area, purpose of use and status. Search by postcode and house number, or by street, house number and town. An address becomes a snapshot.
- **OpenAlex**: descriptions and abstracts of scholarly publications from all over the world, with a link to the full text.
- **Crossref**: descriptions of publications with a DOI; useful for finding an article exactly by its DOI.
- **More sources** (each source explains what is possible when searching): EU law (EUR-Lex), local regulations (CVDR), the Dutch House of Representatives, European tenders (TED); KNMI, Luchtmeetnet and Rijkswaterstaat (measurements as a snapshot); KVK, GLEIF, OpenCorporates, Companies House and SEC EDGAR (companies); arXiv, ClinicalTrials.gov, Europe PMC and PubMed; Europeana, Wikipedia and Wikisource (historical texts in their original language). Some require the organisation's own key.
- **Religion and theology**: Bibles, Marginal notes to the Statenvertaling, Original text and lexicon (Hebrew and Greek), Cross-references, places and people, and Creeds and catechisms work from the TrevantAI computer once the administrator has downloaded them; the search query then goes nowhere. Sefaria (Jewish texts) searches live. See Theological sources: overview.
- **Tip pointing to another source**: if what you are looking for is (also) in another source, a tip appears above the results, with a button that searches there straight away (only if that source is switched on).

Traceable

Each item becomes one source. Every piece of text starts with the path to the article or paragraph, so search results and answers state exactly where something is. Under **View** you can see where the item came from: the source, the identifier (such as an ECLI), the link, the date of retrieval, the search query and the licence. It is a snapshot: the source may have changed since.

Privacy

- The search query is sent to the source itself. Please do not enter client names or other confidential information there.
- TrevantAI never searches these sources on its own; this only happens on this tab.
- As a type of source, these items are called **Public sources**. Items from sources that may contain personal data (Rechtspraak, official publications, council information, the BAG) are processed by the local AI only by default; the administrator can set this differently per source. Questions about the items always go to the local AI.

In a project

If there is an active project, the *Public sources* tab only lists the sources of that project. Above the choice list it says how many there are. Without an active project (in the library), all sources the administrator has switched on are listed.

Where they come from

When a project is created, its sources come from the source profile of the chosen template; a project without a template receives a few very general sources. An existing project received, at the changeover, the sources of its template, plus the sources it already had items or subscriptions from, so that nothing is lost. If a source from the profile is switched on in the installation later, it is automatically switched on in this project, unless the project owner had removed it.

Sources of this project…

With the button *Sources of this project…*, the project owner chooses which sources appear in the choice list. The window only shows sources that are switched on in the installation, by category, with search and filters. In a team project the choice applies to everyone; only the owner can change it, and not while the project is locked. A source that is not listed at all is switched on by the administrator under Admin > Source connections.

Nothing is lost

Removing a source from a project deletes nothing: items already added remain sources of the project, and subscriptions to that source continue. The profile of the template does not change.

The starting point

TrevantAI is a tool for collecting, searching and questioning sources. What is retrieved with it and how it is used is the responsibility of the user. TrevantAI does not put up fences, but points to the rules.

What applies

- *The law*, in particular copyright and the rules on personal data (in the European Union the GDPR).
- *The terms and licences* of the websites, services and sources used, such as the terms of use of a website or video service, or the licence of a public source.
- *The agreements of one's own organisation*.

How TrevantAI helps

- For public sources, the licence or condition is shown with the source, and an export that contains a source with a restriction gives a warning beforehand.
- For every source it can be seen where it comes from and when it was retrieved.
- For web pages that TrevantAI checks automatically, TrevantAI respects the website's robots.txt.

In case of doubt

If in doubt whether something is allowed, it is wise to read the terms of the source or to ask a lawyer for advice. If something in TrevantAI turns out not to comply with the rules, TrevantAI can adjust it in a next version; it can be reported to the administrator or to TrevantAI.

What a subscription is

A subscription is a saved search at one public source, for example new rulings of the Supreme Court in civil law. TrevantAI repeats that search by itself at set times and only fetches what is new since the previous time. Rechtspraak (Dutch case law), Official publications, any website with a feed (RSS, Atom or JSON Feed) and also EMA, openFDA, DailyMed, OpenAIRE, the Rijksmuseum, the Nationaal Archief, Beeld en Geluid (Open Beelden), GDELT, the Wayback Machine, Bluesky and X can be followed now. A source that can be followed has the button *Follow…*; the help topic of that source lists the filters.

Creating a subscription

On the *Public sources* tab, next to *Search*, there is the button *Follow…*: it opens the *Subscription* window, filled in with the search of that moment. On the *Subscriptions* tab one can also start straight away with *New subscription*. The window holds the name, the filters (each with a switch; only a filter that is on counts), what happens to new items, the projects, the date from which the first round fetches (by default 30 days back) and how often it is checked. With *Show what I would get now* the window shows in advance the items the filters would have produced in the past 7 days; nothing is saved then. After *Save* the first round follows immediately.

Following Official publications

For Official publications the filters are: the gazette, the category (for example environmental permit or event permit), the subject, the organisation (while typing, the names the source uses appear), the type of organisation, a radius around an address (100 m to 5 km), *Response period still open* and *Words*. TrevantAI fetches the lists of choices from the source; the names are in Dutch, as the source gives them. For the radius, the address is converted once into coordinates by the PDOK Location Server of the Dutch Land Registry; after that, only those coordinates and the radius go to the source. In New arrivals, an item shows until when responses can be submitted. An old publication that was only updated at the source does not count as a new item.

Following feeds

With Feeds, the address of the feed is the fixed filter; when saving, a website becomes the address of its feed. TrevantAI only requests the feed if something has changed. An item is new if its link has not been seen before; if the feed gives a newer modification date for an item that is already a source, the source shows *Changed at the source on …*. Feeds do not report deletions. *Words* only checks title and summary. The status and the circuit breaker apply per website. Following a podcast or YouTube channel works the same way; a new episode then comes in with the audio and a new video with the subtitles (see Feeds).

New arrivals

By default, new items first go into the list *New arrivals*. There they are not yet a source and go to no AI at all. The list can be searched and filtered; *Add selected* and *Add all shown* turn them into sources, linked to the chosen projects, and *Ignore* removes them from the list (they do not come back). What stays in the list for 90 days expires. Choosing *Add to my sources straight away* skips the list.

Where to find it

Under Explore & add sources there is the *Subscriptions* tab, showing per subscription its status, how often, the last and the next time, and the buttons *New arrivals*, *Edit…*, *Pause*, *Check now* and *Delete*. An added item is an ordinary source among one's own sources and can be linked to any project. The source shows through which subscription it arrived, with the filters of that moment. If the source later reports that an item was deleted or changed, this is shown with the source; the text in TrevantAI stays.

Notifications

In the *Subscription* window, under *Notify*, there is *In TrevantAI (counter in the top bar)*, on by default, and *By e-mail*, off by default. After a round with new items the subscription appears at the Notifications bell in the group *Subscriptions*, with the number; *View* opens *New arrivals* (with *Add straight away*: the new sources). Items that are already a source and were deleted or changed at the source are also reported. The e-mail comes per round, with the titles (at most 25). If a round fails for 24 hours, the group *Problems* shows once "not working since …", with the reason; if the administrator switches the source or following off, or pauses or stops the subscription, straight away. Once it works again, that notification disappears by itself. Problems always appear at the bell; by e-mail if *By e-mail* is on.

Privacy and limits

The filters go to the source in every round: do not use names of clients or other confidential details. Only the user creates a subscription; the AI never does. Rechtspraak has no search on words: with the filter *Words* TrevantAI fetches the new rulings and keeps only those with all the words. The administrator decides which sources and filters exist, how many subscriptions a person may have, how many items arrive per round and how often at most it is checked.

If the source itself is not working

If the source itself does not respond, one notification follows after an hour for all subscriptions to that source together, instead of a notification per subscription. The subscriptions wait and catch up as soon as the source works again. Each subscription then shows the status of the source. See Problems with sources.

For the administrator

Managing source connections

Under **Administration > Source connections** you choose which public sources TrevantAI may query. Only what is switched on here appears for users under Sources > **Public sources**. Everything is switched off by default. In the Single-user edition you are the administrator yourself.

For each source you can see

- whether a key is needed (for some sources a key only gives higher limits);
- whether the source may contain personal data (then added items are processed by the local AI only by default) or is public (then an external AI may process the items, if the AI policy allows cloud AI);
- the licence, with a link to the terms and to the website of the source.

Switch a source on or off with the switch. Every change is recorded in the security log.

Who may process the items?

For each source, under **Who may process items from this source?**, you choose whether only the local AI processes the items, or an external AI too (if the AI policy allows cloud AI). By default it is the local AI only for sources that may contain personal data (such as Rechtspraak and the BAG), and an external AI too for public sources. The choice applies to items added from that moment on; it can still be changed per document afterwards. Questions about the items always go to the local AI. This change is also recorded in the security log.

Keys

TrevantAI is supplied without keys. If a source needs a key (such as Spatial plans and the BAG), please request it from the source yourself; the guide is behind the information sign (ℹ️) next to **Key**. Enter the key at the source and choose **Save**; use **Test** to check it. The key is stored encrypted and never shown again.

Source profiles

Which of these sources are switched on in a project is determined by the source profile of the template the project is created with; the project owner can adjust this per project. A project never switches on a source for the whole organisation. If a source from the profile is switched off, the administrator sees this on the Public sources tab; once it is switched on, it is automatically switched on in the projects that have it in their list. See Source profiles.

Fair use

The sources are free, under their own terms. TrevantAI only requests what a user searches for or adds, and identifies itself with every request. If a source asks to wait, the user receives a message.

Check

With `npm run controleer:koppelingen` (in the backend folder) you can check whether all sources can be reached and respond as expected.

Status

Each source that is switched on shows its status (*Working*, *Slow*, *Intermittent*, *Down since …*, *The source asks to wait*, *Access refused*, *Not used*) with a strip of seven days. Everything together, with the problems of the past 90 days: Admin > Problems.

Source profiles

Two levels

Admin > Source connections shows which public sources are *available* in this installation: with key, terms and AI choice. Which of these are *switched on* in a project is determined by the **source profile** of the template the project is created with. A project never switches on a source for the whole organisation; that remains the administrator's task (in the Single-user edition the user personally).

The source profile of a template

Under **Admin > Templates**, each template has the button *Source profile…*; the template editor has the same button at the top. The window shows all source connections by category, with a search field and the filters *Show*, *Without a key* and *Without personal data*, and a checkbox for each source. Sources that are switched off in the installation can also be ticked; they carry the label *Not switched on*. In a new project they are on the list, but they only appear once the administrator switches them on.

Standard templates

The standard templates have a profile supplied with TrevantAI. A change is saved as this installation's own profile, with the label *Changed in this installation*; the template itself stays untouched. *Back to the TrevantAI profile* undoes the change. If a TrevantAI update adds a source to the profile of a template that has been changed here, that source is not added automatically: a notice appears above the list in Admin > Templates, and the window offers the buttons *Add* and *Do not add*. For an unchanged profile, the new source is simply included.

Own templates

A new own template receives the profile of its starting point (a template or the pathway of a project), or for an empty template the **default profile for own templates**. In the wizard for a new template, *Sources* is the third step. The owner changes the profile of a private template personally; that of a shared template or a standard template can only be changed by those with the right to manage profession templates.

Default profile for own templates

At the top of Admin > Templates is *Default profile for own templates…*. TrevantAI supplies very general sources, without a key and without personal data: Wikipedia, Wikisource, OpenAlex, CBS StatLine and Eurostat. A project without a template also receives this profile. A change applies to templates and projects created afterwards.

Existing projects

Changing a profile does not change existing projects, just as the sections of a template do not change existing projects. The project owner chooses the sources of an individual project: Sources of a project.

Terms of public sources

Each public source has a **type of use**. You see it with the source in **Administration > Source connections** and when searching on the **Public sources** tab.

- **Free to use**: also for commercial use, with attribution where required (e.g. Rechtspraak, Wetten, EUR-Lex, CBS).
- **Free, with attribution and share-alike**: also commercially; anyone who passes on the data itself does so under the same licence (e.g. Wikipedia).
- **Description free, text protected by copyright**: abstracts and full texts belong to authors or publishers; suitable for your own research and quoting (e.g. PubMed, Crossref).
- **Under your own agreement with the source**: works with a key or contract of your organisation, which makes its own agreement with the source (e.g. KVK, OpenCorporates).
- **Non-commercial only**: usable only if the installation is used non-commercially (see *Use of this installation*).

Confirming terms

For sources with substantial restrictions, TrevantAI shows a short summary of the terms and a link to the full text when switching on. Only after ticking "*Our organisation uses this source in accordance with these terms*" is the source switched on. Who confirmed this and when is recorded in the security log. This currently applies to KVK (no passing on of register data to others) and OpenCorporates (a free key only for open-data projects).

When the terms change

If a source changes its terms, an update of TrevantAI adjusts the summary. Until the administrator has confirmed the new terms, nothing new can be added from that source. What has already been added remains.

When a source changes or disappears

A new version of the same source simply keeps working after an update. If a source is replaced by a completely different one or stops, it gets the status *outdated*: nothing new can be added from it, and the explanation refers to any successor. Added items always remain: they are snapshots with date, origin and licence.

Responsibility

Your organisation decides which sources are switched on and is itself responsible for complying with their terms. TrevantAI shows the terms to the best of its knowledge; the full text at the source is decisive.

Exporting outside the organisation

Items that are *free only in the US* are not included in a theme dossier by default (users can deliberately include them). For items with an *ND licence* the text may be included, but no translation or adaptation. For exports made by AI (a chat answer, the whole chat, the final report, a semantic comparison) and for a text comparison, TrevantAI first asks whether to go ahead anyway. In those exports, the raw export and search results, such sources carry a label with a short explanation.

In a *To share* project file, items that are free only in the US or may only be used within the organisation are never included, nor is a translation or adaptation of an item with an ND licence. In a *Complete* export they are included, with a warning beforehand. See the topic *Exporting a project*.

The source library on this computer

Large free sources are stored on the TrevantAI computer itself, in a **source library**. This applies to the theological sources: Bibles with Marginal notes (±390 MB), Original text and lexicon (±85 MB), Cross-references, places and people (±20 MB) and Creeds and catechisms (±9 MB). Searching and adding therefore work without internet.

Installing

Switching a source on under Administration > Source connections installs it: TrevantAI retrieves the data once from the source itself (eBible.org, STEPBible, OpenBible.info, Creeds.json on GitHub). This is a background task; the block *Source library on this computer* shows the progress, and the **Task & Process Monitor** holds the log. Beforehand it shows how much will be retrieved and how much space is free; with too little space, *Install* is disabled.

The buttons

- **Install / Update**: retrieve (again). Searching keeps working meanwhile; the new version only becomes active, all at once, at the end.
- **Free up space**: removes the library from the computer. Items already added to projects remain. Bibles and Marginal notes share one library.
- *This source needs a newer version of the library*: click **Update**; only the missing part is retrieved.

Order: first Bibles, then Original text and lexicon, then Cross-references. That way the Bible texts and Dutch names appear everywhere at once.

Failed? The reason is in the task log (Task & Process Monitor). It is often a temporary network problem: try again later with **Update**.

Overview of the sources

Source	Key	Type of use	Personal data	Where
Law and regulations				
⚖️ Rechtspraak (Dutch case law)	none	Free to use	may contain	at the source
📜 Laws and regulations	none	Free to use	no	at the source
📰 Official publications	none	Free to use	may contain	at the source
🏘️ Local regulations (CVDR)	none	Free to use	no	at the source
EU EUR-Lex (EU law)	none	Free to use	may contain	at the source
GB legislation.gov.uk (United Kingdom legislation)	none	Free to use	no	at the source
Government and politics				
🏛️ Council information (OpenBesluitvorming)	none	Free to use	may contain	at the source
🏛️ House of Representatives (parliamentary papers)	none	Free to use	may contain	at the source
EU European Parliament (documents)	none	Free to use	may contain	at the source
Spatial planning and environment				
🗺️ Spatial plans	required (own key)	Free to use	no	at the source
🏠 BAG (Dutch addresses and buildings)	required (own key)	Free to use	may contain	at the source
🏞️ Environmental documents (Environment and Planning Act, DSO)	required (own key)	Free to use	no	at the source
Environment and climate				
🌦️ KNMI (weather data per station)	required (own key)	Free to use	no	at the source
🌫️ Luchtmeetnet (air quality per monitoring station)	none	Free to use	no	at the source
🌊 Rijkswaterstaat (water level, temperature, waves and discharge)	none	Free to use	no	at the source
Statistics				
📊 CBS StatLine (Statistics Netherlands)	none	Free to use	no	at the source
🩺 RIVM and other third parties (via CBS)	none	Free to use	no	at the source
📈 Eurostat (European statistics)	none	Free to use	no	at the source
Companies and finance				

Source	Key	Type of use	Personal data	Where
🏢 KVK Business Register (Dutch companies)	required (own key)	Under your own agreement with the source	may contain	at the source
🌐 GLEIF (LEI register of legal entities)	none	Free to use	no	at the source
📈 Annual reports of listed companies (ESEF)	none	Description free, text protected by copyright	may contain	at the source
🗂️ OpenCorporates (companies worldwide)	required (own key)	Under your own agreement with the source	may contain	at the source
GB Companies House (UK companies)	required (own key)	Free to use	may contain	at the source
US SEC EDGAR (US listed companies)	none	Free to use	no	at the source
📢 TED (European public tenders)	none	Free to use	no	at the source
🚫 Source connection Sanctions lists (EU, UK, US)	none	Free to use	may contain	local (download)
UN Source connection UN sanctions list (Security Council)	none	Description free, text protected by copyright	may contain	local (download)
🛡️ Source connection OpenSanctions (sanctions and PEPs)	required (own key)	Under your own agreement with the source	may contain	at the source
Science				
🎓 OpenAlex (scholarly publications)	optional	Description free, text protected by copyright	no	at the source
📚 Crossref (publications with a DOI)	none	Description free, text protected by copyright	no	at the source
📄 arXiv (preprints)	none	Free to use	no	at the source
EU OpenAIRE (European open research)	none	Description free, text protected by copyright	no	at the source
🔓 Unpaywall (open versions of articles)	required (own key)	Free to use	no	at the source
Health and medicine				
💉 ClinicalTrials.gov (clinical studies)	none	Free to use	no	at the source
🩺 Europe PMC (biomedical literature)	none	Description free, text protected by copyright	no	at the source
📖 PubMed (medical literature)	optional	Description free, text protected by copyright	no	at the source
🌍 WHO Global Health Observatory	none	Free to use	no	at the source
💊 EMA (European Medicines Agency)	none	Free to use	no	at the source

Source	Key	Type of use	Personal data	Where
💉 openFDA (U.S. FDA)	optional	Free to use	no	at the source
📋 DailyMed (U.S. product labeling)	none	Free to use	no	at the source
Heritage and knowledge				
🖼️ Europeana (European cultural heritage)	required (own key)	Free to use	no	at the source
📚 Wikipedia (encyclopaedia)	none	Free, with attribution and share-alike	may contain	at the source
🪶 Wikisource (historical texts)	none	Free, with attribution and share-alike	no	at the source
🏛️ Dutch National Monuments Register (RCE)	none	Free to use	no	at the source
🖼️ Rijksmuseum (collection)	none	Free to use	no	at the source
🎨 RKD (artists and works of art)	none	Free to use	may contain	at the source
🗄️ Nationaal Archief (archival inventories)	none	Free to use	may contain	at the source
🎞️ Sound & Vision (Open Images)	none	Free to use	may contain	at the source
🕰️ Wayback Machine (web archive)	none	Description free, text protected by copyright	may contain	at the source
News and media				
📰 Source connection Feeds (RSS, Atom, JSON Feed)	none	Description free, text protected by copyright	may contain	at the source
🌐 GDELT (worldwide news)	none	Description free, text protected by copyright	may contain	at the source
🎬 Video sites (Vimeo, Dailymotion, TED)	none	Description free, text protected by copyright	no	at the source
Social media				
🦋 Bluesky (social media)	optional	Description free, text protected by copyright	may contain	at the source
🐘 Mastodon (social media)	optional	Description free, text protected by copyright	may contain	at the source
𝕏 X (formerly Twitter)	none	Description free, text protected by copyright	may contain	at the source
🎵 TikTok (video by link)	none	Description free, text protected by copyright	may contain	at the source
Religion and theology				
📖 Bibles (public domain)	none	Free to use	no	local (download)
📝 Marginal notes to the Statenvertaling	none	Free to use	no	local (download)

Source	Key	Type of use	Personal data	Where
🔤 Original text and lexicon (Hebrew and Greek)	none	Free to use	no	local (download)
🗺️ Cross-references, places and people	none	Free to use	no	local (download)
📜 Creeds and catechisms	none	Free to use	no	local (download)
🕎 Sefaria (Jewish texts)	none	Free, with attribution and share-alike	no	at the source

Law and regulations

⚖️ Rechtspraak (Dutch case law)

Judgments and opinions of Dutch courts from the open data of Rechtspraak.nl: search by ECLI number, court, area of law or period and add the full text as a source.

> Key: none · Type of use: Free to use (Public government information) · Personal data: may contain · Where: at the source

Rechtspraak is the open data of the Dutch judiciary (Rechtspraak.nl). Through this source, TrevantAI retrieves published judgments of district courts, courts of appeal, the Supreme Court (Hoge Raad) and other courts, as well as opinions of the Advocate General. Each item has a permanent identifier: the ECLI number (European Case Law Identifier), for example ECLI:NL:HR:2019:2006.

What can you use this source for?

- Legal: collecting the judgments referred to in a pleading or legal opinion, and querying them side by side in a project.
- Government and policy: checking how courts apply a regulation, for example all judgments of one court in one area of law over a year.
- Journalism: retrieving the judgment in a case in the news, using the ECLI number from the press release.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **Rechtspraak (Dutch case law)**. The fields:

- **ECLI number(s)**: one or more ECLI numbers. You can also paste a text or a link containing ECLI numbers, such as a link from uitspraken.rechtspraak.nl; TrevantAI extracts the numbers, at most 20 at a time. If this field is filled in, the other fields are not used.
- **Court**: a list of the current courts, as supplied by the source. Courts that have been abolished are not included.
- **Area of law**: a list supplied by the source; sub-areas are shown with a dash under their main area.
- **Type**: **Judgments** or **Opinions (Advocate General)**. If you choose nothing, you get both.
- **From (date of judgment)** and **Up to and including**: the period. Only a start date: up to today. Only an end date: from 1900.

Without an ECLI number, at least a court, an area of law or a date is required. Then choose **Search**. The most recent judgments are listed first, 20 at a time; use **More results** to see more.

This source does not support searching for words in the text; the open data do not offer this. Tip: combine court, area of law and a short period, or look up the ECLI number first. Once added, the text is fully searchable in TrevantAI.

What you get

Tick the items you want (at most 20 at a time) and choose **Add to sources**. Retrieval runs in the background; the progress is shown in Task & Process Monitor. Each item becomes a document with:

- the summary (inhoudsindicatie) as the first section;
- then the judgment or opinion per part, as structured by the judiciary, with the breadcrumb trail in the heading (ECLI number › Uitspraak › heading of the part);
- the details: ECLI number, court, date of judgment, date of publication, case number, type, procedure and areas of law, with the link to uitspraken.rechtspraak.nl.

The **Public source** card with the document shows the source, the identifier, the link, the date at the source, **Retrieved on** and **Found with** (the search fields). The document is a **snapshot**: the source may have changed since. If a judgment has no text (yet), it is not added and TrevantAI says so.

Key

No key needed. The administrator only has to switch the source on under **Administration > Source connections**.

Privacy

The search query is sent to the judiciary (data.rechtspraak.nl). You fill in the search fields yourself; the AI never searches this source on its own. Do not enter client names or other confidential information. Judgments are pseudonymised but may still contain names, for example of companies. That is why items from this source are processed by the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**; this applies to items added afterwards. Questions about the items always go to the local AI.

Licence and attribution

The judgments are public government information and may be reused under the terms for the open data of Rechtspraak.nl (the link is shown with the source in Administration > Source connections). TrevantAI stores the attribution **Rechtspraak.nl** with the document. When quoting, state the ECLI number.

Good to know

- The names in the lists come from the source and are in Dutch.
- If a judgment is already in the sources, **Already in the sources** is shown, with **View**.
- Traceable: the ECLI number is a permanent identifier, each fragment carries the breadcrumb trail, and **View at the source** opens the judgment on Rechtspraak.nl.

📜 Laws and regulations

The current text of Dutch laws, orders in council, ministerial regulations and other regulations from the Basiswettenbestand (wetten.overheid.nl): search by title or BWB number and add the text, article by article, as a source.

> Key: none · Type of use: Free to use (Free to use (CC0)) · Personal data: no · Where: at the source

Laws and regulations retrieves the current text of Dutch legislation from the **Basiswettenbestand (BWB)**, the collection behind wetten.overheid.nl, through the government search service (KOOP). It covers acts, orders in council (AMvB), ministerial regulations, Kingdom acts, royal decrees, treaties, policy rules, circulars and regulations of independent administrative bodies. Each regulation has a permanent BWB number, for example BWBR0005537 for the General Administrative Law Act (Algemene wet bestuursrecht).

What can you use this source for?

- Legal: querying the articles of an act alongside a case file, with the correct article in every quotation.
- Government and policy: testing a policy paper against the legislation in force, for example the Environment and Planning Act (Omgevingswet) and its decrees.
- Business and compliance: checking which rules apply to a sector, such as a ministerial regulation with technical requirements.
- Education and science: a legal text as a fixed, traceable source in a study.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **Laws and regulations**. The fields:

- **Title or BWB number** (required): words from the title, such as *Algemene wet bestuursrecht*, or a BWB number (BWBR… or BWBV… for treaties). You can also paste a link from wetten.overheid.nl; TrevantAI extracts the BWB number.
- **Type of regulation**: for example **Act**, **Order in council (AMvB)**, **Ministerial regulation**, **Kingdom act**, **Royal decree**, **Treaty**, **Policy rule** or **Circular**. Empty = all types.

Then choose **Search**. The source only searches the titles (short title, official and unofficial title), not the text of the law itself, and only regulations that are in force **today**. A regulation appears as soon as its title contains one of the words, so use distinctive words or the BWB number. Each regulation is listed once, with type, area of law and BWB number; 20 results at a time, and **More results** shows the next ones. Titles are in Dutch, so search with Dutch words.

What you get

Tick the regulations you want (at most 20 at a time) and choose **Add to sources**. Retrieval runs in the background (Task & Process Monitor). TrevantAI retrieves the version in force today and makes each article a separate section, with the breadcrumb trail in the heading, for example *Algemene wet bestuursrecht › Hoofdstuk 1 › Artikel 1:1*. This way every answer and quotation shows which article the text comes from.

The following are recorded with the document: the BWB number, the type, the area of law, the date from which this version applies, the number of articles and the link to exactly that version on wetten.overheid.nl. The **Public source** card also shows **Retrieved on** and **Found with**. The document is a **snapshot**: if the law is amended later, the document does not change with it.

Key

No key needed. The administrator only has to switch the source on under **Administration > Source connections**.

Privacy

The search query is sent to the government search service (zoekservice.overheid.nl). Do not enter client names or other confidential information. Laws and regulations contain no personal data; that is why an external AI may process the items by default, if the AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The Basiswettenbestand is free to use (CC0): you may copy and reuse the texts without restriction. TrevantAI stores the attribution **wetten.overheid.nl** with the document; when quoting, state the regulation, the article and the date of the version.

Good to know

- Only today's version: earlier and future versions cannot be retrieved through this connection. If a regulation is no longer in force, TrevantAI says so.
- Large acts have hundreds of articles (the General Administrative Law Act more than 570); retrieval may then take a while.
- If a regulation is already in the sources, **Already in the sources** is shown, with **View**.
- Traceable: the BWB number and the date of the version are part of the link; **View at the source** opens the regulation on wetten.overheid.nl.

Official publications

Staatsblad, Staatscourant, municipal, provincial and water authority gazettes, parliamentary papers and proceedings (officielebekendmakingen.nl): search by words or identifier, gazette, organisation and period, and add publications as sources.

> Key: none · Type of use: Free to use (Free to use (CC0)) · Personal data: may contain · Where: at the source

Official publications are the publications in which the Dutch government officially announces rules and decisions, as on officielebekendmakingen.nl. TrevantAI searches them through the KOOP search service (repository.overheid.nl). The source covers the Staatsblad (Bulletin of Acts: laws and decrees), the Staatscourant (Government Gazette: including ministerial regulations and permits), the municipal gazette, the provincial gazette, the water authority gazette, the joint arrangement gazette, the Tractatenblad (treaty series), parliamentary papers, parliamentary proceedings and parliamentary questions. Each publication has a permanent identifier, for example stcrt-2024-26001.

What can you use this source for?

- Legal: finding the decision against which an objection is pending, or the publication with which a rule came into force.
- Government and policy: keeping track of the by-laws and decisions a municipality or province published in a period.
- Journalism: following permits and other decisions in a municipality, for example around a building project.
- Business: checking which ministerial regulations or instructions have been published for a sector.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **Official publications**. The fields:

- **Search words or identifier** (required): one or more words, for example *omgevingsvergunning Dorpsstraat*. Each word must occur in the publication; at most eight words are used. Or a single identifier, such as stcrt-2024-26001; pasting a link from zoek.officielebekendmakingen.nl also works.
- **Gazette**: **Staatsblad (Bulletin of Acts)**, **Staatscourant (Government Gazette)**, **Municipal gazette**, **Provincial gazette**, **Water authority gazette**, **Joint arrangement gazette**, **Treaty series (Tractatenblad)**, **Parliamentary papers**, **Parliamentary proceedings** or **Parliamentary questions**. Empty = all gazettes.
- **Organisation (e.g. municipality)**: the name of the publishing organisation, for example *Utrecht*.
- **From** and **Up to and including**: the publication period.

Then choose **Search**. The most recent publications are listed first, 20 at a time; **More results** shows the next ones. Each result shows the gazette, the organisation, the type of publication and the identifier. Tip: choose a gazette and an organisation to keep the list short; general words such as *besluit* (decision) produce many hits. The publications are in Dutch, so search with Dutch words.

What you get

Tick the publications you want (at most 20 at a time) and choose **Add to sources**. Retrieval runs in the background (Task & Process Monitor). TrevantAI reads the official XML version of the publication, without the fixed gazette header and the metadata. If the publication has articles or chapters, each becomes a section, with the breadcrumb trail (title › chapter › article) in the heading; otherwise it becomes one continuous text.

The following are recorded with the document: the identifier, the gazette, the organisation, the type, the publication date, the link to the publication and, if available, the link to the PDF. The **Public source** card also shows **Retrieved on** and **Found with**. The document is a **snapshot** of the publication.

Key

No key needed. The administrator only has to switch the source on under **Administration > Source connections**.

Privacy

The search query is sent to the government search service (KOOP). Do not enter client names or other confidential information. Official publications may contain personal data, for example the name and address with a permit. That is why items from this source are processed by the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The official publications are free to use (CC0). The GDPR still applies to personal data in a publication: being public does not mean they may be used for any purpose. TrevantAI stores the attribution **officielebekendmakingen.nl** with the document.

Good to know

- The Staatsblad contains laws as adopted, including amending acts. For the current, consolidated text of a law, use the **Laws and regulations** source.
- Parliamentary papers and questions are also in the **House of Representatives (parliamentary papers)** source, where you search by subject and document type.
- If a publication has no text that TrevantAI can read, it is not added and TrevantAI says so; the PDF can still be opened via **View at the source**.
- If a publication is already in the sources, **Already in the sources** is shown, with **View**.
- Traceable: the identifier is permanent and each fragment carries the breadcrumb trail.

🏠 Local regulations (CVDR)

By-laws, policy rules and other regulations of Dutch municipalities, provinces, water authorities and joint arrangements from the CVDR (lokaleregelgeving.overheid.nl, KOOP): search by words or CVDR number and add the regulation, article by article, as a source.

> Key: none · Type of use: Free to use (Public government information) · Personal data: no · Where: at the source

Local regulations (CVDR) is the **Centrale Voorziening Decentrale Regelgeving** (central register of decentralised regulations) of KOOP, the collection behind lokaleregelgeving.overheid.nl. It contains the by-laws, policy rules and other regulations of Dutch municipalities, provinces, water authorities and joint arrangements, such as general local by-laws (APV), a tax by-law or a subsidy scheme. Each version of a regulation has a permanent CVDR number, for example CVDR121269_3 (the part after the underscore is the version).

What can you use this source for?

- Legal: the text of the municipal by-law at the heart of a case, article by article, alongside the case file.
- Government and policy: comparing regulations of different municipalities, for example the dog tax or the rules for terraces.
- Business: checking which local rules apply to a branch, a permit or an event.
- Journalism and science: following changes in local policy through successive versions.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **Local regulations (CVDR)**. The fields:

- **Search words or CVDR number** (required): one or more words, for example *hondenbelasting* (dog tax); each word must occur, and at most eight words are used. Or a single CVDR number, preferably with version (CVDR121269_3). Pasting a link from lokaleregelgeving.overheid.nl also works. The regulations are in Dutch, so search with Dutch words.
- **Organisation (e.g. municipality)**: the name of the municipality, province, water authority or joint arrangement, for example *Utrecht*.
- **Regulations in force only**: **Yes** shows regulations that are in force today; regulations for which the source gives an expiry date before today are left out by TrevantAI. **No, include repealed regulations** also shows earlier versions and repealed regulations. If you choose nothing, all regulations are shown.

Then choose **Search**. Results come 20 at a time; **More results** shows the next ones. Each result shows the organisation, the type of regulation, the CVDR number, the date of entry into force and, if there is one, the expiry date. With **Yes**, a page may show fewer than 20 results, because repealed regulations are left out. Tip: combine a distinctive word with the name of the organisation.

What you get

Tick the regulations you want (at most 20 at a time) and choose **Add to sources**. Retrieval runs in the background (Task & Process Monitor). TrevantAI reads the official XML of the regulation, without the metadata, and makes each article a section, with the breadcrumb trail in the heading (title › chapter › article). If you give a CVDR number without a version, TrevantAI retrieves the highest version the source finds.

The following are recorded with the document: the CVDR number with version, the organisation, the type, the subject, the date of entry into force, the expiry date, whether the regulation was still in force at retrieval, and the link to the regulation on lokaleregelgeving.overheid.nl. The **Public source** card also shows **Retrieved on** and **Found with**. The document is a **snapshot**.

Key

No key needed. The administrator only has to switch the source on under **Administration > Source connections**.

Privacy

The search query is sent to the government search service (KOOP). Do not enter client names or other confidential information. By-laws and policy rules apply to everyone and, apart from a signature, contain no data about individuals. That is why an external AI may process the items by default, if the AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

Local regulations are public government information and free to reuse. TrevantAI stores the attribution **Bron: lokaleregelgeving.overheid.nl (KOOP)** with the document.

Good to know

- With **Regulations in force only**, TrevantAI browses through the hits until there are enough regulations in force (at most 300 hits); expired versions are left out.
- The same regulation may appear in the results with several versions; pay attention to the version number and the dates.
- In a dispute, check which version applied on the relevant date; today's version is not always the right one.
- TrevantAI only retrieves the text from government sites (overheid.nl and officiele-overheidspublicaties.nl).
- If a regulation is already in the sources, **Already in the sources** is shown, with **View**.
- Traceable: the CVDR number with version is permanent; **View at the source** opens exactly that version.

EU EUR-Lex (EU law)

European Union legislation, preparatory acts and case law of the Court of Justice via the Publications Office (EUR-Lex, CELLAR): search by words from the title or CELEX number and add the text, article by article, as a source, in Dutch, English, German, French, Spanish, Italian or Portuguese.

> Key: none · Type of use: Free to use (Free to reuse with acknowledgement of the source (Decision 2011/833/EU)) · Personal data: may contain · Where: at the source

EUR-Lex is the official source of European Union law. TrevantAI uses the public services of the **Publications Office of the EU** (CELLAR), where the EUR-Lex texts come from. They contain EU legislation (regulations, directives, decisions), preparatory acts (such as proposals from the European Commission) and the case law of the Court of Justice. Each item has a permanent CELEX number, for example 32016R0679 for the General Data Protection Regulation (GDPR).

What can you use this source for?

- Legal: querying a regulation or directive article by article alongside a case file, or adding a judgment of the Court of Justice.
- Government and policy: finding out which European rules lie behind a national law, or following a Commission proposal.
- Business and compliance: collecting the European requirements for a product or service, in your customers' language.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **EUR-Lex (EU law)**. The fields:

- **Search words or CELEX number** (required): words from the **title**, for example *data protection*. Each word must occur in the title; at most eight words are used, and only letters, digits and hyphens. Or a single CELEX number, such as 32016R0679 or 62014CJ0362, or a link from eur-lex.europa.eu containing a CELEX number.
- **Type**: **Legislation**, **Case law (Court of Justice)** or **Preparatory acts**. Empty = all types, so also consolidated texts, for example.
- **Language**: **Dutch**, **English**, **German**, **French**, **Spanish**, **Italian** or **Portuguese**. The titles are searched, and the text retrieved, in that language. If you choose nothing, Dutch is used; with a pasted link, the language of the link applies. Choose the language that matches your search words.

Then choose **Search**. The most recent items are listed first, 20 at a time; **More results** shows the next ones. This source does not give a total number of hits. Searching the full text is not possible. Tip: use distinctive words from the official title; very general words may make the search service take too long.

What you get

Tick the items you want (at most 20 at a time) and choose **Add to sources**. Retrieval runs in the background (Task & Process Monitor) and may take a while for large acts; at the check, the GDPR gave 66 sections in just over 20 seconds. TrevantAI retrieves the text in the chosen language and makes each article a section, with the breadcrumb trail in the heading (title › chapter › article). Tables become lines of text.

The following are recorded with the document: the CELEX number, the type, the date, the language and the permanent link at the Publications Office, together with the link to EUR-Lex in that language. The **Public source** card also shows **Retrieved on** and **Found with**. The document is a **snapshot**.

Key

No key needed. The administrator only has to switch the source on under **Administration > Source connections**.

Privacy

The search query is sent to the Publications Office of the European Union. Do not enter client names or other confidential information. Legislation contains no personal data, but judgments and opinions of the Court name the parties, sometimes private individuals. Because the setting applies to the whole source, items from EUR-Lex are processed by the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

Reuse of EUR-Lex is permitted, including for commercial purposes, with acknowledgement of the source (Commission Decision 2011/833/EU). TrevantAI stores the attribution **© Europese Unie, https://eur-lex.europa.eu** (© European Union) with the document. Only the texts published in the Official Journal of the EU are authentic; if in doubt, use that version.

Good to know

- Searching finds whole words in the titles. If that finds nothing, TrevantAI also searches for parts of words; that takes up to twenty seconds.
- The same act in another language is a separate document. This lets you query, for example, the Dutch and the English text side by side.
- Not every item is available as text; sometimes there is only a PDF. TrevantAI then reports that there is no readable text; **View at the source** opens the item on EUR-Lex.
- If an item is already in the sources, **Already in the sources** is shown, with **View**.
- Traceable: the CELEX number and the language are fixed, and each fragment carries the breadcrumb trail.

GB legislation.gov.uk (United Kingdom legislation)

The official legislation of the United Kingdom: find an Act or Statutory Instrument by title, type or year and add the current text, with one section per provision.

> Key: none · Type of use: Free to use (Free to use with attribution (Open Government Licence v3.0)) · Personal data: no · Where: at the source

legislation.gov.uk is the official website for the legislation of the **United Kingdom**. It is managed by **The National Archives**. It contains Acts of the UK Parliament, *Statutory Instruments* and the legislation of Scotland, Wales and Northern Ireland. For many Acts there is a revised version in which later amendments have already been incorporated. TrevantAI uses the site's official data services: a search feed (Atom) and the texts as XML.

What can you use this source for?

- Legal: cite a UK Act or regulation in advice, for example for a contract under English law or a dispute with a UK party.
- Government and policy: compare UK rules with Dutch or European ones, for example after Brexit.

- Business: check which UK rules apply to employers, products or data.
- Research and journalism: compare the current text of an Act with earlier versions.

Searching

- **Words from the title or reference**: words from the title, for example *Equality Act*. You can also enter a reference such as *ukpga/2010/15*, or paste a link to legislation.gov.uk. TrevantAI then retrieves exactly that piece of legislation.
- **Type**: for example a *UK Public General Act*, a *UK Statutory Instrument*, Scottish, Welsh or Northern Irish legislation, or a regulation originating from the EU.
- **Year**: for example *2010*.

Enter words, optionally with type and year. Without words, choose a type and a year to see all legislation of that year. You get 20 results at a time; use **More results** to continue. Each result shows the title, the type, year and number and a short description.

What you get

Tick a result and choose **Add to sources**. TrevantAI retrieves the current version. The item has these parts, each with the breadcrumb in its heading:

- **Gegevens** (details): title, type, reference, date of enactment or making, the version (revised or original), the date from which the text applies and the link to the website.
- **Inleiding** (introduction): the long title, the number and the enacting formula.
- **One section per provision** (*Section*, *Regulation*, *Article* or *Paragraph*), with the structure above it, for example *Equality Act 2010 › Part 2 – Equality: key concepts › Section 4 – The protected characteristics*. Subsections and paragraphs appear as (1), (a) and so on. Below each provision the extent is shown, for example *E+W+S* (England, Wales and Scotland).
- **Schedules**, per paragraph.
- **Bron** (source): the attribution and the licence.

Amending text that changes another Act stays in the provision where it appears. TrevantAI leaves out the site's editorial annotations.

Key

No key is needed. The administrator switches the connection on under **Administration > Source connections**.

Privacy

The search query is sent to legislation.gov.uk. Do not use client names or other confidential information in the query. Legislation normally contains no personal data. That is why, by default, an external AI may also process these items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

The texts are covered by the **Open Government Licence v3.0**, except where otherwise stated. You may use them freely, including commercially, with attribution. With the document, TrevantAI keeps the statement "Contains public sector information licensed under the Open Government Licence v3.0". Legislation

originating from the EU may be reused under Decision 2011/833/EU, also with attribution.

Good to know

- A revised version is not always complete: sometimes recent amendments have not yet been incorporated. This is then stated under **Gegevens**; check on the website.
- Old Acts (before 1963) have a regnal year in the reference, for example *ukpga/Geo5/15-16/20*.
- The site allows at most 1,500 requests per 5 minutes. If the source asks you to wait, try again a little later.
- The labels in the item are in Dutch; the legislation is in English (sometimes also Welsh).
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.
- Retrieving often takes 5 to 15 seconds; a whole act can be very large (hundreds of thousands of characters).
- Under **View** you see the source, the reference, the link, the retrieval date and the search query (**Found with**).

Government and politics

🏛️ Council information (OpenBesluitvorming)

Documents, meetings and motions of Dutch municipal councils, provincial councils and water authorities via OpenBesluitvorming.nl (Open State Foundation): search by words, organisation, type and period, and add council documents as sources.

> Key: none · Type of use: Free to use (Free to use (CC0)) · Personal data: may contain · Where: at the source

Council information comes from **OpenBesluitvorming.nl**, an open collection of the Open State Foundation with the documents of Dutch municipal councils, provincial councils (Provinciale Staten) and water authorities: council proposals, letters, minutes, meeting agendas and motions. OpenBesluitvorming is the successor of Open Raadsinformatie, whose old API stops on 1 November 2026. At the last check, the source covered more than 330 organisations. Which organisations are included, and from when, is decided by OpenBesluitvorming.

What can you use this source for?

- Journalism: finding out what a municipal council discussed and decided on a subject, for example wind turbines or a building project.
- Government and policy: comparing how different municipalities or provinces deal with the same issue.
- Legal: collecting the council decisions and documents behind a by-law or a change of zoning.
- History and science: following decision-making on a local subject over several years.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **Council information (OpenBesluitvorming)**. The fields:

- **Search words** (required): one or more words, for example *windmolens* (wind turbines). OpenBesluitvorming searches the documents themselves. The documents are in Dutch, so search with Dutch words.
- **Organisation**: a list supplied by the source, grouped into **Municipalities**, **Provinces** and **Water authorities**. Empty = all organisations.
- **Type**: **Documents**, **Meetings** or **Motions**. Empty = all types.
- **From** and **Up to and including**: the period.

Then choose **Search**. The results are sorted by relevance, best match first, 20 at a time; **More results** shows the next ones. Each result shows the organisation, the type, the date and a snippet of text. Tip: choose an organisation and a period; a general word across all organisations quickly produces thousands of hits.

What you get

Tick the items you want (at most 20 at a time) and choose **Add to sources**. Retrieval runs in the background (Task & Process Monitor). TrevantAI takes the text as OpenBesluitvorming has extracted it from the document and divides it by heading, with the breadcrumb trail in the heading (organisation › title › heading). Without headings, the text is cut into parts of about 6,000 characters.

A **meeting** has no text of its own. That is why the agenda becomes the text: each agenda item a section, with the names of the documents that belong to it, plus the motions. The documents themselves are not retrieved automatically; look them up separately if needed.

The following are recorded with the document: the identifier at OpenBesluitvorming, the organisation, the type, the date, the meeting the document belongs to and the link to the original file. The **Public source** card also shows **Retrieved on** and **Found with**. The document is a **snapshot**.

Key

No key needed. The administrator only has to switch the source on under **Administration > Source connections**.

Privacy

The search query is sent to OpenBesluitvorming (Open State Foundation). Do not enter client names or other confidential information. Council documents contain names, for example of council members, people who address the council and applicants. That is why items from this source are processed by the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The OpenBesluitvorming data are free to use (CC0). The GDPR still applies to the names of people in the documents. TrevantAI stores the attribution **OpenBesluitvorming (Open State Foundation)** with the document.

Good to know

- The names in the **Organisation** list come from the source and are in Dutch.
- Searching by identifier is not possible here; search by words and narrow down by organisation and period.
- OpenBesluitvorming asks for fair use. If you search in quick succession, the source may ask you to wait a moment; then try again in a minute.
- The text has been extracted from the original files; the layout of tables may be lost. If in doubt, open the original via **View at the source**.
- If an item is already in the sources, **Already in the sources** is shown, with **View**.

🏛 House of Representatives (parliamentary papers)

Motions, amendments, letters to parliament, parliamentary questions, bills and reports from the data warehouse of the Dutch House of Representatives: search by words from the subject, document type and period, and add the document with the text of the PDF as a source.

> Key: none · Type of use: Free to use (Free to use (CC0)) · Personal data: may contain · Where: at the source

House of Representatives (parliamentary papers) retrieves documents from the **Gegevensmagazijn** (data warehouse) of the Tweede Kamer der Staten-Generaal, the Dutch House of Representatives, via its open data portal (opendata.tweedekamer.nl). It contains, among other things, motions, amendments, letters from the

government, written questions and the answers to them, bills, explanatory memoranda, memoranda of amendment, reports of committee debates and written consultations, and reports. Each document has a permanent identifier (Id) and usually a document number.

What can you use this source for?

- Government and policy: collecting the motions and letters to parliament on a dossier, for example on nitrogen (stikstof), and querying them together.
- Legal: looking up the explanatory memorandum and the memoranda of amendment for an act, to interpret an article.
- Journalism: finding out which written questions have been asked on a subject and what the government answered.
- Science: following the parliamentary debate on a subject over several years.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **House of Representatives (parliamentary papers)**. The fields:

- **Search words** (required): words from the **subject** or **title** of the document, for example *stikstof*. Each word must occur; at most eight words are used. The text of the documents themselves is not searched. The documents are in Dutch, so search with Dutch words. You can also paste the identifier (Id) of a single document, or a link containing that identifier.
- **Document type**: including **Motion**, **Amendment**, **Letter from the government**, **Written questions**, **Answer to written questions**, **Bill**, **Explanatory memorandum**, **Memorandum of amendment**, **Report of a committee debate**, **Report of a written consultation** and **Report**. Empty = all types.
- **From** and **Up to and including**: the period, based on the date of the document.

Then choose **Search**. The most recent documents are listed first, 20 at a time, with the total number of hits; **More results** shows the next ones. Each result shows the subject, the type, the document number and the parliamentary year. Deleted documents are not shown. Tip: start with one distinctive word and a document type, then refine with a period.

What you get

Tick the documents you want (at most 20 at a time) and choose **Add to sources**. Retrieval runs in the background (Task & Process Monitor). Each document first gets a **Gegevens** (details) section: type, title, subject, document number, parliamentary year, date and the source. If the file is a PDF, the text follows **page by page**, with the breadcrumb trail in the heading (subject › Pagina 3).

- A PDF without readable text, such as a scan: only the details, with the note that the document can be opened via the link.
- A file other than a PDF: only the details; TrevantAI does not read that file here.
- No file: only the details.

The following are recorded with the document: the identifier, the type, the document number, the parliamentary year, the date, the file type and whether the full text was included. The **Public source** card also shows **Retrieved on** and **Found with**. The document is a **snapshot**.

Key

No key needed. The administrator only has to switch the source on under **Administration > Source connections**.

Privacy

The search query is sent to the data warehouse of the House of Representatives. Do not enter client names or other confidential information. Parliamentary papers name members of parliament and ministers, and letters from citizens and reports of hearings contain names of private individuals. That is why items from this source are processed by the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The open data of the House of Representatives are free of rights (CC0). The GDPR still applies to the names of people. TrevantAI stores the attribution **Tweede Kamer der Staten-Generaal (opendata.tweedekamer.nl)** with the document.

Good to know

- Parliamentary papers are also in the **Official publications** source (gazette **Parliamentary papers**), where you can search for words in the text.
- Retrieving a large PDF may take a while.
- If a document is already in the sources, **Already in the sources** is shown, with **View**.
- Traceable: the identifier and the document number are permanent; **View at the source** opens the document on tweedekamer.nl.

EU European Parliament (documents)

Adopted texts, reports, motions for resolutions and parliamentary questions from the European Parliament's Open Data Portal, with the full text in the chosen language where the Parliament provides it.

> Key: none · Type of use: Free to use (Free to use with attribution (CC BY 4.0)) · Personal data: may contain · Where: at the source

The **European Parliament** publishes its documents as open data in its **Open Data Portal** (data.europarl.europa.eu). TrevantAI uses the portal's official API (version 2). With it you retrieve adopted texts, committee reports, motions for resolutions and parliamentary questions. A document comes with its details and, where the Parliament provides a file, with the full text in one of TrevantAI's seven languages.

What can you use this source for?

- Government and policy: check what Parliament adopted on a subject, for example its first-reading position on a Commission proposal.
- Legal: compare the text of a legislative resolution or amendments with the final EU legislation (EUR-Lex).
- Journalism: search questions by Members of the European Parliament and the reports of a given year.
- Research: follow European decision-making over a period of time.

Searching

- **Search words or document number**: words such as *climate*, or a number such as *TA-9-2024-0123*. A link to a document on europarl.europa.eu also works; TrevantAI then retrieves exactly that document.
- **Type**: adopted texts (default), reports, motions for resolutions, joint motions for resolutions, written questions or oral questions.
- **Search in**: the title (default) or the full text. This only works for adopted texts.
- **Year**: for example *2024*.
- **Sitting from** and **Sitting up to**: a range of plenary sittings, for adopted texts only.
- **Language**: the language of the titles and the text. The words are searched in that language.

Important: only for **adopted texts** can the Parliament itself search by words. For the other types the source can only search by year and type, so choose a year. TrevantAI retrieves twenty documents per page and shows only those whose title contains your words. Use **More results** to continue. This keeps the number of requests limited.

What you get

Tick a result and choose **Add to sources**. The item is called *reference – title*, for example *T9-0123/2024 – Legislative resolution …*. Each section has the breadcrumb in its heading:

- **Gegevens** (details): title, document number, reference, type, date, the procedure (for example *2021-0211*), the language, a link to the reading version and the retrieval date.
- **The text**: where the Parliament provides an HTML file, in sections per heading (for example *› Article 1 – Subject matter*); otherwise from the PDF file, per page.
- **Bron** (source): the attribution and the licence.

If the Parliament does not (yet) provide a file in the chosen language, this is stated under **Gegevens** and you only get the details with the link.

Key

No key is needed. The administrator switches the connection on under **Administration > Source connections**.

Privacy

The search query is sent to the European Parliament. Do not use client names or other confidential information in the query. The documents contain names of Members of the European Parliament, such as authors, rapporteurs and those asking questions, and sometimes of other people. That is why, by default, only the local AI processes these items. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. TrevantAI does not copy contact details.

Licence and attribution

The Parliament allows reuse of its open data and documents under **CC BY 4.0**, including commercially, with attribution. With the document, TrevantAI keeps the attribution "© Europese Unie – Bron: Europees Parlement" (© European Union – Source: European Parliament). Logos and material from others (such as photographs) are not covered by this licence.

Good to know

- The source allows 500 requests per 5 minutes. Searching reports or questions takes more requests, because titles are retrieved per document. If the source asks you to wait, try again a little later.
- New documents sometimes appear without a title or text at first; try again later.
- Adopted texts are not the final EU legislation. You find that in EUR-Lex.
- The labels in the item are in Dutch; titles and text are in the chosen language.
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.
- For reports, resolutions and questions, TrevantAI only searches the words in the titles of the page shown. If nothing is there, **More results** helps. That is about twenty requests to the source per page, so it takes a while.
- Under **View** you see the source, the document number, the link, the retrieval date and the search query (**Found with**).

Spatial planning and environment

🗺️ Spatial plans

Look up a zoning plan or other spatial plan through the Dutch government's Spatial Plans API with a link or plan number and add its rules per article and its explanatory statement to your sources.

> Key: required (own key) · Type of use: Free to use (Public government information) · Personal data: no · Where: at the source

Spatial plans are the plans in which the Dutch authorities set out what is allowed where: above all **zoning plans** (bestemmingsplannen), but also structure visions, management ordinances and other plans published on **ruimtelijkeplannen.nl**. TrevantAI retrieves them through the official **Ruimtelijke Plannen API** (Spatial Plans API, version 4) of the **Digitaal Stelsel Omgevingswet (DSO)**, the digital system of the Dutch Environment and Planning Act, managed by the government (Informatiehuis Ruimte, Kadaster). One plan becomes one source, with the **rules per article** and the **explanatory statement**. Since the Environment and Planning Act took effect (1 January 2024), they form part of the municipality's environment plan (omgevingsplan).

What can you use this source for?

- **Legal**: checking which land use and building rules apply to a plot, in an objection, an enforcement case or a purchase agreement.
- **Government and policy**: quoting or comparing the rules of a plan for a new policy proposal.
- **Business and real estate**: investigating the building and use options of a site before a purchase or investment.
- **Journalism**: showing exactly what a controversial plan allows, with the article cited.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **Spatial plans**. This source has its own screen: a field for a link or plan number and a button to look up the plan.

- Paste a **link** to the plan from Planviewer, ruimtelijkeplannen.nl or the Omgevingsloket (the Dutch environment portal). TrevantAI takes the plan number from the link.
- Or enter the **plan number** (IMRO code). It has the form *NL.IMRO.[code].[plan]-[version]*: the code is that of the authority (for a municipality, the municipal code), and the letters in the plan part often indicate the type (*BP* for bestemmingsplan, zoning plan). The version shows whether it is a draft (*ON*…) or an adopted plan.

TrevantAI then shows the plan: name, type, status and date, the authority and the plan number, and, where relevant, that since 1 January 2024 the plan forms part of the environment plan or no longer applies. Choose **Add to sources**; retrieval runs in the background (Task & Process Monitor). If a project is active, the plan is linked to it. If you paste a plan link when importing a web page, TrevantAI also retrieves the plan through the API.

Searching by words or by address is not (yet) possible.

What you get

- Each **article** (and every other text with content) becomes a separate section. A text fragment never spans two articles and always starts with the breadcrumb trail, for example *Regels › Hoofdstuk 2 › Artikel 3 Wonen › 3.1 Bestemmingsomschrijving* (Rules › Chapter 2 › Article 3 Residential › 3.1 Description of use).
- The **explanatory statement**: if it is not among the texts of the API but on a separate page on ruimtelijkeplannen.nl, TrevantAI reads that page; if that fails, adding continues with the rules.
- The **plan data**: plan number, name, type, status and date, authority, end of legal validity and the date of retrieval; shown under **View**.

The map (verbeelding) and the annexes are not included. The item is a **snapshot**: a later version of the plan is not retrieved automatically. At most 2,000 texts are read per plan.

Key

The Ruimtelijke Plannen API requires your organisation's **own key**: free, to be requested on the Ontwikkelaarsportaal Omgevingswet (the developer portal of the Environment and Planning Act). TrevantAI does not supply any keys. The administrator enters it under **Administration > Source connections**, where the connection itself must also be switched on. Without a key, the screen tells you what to do. How to request the key is explained in the Help Centre topic **Key for the Spatial Plans API**.

Privacy

The link or plan number is sent to the Ruimtelijke Plannen API. Spatial plans are public government documents. That is why the same defaults apply to these items as to web pages: an external AI may process them if your organisation's AI policy allows cloud AI. This source has its own import and does not (yet) follow the choice **Who may process items from this source?** in **Administration > Source connections**; it can, however, be set for the source type **Public sources (source connections, such as case law and spatial plans)** and per document. Questions about the items always go to the local AI.

Licence and attribution

Spatial plans are **public government information**. You may quote from them and use them in your work, provided you credit the source. Requests are subject to the fair use policy of the Ontwikkelaarsportaal Omgevingswet; TrevantAI only requests what you look up or add. With the document, TrevantAI stores the attribution **Ruimtelijke Plannen API (DSO)**.

Good to know

- Only the plans on ruimtelijkeplannen.nl are available. The new rules of the environment plan under the Environment and Planning Act come through another API and are not (yet) included here; please consult the Omgevingsloket for those.
- Not every plan provides its texts through the API. For some plans they exist only as separate files; TrevantAI then reports that the plan has no text. A plan that does not (or no longer) exist produces the message that it was not found; in that case, look for the current version in the plan's file.
- If the API asks you to wait, please try again a little later.
- Adding the same plan again does not create a copy: the existing item is used and, if needed, linked to the project.

🏠 BAG (Dutch addresses and buildings)

Look up an address in the Kadaster's Basic Registration of Addresses and Buildings and add the unit and building (year of construction, floor area, purpose of use, status) to your sources as a snapshot.

> Key: required (own key) · Type of use: Free to use (Public domain (Public Domain Mark 1.0)) · Personal data: may contain · Where: at the source

The **Basic Registration of Addresses and Buildings (Basisregistratie Adressen en Gebouwen, BAG)** is the official register of all addresses and buildings in the Netherlands. Municipalities maintain the data; the **Kadaster** (the Dutch land registry) manages the national facility. For each address it records which unit it concerns (a dwelling or other unit, a pitch or a mooring), with status, purpose of use and floor area, and which building it is in, with the original year of construction. TrevantAI retrieves the data through the Kadaster's **BAG API Individuele Bevragingen** and adds an address as a **snapshot**.

What can you use this source for?

- **Legal**: in a rental or purchase dispute, checking which purpose of use and which floor area are registered for a unit.
- **Government and policy**: the year of construction and status of buildings for a sustainability or redevelopment plan.
- **Business and finance**: checking a business address or collateral (does the address exist, is it in use, what is the floor area).
- **Heritage and history**: the year of construction of a building as a starting point for historical research.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **BAG (Dutch addresses and buildings)**. There are two ways:

- **House number** with **Postcode** (for example *701* and *7334 DP*; the postcode may be written with or without a space), or **House number** with **Street** and **Town** (for example *Laan van Westenenk* and *Apeldoorn*). Enter the street and town in full.
- **House letter** (one letter) and **Addition** (at most four letters or digits) are optional. With a house letter or addition, TrevantAI looks for exactly that address; without them you get all addresses at that house number (such as 12, 12A, 12-1), so that you can choose.
- **Or: address ID (nummeraanduiding)**: the BAG address number, 16 digits, for example *0200200000075716*. You may also paste a link or text containing that number, such as a link from the BAG Viewer.

Choose **Search**. You see at most 50 addresses, with the type of unit, the purpose of use, the floor area, the year of construction and the status. There is no **More results**: narrow the query if needed. **View at the source** opens the address in the Kadaster's BAG Viewer. Tick the addresses and choose **Add to sources**.

What you get

Each address becomes one source, titled *BAG: street house number, postcode town*, with these parts:

- **Adres** (address): street, house number, postcode, town, the address ID and the IDs of the public space and the town; and whether it is a secondary address.

- **Verblijfsobject** (unit), or **Standplaats** (pitch) or **Ligplaats** (mooring): the ID, the status, the purpose of use, the usable floor area, the location in RD coordinates and whether the unit was *geconstateerd* (recorded without a permit or notification).
- **Pand** or **Panden** (building or buildings, at most five): building ID, original year of construction, status, the source document (number and date) and what is under investigation.
- **Bron** (source): which data the municipality is investigating, and the attribution.

Every text fragment starts with the breadcrumb trail, such as *Laan van Westenenk 701, 7334 DP Apeldoorn › Pand*. TrevantAI also records the address ID (identifier), the link to the BAG Viewer, the date of retrieval and the search query. Not included are the shape of the building on the map and the history of the data.

Key

The BAG requires your organisation's **own key**. It is free and is requested from the Kadaster; TrevantAI does not supply any keys. The administrator enters it under **Administration > Source connections**. How to request the key is explained in the Help Centre topic **Key for the BAG API**.

Privacy

The address you search for is sent to the Kadaster. Please do not search for clients' addresses unless necessary. The BAG itself contains no names, but an address in an inquiry into a person does become personal data. That is why items from the BAG are processed **by the local AI only** by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The data from the BAG are covered by the **Creative Commons Public Domain Mark 1.0**: they are in the public domain and free to use, without conditions. TrevantAI nevertheless stores the attribution **Bron: BAG, Kadaster** (Source: BAG, Kadaster) with the document, so that it is always clear where the data come from.

Good to know

- The data are those of the day of retrieval. Municipalities update the BAG continuously; the current state is shown in the BAG Viewer. If an address is already in the sources, you see **Already in the sources** and the existing item is used.
- *In onderzoek* (under investigation) means that the municipality doubts a piece of data; use it with care.
- The floor area is the usable floor area of the unit, not that of the plot. The year of construction is the original year of construction of the building.
- An address can lie in more than one building; TrevantAI includes at most five.

🏞️ Environmental documents (Environment and Planning Act, DSO)

Environmental plans, visions, regulations and programmes under the Dutch Environment and Planning Act from the DSO: search by title or authority and add the regulation in force with articles, table of contents and breadcrumb.

> Key: required (own key) · Type of use: Free to use (Public government information) · Personal data: no · Where: at the source

The Dutch **Environment and Planning Act** (Omgevingswet) has applied since 1 January 2024. Municipalities, provinces, water authorities and central government publish their rules for the physical living environment in the **Digital System for the Environment and Planning Act (DSO)**: the municipal environmental plan, the environmental vision, the provincial environmental regulation, the water authority regulation and programmes. The Omgevingsloket shows them under "Regels op de kaart". TrevantAI uses two official DSO APIs: **Omgevingsinformatie Ontsluiten** to search and **Omgevingsdocumenten Presenteren (Ozon)** for the text.

What can you use this source for?

- Government and policy: read and compare the rules of your own or another environmental plan.
- Legal: record the rules in force for a permit application, objection or advice, with article numbers.
- Journalism: check what a municipality says about housing, green space or energy in its environmental vision.
- Research: compare environmental visions or regulations of several authorities.

Searching

- **Title, municipality, province or water authority** (required): for example *Omgevingsplan gemeente Utrecht* or *omgevingsverordening Gelderland*. The search covers the title and the identifier. You can also enter an identifier (*/akn/nl/act/gm0344/2020/omgevingsplan*) or a link from the Omgevingsloket.
- **Level of government** (choice): municipality, province, water authority or central government.
- **Type of document** (choice): environmental plan, environmental vision, environmental regulation, water authority regulation or programme. TrevantAI filters on this after searching; a page may then show fewer than 20 results.
- **Chapter only (number)**: for example *3*. TrevantAI then retrieves only that chapter. Useful for large environmental plans.

Choose **Search**. Each result shows the type, the competent authority and the date from which the regulation applies. Use **More results** to see more. Old zoning plans (IMRO) are not found here; use the **Ruimtelijke Plannen** connection for those.

What you get

Tick a regulation and choose **Add to sources**. The item is a **snapshot** of the regulation in force on that day. Each section has the breadcrumb in its heading, for example *Omgevingsplan gemeente Utrecht › Hoofdstuk 3 Activiteiten › Afdeling 3.1 … › Artikel 3.1 Bouwhoogte*:

- **Gegevens** (details): title, type, competent authority, identifier, version, date in force and the link to the publication.
- **Inhoudsopgave** (table of contents): all chapters, divisions, articles and annexes.
- **Articles**: one section per article, with its paragraphs, lists and tables. For an environmental vision or programme, one section per text block.
- **Bron** (source): origin and where to find the current text.

Environmental plans with the temporary part (the "bruidsschat") are very large. TrevantAI therefore includes at most 600 articles and text blocks and 1.5 million characters. If there is more, the item states which part is missing. Retrieve that part separately with **Chapter only**. The table of contents always shows the whole regulation. Maps, areas of application and images are not included.

Key

An own DSO key is required. The guide is under **Key for environmental documents (DSO)**. The administrator enters the key under **Administration > Source connections**.

Privacy

The search query is sent to the DSO. Environmental documents contain rules, not data about individuals. Therefore an external AI may also process these items by default, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

Environmental documents are **public government information** and may be reused freely. TrevantAI keeps the attribution "Bron: Digitaal Stelsel Omgevingswet (Ozon)" with the document.

Good to know

- The item shows the regulation as it applies on the day of retrieval. If the municipality changes the plan, add it again.
- Drafts and environmental permits are not retrieved.
- Whether a rule applies at a particular location depends on the areas on the map. Check this in the Omgevingsloket.
- The text of the regulations is in Dutch.
- **View** shows the source, the link to Regels op de kaart, the date of retrieval and the search (**Found with**).

Environment and climate

🌦️ KNMI (weather data per station)

Add validated daily values from a KNMI weather station (temperature, precipitation, sunshine, wind) over a chosen period to your sources as a table.

> Key: required (own key) · Type of use: Free to use (Free to use with attribution (CC BY 4.0)) · Personal data: no · Where: at the source

The **KNMI** (Koninklijk Nederlands Meteorologisch Instituut, the Royal Netherlands Meteorological Institute) measures the weather at weather stations throughout the Netherlands. Through the **EDR API** of the **KNMI Data Platform**, TrevantAI retrieves the **validated daily values** of one station: temperature, precipitation, sunshine and wind. Validated means that the KNMI first checked and completed the hourly values before summarising them into daily values. One station over a chosen period is added to your sources as a **snapshot** (table).

What can you use this source for?

- **Legal and insurance**: in a damage claim, checking how strong the wind was or how much rain fell on the day of the damage.
- **Business**: substantiating delays in construction or a poor harvest with the precipitation and temperature of those weeks.
- **Journalism**: checking a report about heat, drought or a storm against the figures of the nearest station.
- **Science and health care**: setting weather data alongside other data, for example the temperature during a period with many hospital admissions.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **KNMI (weather data per station)**.

- **Station (name or number)** (required): the name or part of it, for example *De Bilt*, or the station number, for example *6260*. Capitals do not matter.
- **From date** and **Up to and including date** (required): the period, as YYYY-MM-DD. One item covers **at most 366 days**; the end date may not be before the start date.
- **Quantities**: **Temperature (mean, minimum, maximum)**, **Precipitation**, **Sunshine**, **Wind** or **All four**. Without a choice, you get all four.

Choose **Search**. You see the matching stations, at most 20 at a time, sorted by name, with the station number, the quantities that will be retrieved and the location. Use **More results** to see further results. Tick one or more stations and choose **Add to sources**; each station becomes a separate item.

Tips: if you need more than a year, add several periods. Choose the station closest to the place you are investigating; the location is shown with each result.

What you get

Each item is titled, for example, *KNMI De Bilt: dagwaarden 2026-09-01 t/m 2026-09-30* (daily values from … to …) and has these parts:

- **Beschrijving** (description): the station with its number, the location (and altitude if the KNMI provides it), the period with the number of days with values, the description and unit of each column, the collection, the date of retrieval, and an explanation of the validation and the limitations.
- **Gegevens (deel 1, 2, …)** (data, part 1, 2, …): a table with one row per day, the date and one column per quantity with its unit, 120 days per part.
- **Bron** (source): the attribution and the licence.

The columns are named as the KNMI names the quantities (codes such as TG, TN and TX for the mean, lowest and highest temperature); the description explains each column. Every text fragment starts with the breadcrumb trail, such as *KNMI De Bilt: dagwaarden … › Gegevens (deel 1)*. The identifier records the whole query (station, start, end and quantity), together with the link, the date of retrieval and the search query.

Key

The KNMI requires your **own key** for the EDR API. It is free and is requested from the KNMI Data Platform; TrevantAI does not supply any keys. The administrator enters it under **Administration > Source connections**. How this works, and what can be done with the KNMI's public anonymous key, is explained in the Help Centre topic **Key for the KNMI EDR API**.

Privacy

The station and period you search for are sent to the KNMI. Please do not enter client names or other confidential information. Measurements from weather stations contain no personal data. That is why, by default, an external AI may also process the items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The KNMI makes the data available under **Creative Commons Attribution 4.0 (CC BY 4.0)**: you may use and share them freely, including in publications, provided you credit the source. TrevantAI stores the attribution **Bron: KNMI** (Source: KNMI) with the document.

Good to know

- Validated daily values are only available after **1 to 10 working days**. The most recent days are therefore often missing; if the station has no values for the whole period, TrevantAI says so.
- Validated values may still be adjusted later. The snapshot does not change with them.
- The KNMI describes these series as **not homogeneous**: stations have been moved and measuring techniques have changed. Do not use them for long-term trend analysis without correction.
- Dates are shown as YYYY-MM-DD, numbers in Dutch notation (decimal comma).
- This connection has not yet been tested against the real EDR API. In particular, the column names may differ from what is described here; the description of each item gives the correct explanation.
- An item with the same station, the same period and the same quantities that is already in the sources is shown as **Already in the sources**.

Luchtmeetnet (air quality per monitoring station)

Add hourly air quality values from a Dutch monitoring station (nitrogen dioxide, particulate matter, ozone, ammonia, sulphur dioxide) over at most 31 days to your sources as a table.

> Key: none · Type of use: Free to use (Public government information) · Personal data: no · Where: at the source

Luchtmeetnet (the Dutch air quality monitoring network) publishes the air quality measurements for the Netherlands. It is a partnership of the **RIVM** (National Institute for Public Health and the Environment), the **GGD Amsterdam** (public health service), **DCMR Milieudienst Rijnmond** (environmental agency) and the **provinces**, which together operate monitoring stations throughout the country. For each station and each hour, it records how much of a substance was in the air. Through the Luchtmeetnet open API, TrevantAI retrieves the **hourly values** of one station and one substance and adds them to your sources as a **snapshot** (table).

What can you use this source for?

- **Legal**: in an objection to a permit or a zoning change, showing the measured concentrations of nitrogen dioxide or ammonia nearby.
- **Government and policy**: checking whether a low-emission zone or a traffic measure is visible in the measurements.
- **Journalism**: checking a report about smog or poor air near a school against the figures of the nearest station.
- **Medical and science**: setting air quality alongside complaints or admissions, for example asthma during a period with high levels of particulate matter.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **Luchtmeetnet (air quality per monitoring station)**.

- **Place or monitoring station** (required): (part of) the name of the place or the monitoring location, for example *Amsterdam*, or the number of a station, for example *NL49007*. Capitals do not matter.
- **Substance** (required): **Nitrogen dioxide (NO_2)**, **Particulate matter (PM10)**, **Fine particulate matter (PM2.5)**, **Ozone (O_3)**, **Ammonia (NH_3)** or **Sulphur dioxide (SO_2)**.
- **From date** and **Up to and including date** (required): the period, **at most 31 days**.

Choose **Search**. You see the matching stations, at most 20 at a time, sorted by place, with the station number and the substance. Use **More results** to see further results. Tick the stations and choose **Add to sources**; each station becomes a separate item.

Tips: not every station measures every substance. The list shows all stations matching the place; if a station does not measure the chosen substance, adding it will say so. If you need more than a month, add several periods.

What you get

Each item is titled, for example, *Luchtmeetnet Amsterdam-…: NO2 2026-09-01 t/m 2026-09-07* (from … to …) and has these parts:

- **Beschrijving** (description): the monitoring station with its number, municipality and province, operator, type of station, in use since, location, all substances the station measures and the description of the station; the substance and the unit (µg/m³); the period with the number of hourly values, the mean, the lowest and the highest value; the date of retrieval and the warning that the measurements have not been validated.
- **Gegevens (deel 1, 2, …)** (data, part 1, 2, …): a table with the time and the value for each hour, one week (168 hours) per part.
- **Bron** (source): the attribution and the terms.

Every text fragment starts with the breadcrumb trail, such as *Luchtmeetnet … › Gegevens (deel 2)*. The identifier records the whole query (station, start, end and substance), together with the link, the date of retrieval and the search query.

Key

Luchtmeetnet does not require a key.

Privacy

The place, the substance and the period are sent to Luchtmeetnet. Please do not enter client names or other confidential information. Measurements from fixed monitoring stations contain no personal data. That is why, by default, an external AI may also process the items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

Luchtmeetnet makes the measurements available as **open data**: free to use, with attribution. TrevantAI stores the attribution **Bron: Luchtmeetnet (RIVM, GGD, DCMR, provincies)** (Source: Luchtmeetnet: RIVM, GGD, DCMR, provinces) with the document.

Good to know

- The measurements have **not been validated**. They may still be corrected or rejected later. The RIVM publishes validated series later (data.rivm.nl); use those for official conclusions.
- All times are in **UTC**, as the source gives them. Dutch time is one hour later in winter and two hours later in summer.
- Luchtmeetnet allows at most 100 requests per 5 minutes. TrevantAI stays well below that and waits if necessary; when it is busy, adding may therefore take longer.
- For each item, TrevantAI retrieves at most 60 pages of measurements. If the series is incomplete as a result, *onvolledig* (incomplete) appears in the summary and in the description; choose a shorter period in that case.
- Numbers are in Dutch notation (decimal comma). An item with the same station, the same period and the same substance that is already in the sources is shown as **Already in the sources**.

🌊 Rijkswaterstaat (water level, temperature, waves and discharge)

Add Rijkswaterstaat measurements at one monitoring location (water level, water temperature, wave height or discharge) over a chosen period to your sources as a table.

> Key: none · Type of use: Free to use (Free to use (CC0)) · Personal data: no · Where: at the source

Rijkswaterstaat (the Dutch Directorate-General for Public Works and Water Management) manages the main waters of the Netherlands and takes measurements at many places in rivers, lakes, canals and the North Sea. Through Rijkswaterstaat's **Waterwebservices** (the Data Distributielaag, the same data as on Waterinfo), TrevantAI retrieves the measurements of one **monitoring location** and one **quantity**: water level, water temperature, wave height or discharge. One location over a chosen period is added to your sources as a **snapshot** (table).

What can you use this source for?

- **Legal**: in a case of flood damage or a shipping accident, showing the water level on that day and at that hour.
- **Government and policy**: recording a period of drought or low water, for example the discharge of the Rhine at Lobith.
- **Business**: substantiating delays in inland shipping due to low water or in offshore work due to high waves.
- **Science and journalism**: setting water temperature and water levels alongside ecological data or a news report.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **Rijkswaterstaat (water level, temperature, waves and discharge)**.

- **Monitoring location** (required): (part of) the name of the location, for example *Hoek van Holland* or *Lobith*, or the code of the location. Capitals do not matter.
- **Quantity** (required): **Water level (relative to NAP)**, **Water temperature**, **Wave height (significant)** or **Discharge**.
- **From date** and **Up to and including date** (required): the period, **at most 366 days**.

Choose **Search**. You only see locations where the chosen quantity is measured, at most 20 at a time, sorted by name, with the code, the quantity, the unit and the reference (such as NAP, the Dutch ordnance datum). Use **More results** to see further results. Tick the locations and choose **Add to sources**; each location becomes a separate item.

Tips: if you cannot find a place, try a shorter name or another quantity; not every location measures everything. For detail, choose a short period: up to 31 days, you get all measurements.

What you get

Each item is titled, for example, *Rijkswaterstaat Hoek van Holland: waterhoogte 2026-09-01 t/m 2026-09-07* (water level from … to …) and has these parts:

- **Beschrijving** (description): the monitoring location with its code, the location with the coordinate system, the quantity, the unit, the reference and the compartment (for example surface water), the period with the number of measurements, the mean, the lowest and the highest value, how many measurements have which status, and the date of retrieval.
- **Gegevens (deel 1, 2, …)** (data, part 1, 2, …): for a period **of up to 31 days**, all measurements with time, value and status (144 per part); for a **longer period**, the daily mean, minimum, maximum and number of measurements (120 days per part).
- **Bron** (source): the attribution and the licence.

Every text fragment starts with the breadcrumb trail, such as *Rijkswaterstaat Hoek van Holland: waterhoogte … › Gegevens (deel 3)*. The identifier records the whole query (location, start, end and quantity), together with the link, the date of retrieval and the search query. Only actual measurements are included, no forecasts or astronomical tide. Measurements with quality code 99 (invalid) and missing values are left out.

Key

The Rijkswaterstaat data do not require a key.

Privacy

The location, the quantity and the period are sent to Rijkswaterstaat. Please do not enter client names or other confidential information. Measurements at fixed monitoring locations contain no personal data. That is why, by default, an external AI may also process the items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

Rijkswaterstaat publishes these open data under **CC0 1.0**: free to use, without conditions. TrevantAI nevertheless stores the attribution **Bron: Rijkswaterstaat** (Source: Rijkswaterstaat) with the document, so that it is always clear where the figures come from.

Good to know

- A place sometimes also lists sub-locations; these do not always have measurements for the chosen period. When in doubt, choose the main location.
- **Recent measurements** have often not yet been checked and may be corrected later. The description counts how many measurements have each status.
- How often measurements are taken differs per location and quantity; at Hoek van Holland there is a water level measurement every ten minutes. A long period with all measurements therefore becomes a large item; that is why TrevantAI summarises per day above 31 days.
- Times are shown as Rijkswaterstaat gives them, with the time zone. TrevantAI requests the period in Dutch standard time (UTC+1).
- If Rijkswaterstaat refuses a query because there are too many measurements, choose a shorter period.
- Numbers are in Dutch notation (decimal comma). An item with the same location, period and quantity that is already in the sources is shown as **Already in the sources**.

Statistics

📊 CBS StatLine (Statistics Netherlands)

Search the statistical tables of Statistics Netherlands (CBS) by words in the title and add a table to your sources as a traceable snapshot.

> Key: none · Type of use: Free to use (Free to use with attribution (CC BY 4.0)) · Personal data: no · Where: at the source

CBS StatLine is the database of statistical tables of **Statistics Netherlands (Centraal Bureau voor de Statistiek, CBS)**, the national statistics office of the Netherlands. The tables cover subjects such as population, economy, labour, housing, health and environment, often per period and per region. TrevantAI retrieves the tables through the CBS open data service (OData v4) and adds a table as a **snapshot**: the figures as they stood at CBS on the day of retrieval.

What can you use this source for?

- **Government and policy**: key figures on population or housing stock support a policy memo or council proposal.
- **Journalism**: checking a claim in the news against the official figures, with the table as a traceable source.
- **Business and finance**: economic figures for a market survey or a business plan.
- **Science**: background figures in a research report, with table number and date of retrieval.

Searching

Open **Sources > Explore & add sources > Public sources** and, under **Source**, choose **CBS StatLine (Statistics Netherlands)**.

- **Search words** (required): words from the title of the table, for example *bevolking kerncijfers* (population key figures). A table is found if all words occur in its title. Capitals do not matter, and part of a word also counts: *woning* also finds *woningvoorraad*. At most six words are used.
- **Tables**: choose **Current tables only** or **Discontinued tables only** (tables that CBS no longer updates). Without a choice, TrevantAI searches all tables.

Choose **Search**. You see at most 20 tables at a time, most recently changed first, with the table number, the status, the number of values and the date of the last change. Use **More results** to see further results; **View at the source** opens the table on opendata.cbs.nl. Tick the tables you want and choose **Add to sources**.

Tips: only the titles are searched, not the figures or the explanatory notes. Start with two or three distinctive words and add a word if too many tables appear. Words such as *kerncijfers* (key figures) appear in the titles of many overview tables.

What you get

Each table becomes one source, with these parts:

- **Beschrijving** (description): the CBS summary and explanation, the period, the licence and, for a shortened table, how many values have been included.
- **Onderwerpen** (subjects): the quantities in the table, with unit and description.

- **Gegevens (deel 1, 2, …)** (data, part 1, 2, …): each value as a readable line, for example *region · period: subject = value unit*, 120 lines per part. Empty cells are skipped.

Every text fragment starts with the breadcrumb trail, such as *title › Gegevens (deel 2)*. With the document, TrevantAI records: the table number (identifier), the link, the date of the last change at CBS, the period, the status, the number of values, the date of retrieval, the search query (**Found with**) and the licence. A table is included up to **20,000 values**; if it is larger, the description states that it has been shortened.

Key

CBS StatLine does not require a key.

Privacy

The search words are sent to CBS. Please do not enter client names or other confidential information there. The tables contain statistical figures and no data about individuals. That is why, by default, an external AI may also process the items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

CBS publishes the tables under **Creative Commons Attribution 4.0 (CC BY 4.0)**: you may use and share the figures freely, including in your own publications, provided you credit the source. TrevantAI stores the attribution **Bron: CBS** (Source: CBS) with the document; please include it when you quote figures.

Good to know

- A table per municipality or district is large; adding it can then take ten seconds or longer.
- The titles are in Dutch, so search with Dutch words.
- Very large tables are cut off after 20,000 values. In that case, look for a smaller table on the same subject or view the full table at the source.
- Numbers are written in Dutch notation (comma as decimal separator).
- The snapshot does not change when CBS later updates the table. If a table is already in the sources, you see **Already in the sources** with **View**; the existing item is then used.
- The table number, the link and the date of retrieval keep every figure traceable.

RIVM and other third parties (via CBS)

Statistical tables from the RIVM and other organisations, published by CBS: search by words from the title and add a table as a snapshot with description, topics and values.

> Key: none · Type of use: Free to use (Free to use with attribution (CC BY 4.0)) · Personal data: no · Where: at the source

Statistics Netherlands (**CBS**) publishes not only its own figures. A separate catalogue holds statistical tables **on behalf of other organisations** ("third parties"). The most important one is the **RIVM** (National Institute for Public Health and the Environment): health by district and neighbourhood, life expectancy, causes of death, vaccination coverage, food consumption and the cost of illness. You will also find the Youth Monitor,

the long-term care monitor, the labour market in care and welfare, police figures and the finances of municipalities, provinces and water authorities (IV3). TrevantAI uses the public OData service of this catalogue (dataderden.cbs.nl).

What can you use this source for?

- Medical and care: compare health, lifestyle and mortality in a region with the national picture.
- Government and policy: use figures per municipality, district or neighbourhood in a policy note or a health section.
- Journalism: check how vaccination coverage or life expectancy develops in a municipality.
- Research: record tables from the RIVM and the regional health services (GGD) as a traceable source for an analysis.

Searching

- **Words from the title or table number** (required): one or more words from the title of the table, for example *levensverwachting regio* or *vaccinatie*. Most titles are in Dutch. All words must appear in the title; capitals do not matter. You can also enter a table number such as *50152NED*, or a link containing that number.
- **Provider** (choice, optional): for example RIVM, Youth Monitor, long-term care monitor, Police or IV3. Leave it empty to search the whole catalogue.

Choose **Search**. Each result shows the table number, the provider, the period, the number of rows and the original source. The date is that of the last change. The newest tables come first. You get 20 results at a time; use **More results** to see more. Tip: many RIVM tables exist in several editions. Choose the newest unless you need an older regional division.

What you get

Tick a table and choose **Add to sources**. The item is a **snapshot** with the date of retrieval. Each section has the breadcrumb in its heading, for example *Title of the table › Gegevens (deel 2)*:

- **Beschrijving** (description): the explanation of the table, the period and the provider.
- **Onderwerpen** (topics): what was measured, with unit and explanation, and the breakdown (for example sex, age, region, period).
- **Gegevens** (data): the values as readable lines, such as *Mannen · 0 jaar · Waarde · Amsterdam · 2024: Levensverwachting = 80,1 jaren*. Each section holds at most 120 lines.
- **Bron** (source): the original source, the licence and where to find the current table.

Large tables (for example by district and neighbourhood) are **cut off** at 20,000 values. The item then states how many rows are included and how many the table has.

Key

No key is needed. The administrator only has to switch the connection on under **Administration > Source connections**.

Privacy

The search query is sent to CBS. The tables contain no data about individuals: they are figures about groups, and the source suppresses small numbers. Therefore an external AI may also process these items by default, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

Most tables may be used freely under **CC BY 4.0**, including commercially, with attribution. Name the original source (for example the RIVM) and CBS. TrevantAI keeps the attribution "Bron: RIVM of andere derde partij, via CBS" with the document and puts the exact source in the item. If a table only states copyright (©) and no CC BY 4.0, for example for the Police, that table is **not** under CC BY 4.0. TrevantAI records this per item, already says so in the search results and puts a clear sentence in the item. Quoting with attribution is then allowed; further reuse is subject to the source's conditions.

Good to know

- The own tables of CBS (StatLine) are in the **CBS StatLine** connection; this connection only covers third-party tables.
- Empty values and values the source keeps confidential are skipped.
- Figures from surveys (such as the Health Monitor) often come with a lower and an upper bound of the 95% interval. Use them when comparing.
- Municipal boundaries change through mergers. Note the year of the regional division in the title.
- The labels in the item are in Dutch; topic names appear as the source gives them. A few tables (for example the food consumption survey) are in English.
- **View** shows the source, the table number, the link, the date of retrieval and the search (**Found with**).

📈 Eurostat (European statistics)

The statistics of the European Union: find a dataset by words or code, choose country and period, and add the figures as a snapshot in table form.

> Key: none · Type of use: Free to use (Free to reuse with acknowledgement of the source (Decision 2011/833/EU)) · Personal data: no · Where: at the source

Eurostat is the statistical office of the **European Union**. It collects figures from the national statistical institutes (such as Statistics Netherlands, CBS) and makes them comparable. It offers thousands of datasets on the economy, population, labour, education, health, agriculture, energy, the environment and more, by country and often by region. TrevantAI uses two official Eurostat services: the table of contents (to search by words) and the statistics API (to retrieve the figures).

What can you use this source for?

- Government and policy: compare the Netherlands with other EU countries, for example unemployment, energy prices or health spending.
- Business and finance: support market data by country, such as GDP, inflation or trade, with official figures.
- Journalism: check a claim about "Europe" against the figures themselves.
- Research: record a series covering several years and countries as a source.

Searching

- **Search words or dataset code** (required): words from the name, for example *unemployment* or *gross domestic product*. The names are in English, French or German. You can also enter the code, for example *nama_10_gdp*; it is then listed first.
- **Country**: the EU as a whole or a country, such as the Netherlands or Belgium. Leave it empty for all countries and areas.
- **From year** and **Up to year**: for example *2015* and *2024*. Leave them empty for the last ten periods.
- **Language of the names**: English, French or German; this applies to the names in the table of contents and in the table.

Country and period apply to the item you add. Each result shows the code, its place in Eurostat's table of contents, the period covered and the number of values. The date is that of the last update. You get 20 results at a time.

What you get

Tick a result and choose **Add to sources**. The item is a **snapshot** with the retrieval date. It is called *dataset name (code)* and has these parts, each with the breadcrumb in its heading:

- **Beschrijving** (description): the query (country, period, language), the date of Eurostat's last update, the fixed choices (for example *Unit of measure: percentage of population in the labour force*) and the number of values.
- **Gegevens (deel 1, 2, …)** (data, part 1, 2, …): a table with a column for each varying characteristic (for example sex, country, year), the value and, where present, a Eurostat flag such as *p* (provisional) or *e* (estimated). Each part has 120 lines and the column names at the top.
- **Bron** (source): the attribution and the conditions.

An item contains at most **20,000 values**. If the selection is larger, the description states that the item has been truncated. Then choose a country or a shorter period.

Key

No key is needed. The administrator switches the connection on under **Administration > Source connections**.

Privacy

The search words stay in TrevantAI: TrevantAI retrieves Eurostat's table of contents and searches it itself. Only the dataset code, the country and the period are sent to Eurostat. Eurostat publishes only aggregated figures, no personal data. That is why, by default, an external AI may also process these items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

Eurostat allows reuse, including commercial reuse, without a separate licence, on two conditions: name Eurostat as the source, and state clearly if you have modified the figures. With the document, TrevantAI keeps the attribution "Bron: Eurostat" (Source: Eurostat). There are exceptions: figures on countries outside the EU, EFTA and the candidate countries (for example the US, Japan or China) and figures from other sources may only be reused **non-commercially**. That is why TrevantAI checks how the installation is used

(**Settings > General**). With commercial use, TrevantAI leaves out the values for such countries and the description names the countries left out; choosing such a country yourself, such as the US or Japan, then gives a message. With non-commercial use they remain and the description names them.

Good to know

- The table of contents is a large file. TrevantAI retrieves it at most once a day, so the first search may take a little longer.
- If Eurostat considers a selection too large, it prepares the file later. You then get a message; try again after a few minutes or choose less.
- If Eurostat has no figures for the chosen country or period, you get a message.
- For Dutch figures in more detail, CBS is often better.
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.
- A large table without a chosen country is sometimes too large for Eurostat. TrevantAI then only retrieves the latest period and says so in the description; choose a country for more periods.
- Under **View** you see the source, the code, the link, the retrieval date and the search query (**Found with**).

Companies and finance

🏢 KVK Business Register (Dutch companies)

The Business Register of the Netherlands Chamber of Commerce: look up a Dutch business or legal entity by name, KVK number or town and add the entry as a snapshot.

> Key: required (own key) · Type of use: Under your own agreement with the source (Subject to the KVK terms of use; requests are charged) · Personal data: may contain · Where: at the source

The **Business Register** (Handelsregister) of the Netherlands Chamber of Commerce (KVK) is the official register of businesses and legal entities in the Netherlands, from sole proprietorships and general partnerships (vof) to BVs, NVs, foundations and associations. TrevantAI uses the official KVK APIs **Zoeken** (search) and **Basisprofiel** (basic profile). You look up an entry and add it as a **snapshot**: name, legal form, trade names, the address of the main branch and the activities, as they stood in the register on the day of retrieval.

What can you use this source for?

- Legal: establishing the exact name of a counterparty, its legal form and since when it has been registered.
- Journalism: checking which trade names a business operates under and whether it is still active.
- Government and policy: recording the entry and the registered activities (SBI codes) for a permit or a grant.
- Business: keeping the register data of that moment on file for a new customer or supplier.

Searching

There are three search fields; fill in at least one.

- **Name or trade name**: for example *Koninklijke Philips*.
- **KVK number**: 8 digits, for example *17001910*. Spaces and dots are allowed ("KVK 6875 0110"); TrevantAI extracts the 8 digits from pasted text or a link. A different number of digits is refused.
- **Town**: for example *Amsterdam*. Combine a town with a name to tell companies with the same name apart.

Choose **Search**. Each result shows the name, the KVK number, the type of entry (main branch, secondary branch or legal entity), the town and, where relevant, "expired name" or "no longer active". Each KVK number appears once in the list. You get 20 results at a time; use **More results** to page further, up to 100 pages at most. If you are looking for one particular company, use the KVK number.

What you get

Tick a result and choose **Add to sources**. The item is called "Name (KVK number)". Every section has the breadcrumb in its heading, for example *Name (KVK 17001910) › Gegevens*. The section names are in Dutch:

- **Gegevens** (details): name, statutory name, KVK number, RSIN, legal form, trade names, date of registration, start date (and end date, if any), the number of people working at the main branch, and websites.

- **Hoofdvestiging en adres** (main branch and address): branch number, first trade name, whether it is a commercial branch, and the addresses.
- **Activiteiten** (activities): the SBI codes with their descriptions; the main activity is marked.
- **Bron** (source): that it is a snapshot and where the current entry can be found.

Deliberately not included: a **shielded address** (the item then only states that it is shielded), **officers** such as directors and partners (they are not in the Basisprofiel), and secondary branches or other profiles. The search results show only the town, not the street.

Key

Your own key is required. Your organisation takes out a subscription with KVK itself; TrevantAI does not supply any keys. The administrator enters the key under **Administration > Source connections**. How to request it is explained in the Help centre under *Key for the KVK Business Register*. The subscription costs €6.40 per month; searching is free, each company added costs €0.02.

Privacy

The search query is sent to KVK. Do not use client names or other confidential information in the query. The register may contain personal data: for a sole proprietorship or general partnership, the business name is often the name of a person, and the address is sometimes a home address. That is why these items are processed by the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The data are not open data. You use them in accordance with the KVK terms of use and your organisation's agreement with KVK. Each request for a Basisprofiel is charged; that is why TrevantAI only retrieves what you choose. With the document, TrevantAI keeps the attribution "Bron: KVK Handelsregister" (Source: KVK Business Register).

Good to know

- The item is a snapshot with the date of retrieval. The entry may change afterwards; you can see the current state with **View at the source**.
- The headings and labels in the item are in Dutch.
- If a company is already in your sources, you will see **Already in the sources**.
- The public KVK test key only works with fictitious data and is refused.
- This connection has not yet been checked against the real register. If something is missing, such as an address, check at KVK.
- Under **View** you can see where the item came from: the source, the KVK number, the link, the date of retrieval, the search query (**Found with**) and the licence.

🌐 GLEIF (LEI register of legal entities)

The global index of Legal Entity Identifiers: look up a legal entity or fund by name or LEI and add its name, addresses, parent companies and LEI registration as a snapshot.

> Key: none · Type of use: Free to use (Free to use (CC0)) · Personal data: no · Where: at the source

GLEIF (Global Legal Entity Identifier Foundation) maintains the global index of **Legal Entity Identifiers (LEI)**. A LEI is a 20-character code (ISO 17442) that legal entities and funds need for, among other things, financial transactions and reports to supervisory authorities. Each LEI comes with the legal name, the legal form, the addresses, the number in the national register and, if reported, the direct and the ultimate parent company. TrevantAI uses the public GLEIF API.

What can you use this source for?

- Business and finance: establishing, for customer due diligence (KYC) or compliance, which legal entity is behind a LEI and which group it belongs to.
- Journalism: tracing the group structure of a business through the direct and ultimate parent.
- Legal: recording the correct legal name, registration number and registered address of a foreign counterparty.
- Science: identifying legal entities in different countries unambiguously in a dataset.

Searching

- **Name or LEI** (required): a name, for example *ING Groep*, or a 20-character LEI. TrevantAI also recognises a LEI in pasted text or a link and then requests that one record directly; if the LEI is unknown, the list stays empty. A name is searched for in the data of the records (free text).
- **Country (registered address)** (choice, optional): Netherlands, Belgium, Germany, France, United Kingdom, United States, Luxembourg or Ireland. Leave it empty to search in all countries.

Choose **Search**. Each result shows the name, the LEI, town and country, the registration number and the status; the date is that of the initial registration. You get 20 results at a time; use **More results** to page further. Tip: if a common name gives too many hits, choose a country or use the LEI.

What you get

Tick a result and choose **Add to sources**. The item is called "Name (LEI code)" and is a **snapshot** with the date of retrieval. Every section has the breadcrumb in its heading, for example *Name (LEI …) › Adressen*. The section names are in Dutch:

- **Gegevens** (details): legal name, other names, LEI, legal form (ELF code), jurisdiction, category, status, registration number, register (GLEIF code) and date of incorporation.
- **Adressen** (addresses): the registered address and the headquarters address.
- **Moedermaatschappijen** (parent companies): the direct and the ultimate parent with their LEI, or "none reported to GLEIF".
- **Registratie van de LEI** (LEI registration): status (for example issued or lapsed), initial registration, last update, next renewal, the managing issuer (LOU) and the level to which the data have been corroborated.
- **Bron** (source): licence and where the current record can be found.

Key

No key is needed. The administrator only has to switch the connection on under **Administration > Source connections**.

Privacy

The search query is sent to GLEIF. Do not use client names or other confidential information in the query. A LEI exists only for legal entities and funds; the records that TrevantAI retrieves contain no names of persons. That is why, by default, an external AI may also process these items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The LEI data are freely available under **CC0 1.0**: you may use them without restrictions, including commercially. With the document, TrevantAI keeps the attribution "Bron: GLEIF (LEI-gegevens)" (Source: GLEIF, LEI data). Do not mention GLEIF in a way that suggests GLEIF endorses your work or TrevantAI.

Good to know

- Small companies often have no LEI. If you cannot find a Dutch company, try the KVK Business Register.
- A LEI with the status "lapsed" has not been renewed in time; the data may then be out of date.
- Parent companies only appear if they have been reported to GLEIF; "none reported" does not mean there is no parent.
- GLEIF allows 60 requests per minute; when adding, TrevantAI makes three requests per item. If the source asks you to wait, try again a little later.
- The labels in the item are in Dutch; names and addresses appear as GLEIF provides them.
- If a record is already in your sources, you will see **Already in the sources**. Under **View** you will find the source, the LEI, the link to search.gleif.org, the date of retrieval and the search query (**Found with**).

📈 Annual reports of listed companies (ESEF)

Annual reports of listed companies in Europe via filings.xbrl.org: search by name, LEI, country and financial year and add details, key figures and the text of the report as a snapshot.

> Key: none · Type of use: Description free, text protected by copyright (Made public by the company; copyright held by the company) · Personal data: may contain · Where: at the source

Listed companies in the European Union must publish their annual financial report digitally in the **European Single Electronic Format (ESEF)**. This is a web page (XHTML) in which the figures are tagged with **XBRL** labels. **filings.xbrl.org**, a service of XBRL International, collects these reports from the official storage mechanisms of the countries and makes them searchable. TrevantAI uses the public API of filings.xbrl.org.

What can you use this source for?

- Business and finance: compare the key figures and the risk section of a company and its competitors.
- Journalism: check what a company writes about remuneration, sustainability or lawsuits.
- Legal: record the text of the annual report as a source for a case file.
- Research: compare annual reports across several years or countries.

Searching

- **Company name or LEI**: part of the name, for example *Philips*, or a LEI of 20 characters.
- **Country** (choice): the country where the report was filed. Germany and Ireland are not on filings.xbrl.org.
- **Financial year (e.g. 2025)**: the year in which the financial year ends, as four digits.

Enter a name or LEI, or choose at least a country. Choose **Search**. Each result shows the company, the LEI, the country and the end of the financial year; the newest report comes first. You get 20 results at a time; use **More results** to see more.

What you get

Tick a report and choose **Add to sources**. The item is called "Company – jaarverslag year" and is a **snapshot** with the date of retrieval. Each section has the breadcrumb in its heading:

- **Gegevens** (details): company, LEI, country, end of the financial year, date added and the result of the technical check.
- **Kerncijfers** (key figures): from the XBRL data, for example revenue, operating result, net result, earnings per share, total assets, equity and cash. Only figures for the company as a whole, with the latest and the previous year.
- **The text of the report**: per heading if the report has headings. Many reports are converted PDFs without real headings; the text then follows in numbered parts.
- **Bron** (source): origin and conditions.

An annual report can run to hundreds of pages. TrevantAI therefore includes at most 250 text parts and 1.5 million characters. If there is more, the item says so. If the file is too large to read, you only get the details and key figures, with a notice.

Key

No key is needed. The administrator only has to switch the connection on under **Administration > Source connections**.

Privacy

The search query is sent to filings.xbrl.org. Annual reports name directors and supervisory board members, often with their remuneration. Items from this source are therefore processed **by the local AI only** by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

filings.xbrl.org names no licence and imposes no restrictions. The annual report itself was **made public by the company**; the **copyright** in the text belongs to the company. You may use figures and facts freely. You may quote from the text with attribution, but do not copy the text as a whole or distribute it. TrevantAI keeps the attribution "Bron: jaarverslag van de onderneming (ESEF), via filings.xbrl.org" with the document.

Good to know

- An annual report is large (up to tens of megabytes); adding it takes up to a minute.
- filings.xbrl.org is not complete. The official copy is held by the national storage mechanism, in the Netherlands by the AFM.
- The key figures are the standard IFRS concepts. If a company uses its own concepts, a figure may be missing; check the text in that case.
- Amounts are in the currency of the report, without rounding.
- The language of the item is that of the report; the labels are in Dutch.

- **View** shows the source, the link to the filings.xbrl.org viewer, the date of retrieval and the search (**Found with**).

🗂️ OpenCorporates (companies worldwide)

The largest open database of companies, with data from more than 140 business registers: look up a company by name or identifier and add its register data as a snapshot.

> Key: required (own key) · Type of use: Under your own agreement with the source (Subject to the licence of the customer's account (ODbL for free use)) · Personal data: may contain · Where: at the source

OpenCorporates is the largest open database of companies in the world. The organisation collects the data from more than 140 business registers and makes them searchable in one consistent way: name, registration number, status, incorporation, address, activities and officers. TrevantAI uses the official OpenCorporates API (version 0.4). For each company you can also see which original register the data refer to.

What can you use this source for?

- Journalism: checking in one go in which countries a company or a network of companies is registered.
- Legal: recording the register data of a foreign counterparty, with the link to the original register.
- Business and compliance: checking the status, date of incorporation and previous names of a foreign customer or supplier.
- Science: comparing companies from different countries in research on businesses or corporate governance.

Searching

- **Company name or identifier** (required): a name, for example *Koninklijke Philips*, or an OpenCorporates identifier in the form *jurisdiction/number*, for example *nl/17001910* or *us_de/2927442*. A link to opencorporates.com/companies/… also works. TrevantAI requests an identifier directly; if OpenCorporates does not know it, TrevantAI searches for it as a name (so a name with a slash also works).
- **Jurisdiction** (choice, optional): Netherlands, Belgium, Germany, United Kingdom, United States: Delaware, Luxembourg, Ireland or France. Empty means all jurisdictions.

Choose **Search**. Each result shows the identifier, the type, the status, "not active" if applicable, and the country; the date is the date of incorporation. You get 20 results at a time; use **More results** to page further, up to 100 pages at most. Tip: choose a jurisdiction if a name occurs in many countries.

What you get

Tick a result and choose **Add to sources**. The item is called "Name (identifier)" and is a **snapshot** with the date of retrieval. Every section has the breadcrumb in its heading, for example *Name (nl/17001910) › Gegevens*. The section names are in Dutch:

- **Gegevens** (details): name, jurisdiction, registration number, type, status, foreign branch if applicable, incorporation, dissolution, previous names and the register the data come from, with the date on which OpenCorporates retrieved them.
- **Adres** (address): the registered address.
- **Activiteiten** (activities): the industry codes with their descriptions, if available.
- **Functionarissen** (officers): at most 50, only name, position, start and end date.

- **Bron** (source): licence, the link to the original register and that it is a snapshot.

Deliberately not included: the address, date of birth and nationality of officers, even if the source provides them.

Key

Your own key (API token) is required. Your organisation opens an account with OpenCorporates itself; TrevantAI does not supply any keys. The API is paid; free use is possible through OpenCorporates' programmes for projects in the public interest, such as journalism, NGOs and research. The administrator enters the key under **Administration > Source connections**. How to request it is explained in the Help centre under *Key for OpenCorporates*.

Privacy

The search query is sent to OpenCorporates. Do not use client names or other confidential information in the query. The data contain names of officers, and for sole proprietorships the company name is often the name of a person. That is why these items are processed by the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The licence of your organisation's account applies. For free use this is the **Open Database License (ODbL)**: you may reuse the data with attribution, and whatever you derive from them and distribute falls under the same licence. For a paid account, the terms of that account apply. With the document, TrevantAI keeps the attribution "Bron: OpenCorporates" (Source: OpenCorporates).

Good to know

- Which data are available varies greatly by jurisdiction; not every register supplies officers or activities.
- Every account has a daily and a monthly limit. Once it is reached, OpenCorporates refuses further requests with the same message as for an incorrect key. So when you see that message, also check the account’s usage.
- This connection has not yet been checked against the real source; officers and activities in particular may turn out differently. If in doubt, check with **View at the source** or in the original register.
- The labels in the item are in Dutch. If a company is already in your sources, you will see **Already in the sources**. Under **View** you will find the source, the identifier, the link, the date of retrieval and the search query (**Found with**).

GB Companies House (UK companies)

The UK business register: look up a company by name or company number and add its details, registered office, activities, officers and recent filings as a snapshot.

> Key: required (own key) · Type of use: Free to use (Public government information) · Personal data: may contain · Where: at the source

Companies House is the official business register of the United Kingdom (England and Wales, Scotland and Northern Ireland). Every UK company is listed, with its company number, status, registered office, activities, officers and the documents it has filed, such as annual accounts. TrevantAI uses the official **Companies**

House Public Data API.

What can you use this source for?

- Legal: recording the correct name, company number and registered office of a UK counterparty, and checking whether there is an insolvency history.
- Journalism: seeing who is or was a director, and since when.
- Business and finance: checking for a UK customer or supplier whether the accounts were filed on time and which documents were filed recently.
- Science: following UK companies in research on businesses or corporate governance.

Searching

There is one search field: **Company name or company number** (required).

- A name, for example *Tesco*.
- A company number of 8 characters, for example *00445790* or *SC123456*. TrevantAI pads a number of 6 or 7 digits with leading zeros. A link to the company on find-and-update.company-information.service.gov.uk (…/company/00445790) also works.

TrevantAI requests a number or link directly. If there is no company with that number, TrevantAI simply searches for it. Each result shows the number, the type, the status (for example active, dissolved or in liquidation) and the town; the date is the date of incorporation. You get 20 results at a time; use **More results** to page further.

What you get

Tick a result and choose **Add to sources**. The item is called "Name (company number)" and is a **snapshot** with the date of retrieval. Every section has the breadcrumb in its heading, for example *Name (00445790) › Functionarissen*. The section names are in Dutch:

- **Gegevens** (details): name, company number, type, status, jurisdiction, incorporation, dissolution, previous names, the date of the last accounts, when the next accounts are due, the last confirmation statement, and whether an insolvency history or charges have been registered.
- **Statutair adres** (registered office): with a warning if the address is in dispute or if post cannot be delivered there.
- **Activiteiten** (activities): the SIC codes (UK SIC 2007).
- **Functionarissen** (officers): at most 50, with name, role and the dates of appointment and resignation.
- **Recente deponeringen** (recent filings): at most 25, as a table with date, category, form, description and a link to the document (PDF).
- **Bron** (source): that it is a snapshot and where the current register can be found.

Deliberately not included: date of birth, address, nationality, occupation and country of residence of officers, even if Companies House (partly) provides them. The search results show only the town, not an address. The filed documents themselves are not read; only the link is included.

Key

Your own key is required. It is free; your organisation creates it itself with an account at Companies House. TrevantAI does not supply any keys. The administrator enters the key under **Administration > Source connections**. How to create it is explained in the Help centre under *Key for Companies House*.

Privacy

The search query is sent to Companies House. Do not use client names or other confidential information in the query. The register names officers, so the items contain personal data. That is why these items are processed by the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The register is public. According to Companies House itself, there are no rules on reuse of the public register; an explicit open licence for the API data is not known. Therefore always cite the source, and take care with the names of persons. With the document, TrevantAI keeps the attribution "Bron: Companies House" (Source: Companies House).

Good to know

- Companies House allows 600 requests per five minutes. When adding, TrevantAI makes three requests per company (profile, officers, filings). If the source asks you to wait, try again a little later.
- The item contains the 25 most recent filings; older ones are available at Companies House.
- The labels in the item are in Dutch; values such as the type and the description of a filing are in English.
- This connection has not yet been checked against the real source. If something seems to be missing, check with **View at the source**.
- If a company is already in your sources, you will see **Already in the sources**. Under **View** you will find the source, the company number, the link, the date of retrieval and the search query (**Found with**).

US SEC EDGAR (US listed companies)

The public filings with the US Securities and Exchange Commission: look up a listed company by name, ticker or CIK and add its company data and recent filings with links as a snapshot.

> Key: none · Type of use: Free to use (Public government information) · Personal data: no · Where: at the source

EDGAR is the public archive of the **U.S. Securities and Exchange Commission (SEC)**, the US securities regulator. Listed companies file their mandatory documents there, such as the annual report (10-K), quarterly reports (10-Q) and reports of major events (8-K). TrevantAI uses the SEC's official EDGAR APIs and adds, for each company, the company data and the recent filings, with a link to each document.

What can you use this source for?

- Business and finance: quickly seeing which annual and quarterly reports a US company has filed recently, with the link to the document.
- Journalism: checking when a company filed an 8-K report, for example around an acquisition or a change of management.
- Legal: recording the official name, the state of incorporation and the previous names of a US counterparty.
- Science: keeping a fixed, traceable list of filings for research on companies or financial reporting.

Searching

There is one search field: **Company name, ticker or CIK** (required).

- A name, for example *Apple*. All the words you enter must occur in the name.
- A ticker, for example *AAPL*. An exact ticker appears at the top.
- A CIK (the SEC's permanent number, up to 10 digits), for example *320193*, or a link containing CIK=… or /edgar/data/…/.

Choose **Search**. Each result shows the name, the CIK and the tickers; a company with several share classes appears once in the list, with all its tickers. When searching by name, shorter names come first. You get 20 results at a time; use **More results** to page further.

What you get

Tick a result and choose **Add to sources**. The item is called "Name (CIK number)" and is a **snapshot** with the date of retrieval. Every section has the breadcrumb in its heading, for example *Name (CIK 0000320193) › Recente deponeringen*. The section names are in Dutch:

- **Bedrijfsgegevens** (company data): name, CIK, type, tickers, exchanges, SIC code with description, filer category, EIN, state of incorporation, fiscal year end (day-month), business address, mailing address, website and previous names.
- **Recente deponeringen** (recent filings): the last 40 as a table, with date, form, period, description, accession number and a link to the document.
- **Bron** (source): that it is a snapshot and that all filings are in EDGAR.

The documents themselves are not retrieved or read; only the link is included. If you want the content of, say, an annual report as a source, open the link and add the document separately.

Key

No key is needed. The administrator only has to switch the connection on under **Administration > Source connections**. The SEC does ask that every request includes a name and contact address; TrevantAI sends its own name and address, never those of a user.

Privacy

Do not use client names or other confidential information in the query. With this source, the query itself is not sent to the SEC: at most once a day TrevantAI retrieves the public list of listed companies and searches in it. Only when adding does TrevantAI request the data of that one company (by CIK) from the SEC. The items contain company data and no persons; a CIK that belongs to a person is refused. That is why, by default, an external AI may also process these items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

Information on sec.gov is public information of the US government and free to use. The SEC only asks for fair use of the service (at most 10 requests per second). With the document, TrevantAI keeps the attribution "Bron: SEC EDGAR" (Source: SEC EDGAR).

Good to know

- Only companies with a stock listing (ticker) are in the search list. Funds, private individuals and companies without a ticker cannot be found here.
- A new listing may appear in the search list up to a day later, because the list is refreshed once a day.
- The item contains the 40 most recent filings; older ones are in EDGAR (**View at the source**).
- The labels in the item are in Dutch; the values (such as the description of a filing) are in English.
- If a company is already in your sources, you will see **Already in the sources**. Under **View** you will find the source, the CIK, the link, the date of retrieval and the search query (**Found with**).

📢 TED (European public tenders)

Public procurement notices from the Supplement to the Official Journal of the European Union: search by words or publication number, country and date and add a notice as a source.

> Key: none · Type of use: Free to use (Free to reuse with acknowledgement of the source (Decision 2011/833/EU)) · Personal data: no · Where: at the source

TED (Tenders Electronic Daily) is the online version of the Supplement to the Official Journal of the European Union. The Publications Office of the EU publishes there the public procurement notices from the whole European Union: contract notices, award notices and corrigenda, with the contracting authority, the description, the lots, the CPV codes, the estimated value and the deadlines. TrevantAI uses the TED Search API (version 3). In the check against the real source, the word "onderhoud" (maintenance) alone gave more than 43,000 notices for the Netherlands.

What can you use this source for?

- Business: finding contracts your organisation can bid for, and recording the deadlines.
- Government and policy: seeing how other authorities described a similar contract and divided it into lots.
- Journalism: following which contracts a government puts out to tender and to whom they are awarded.
- Legal: keeping the official notice as evidence in a dispute about a tender.

Searching

- **Search words or publication number** (required): one or more words, for example *onderhoud bruggen* (bridge maintenance). Each word must occur in the notice; the full text is searched. At most 8 words count, and only letters, digits and hyphens; quotation marks and punctuation are removed, so searching for an exact phrase is not possible. You can also enter a publication number, such as *123456-2024*, or paste a link to the notice on ted.europa.eu; that exact notice is then requested.
- **Country of the contracting authority** (choice): All countries, Netherlands, Belgium, Germany, France or Luxembourg.
- **Published from** (date, optional): only notices from this date. There is no end date.

Choose **Search**. The newest notices appear at the top. Each result shows the title, the contracting authority, town and country, the type of notice, the publication number and the publication date. You get 20 results at a time; use **More results** to page further. TED returns at most 15,000 results per query; in that case make the query more specific with a country, a date or more words. If TED refuses a query, try other words.

What you get

Tick a notice and choose **Add to sources**. The item has the title of the notice. Every section has the breadcrumb in its heading, for example *Title › Data en termijnen*. The section names are in Dutch:

- **Opdracht** (contract): title, type of notice, nature of the contract, CPV codes, estimated value and the description of the procedure.
- **Aanbestedende dienst** (contracting authority): name, town and country.
- **Percelen en omschrijving** (lots and description): the description, per lot if there are several lots.
- **Data en termijnen** (dates and deadlines): publication date and deadline for submitting tenders.
- **Bron** (source): the publication number, the link to ted.europa.eu and the attribution "© Europese Unie" (© European Union).

A section without content is left out. Deliberately not included: contact persons, email addresses and telephone numbers.

Key

No key is needed. The administrator only has to switch the connection on under **Administration > Source connections**.

Privacy

The search query is sent to TED (the Publications Office of the EU). Do not use client names or other confidential information in the query. Notices concern contracts of organisations, and TrevantAI does not request any contact details. That is why, by default, an external AI may also process these items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

TED data are free to reuse, including commercially, with acknowledgement of the source (European Commission Decision 2011/833/EU). With the document, TrevantAI keeps the attribution "© Europese Unie, https://ted.europa.eu" (© European Union).

Good to know

- A notice is often available in several languages. TrevantAI takes the Dutch text if there is one, otherwise the English, otherwise the first available language. The labels in the item are in Dutch.
- A corrigendum appears in TED as a separate notice with its own publication number. So also look for later notices about the same contract.
- The item shows the notice as TED provided it at the time of retrieval; the date of retrieval is recorded with it. For the complete notice with all form sections, use **View at the source**.
- If a notice is already in your sources, you will see **Already in the sources**. Under **View** you will find the source, the publication number, the link, the publication date, the date of retrieval and the search query (**Found with**).

🚫 Source connection Sanctions lists (EU, UK, US)

The official sanctions lists of the EU, the UK and the US on this computer: search for a name without it leaving the computer, and add an entry as a snapshot.

> Key: none · Type of use: Free to use (Official sanctions lists of the EU (Decision 2011/833/EU), the UK (Open Government Licence v3.0) and the US (public domain): free to reuse with attribution) · Personal data: may contain · Where: local (download)

A **sanctions list** names persons, organisations, vessels and aircraft against which a government has taken measures, such as freezing assets or a travel ban. TrevantAI retrieves three official lists from the publishers themselves and keeps them as a **source library on this computer**:

- **EU**: the European Commission's consolidated list of financial sanctions (FSF).
- **UK**: the UK Sanctions List of the Foreign, Commonwealth & Development Office (since 28 January 2026 the only UK list; the OFSI list is no longer updated).
- **US**: the SDN list of OFAC (U.S. Department of the Treasury).

What can you use this source for?

- Business and finance: checking during customer due diligence (KYC) or compliance whether a name appears on a sanctions list.
- Legal: recording which measure applies, since when and under which programme or regime.
- Journalism: checking whether a person, company or vessel is sanctioned, and by whom.
- Government and policy: comparing the sanctions regimes of different countries.

Searching

- **Name** (required): the name of a person, organisation or vessel. Capitals and accents do not matter (*Muller* finds *Müller*). The beginning of a word also works (*Müll*), as does part of a name (*idorov*), a name without spaces or apostrophe (*alqaida*) and a reference number from the list (for example *EU.27.28* or *QDi.001*). TrevantAI searches all names and aliases, including the original script (for example Cyrillic or Arabic) where the source provides it.
- **List** (choice): EU, UK or US. Leave it empty to search all three.
- **Type** (choice): persons, organisations or vessels.

There is **no automatic 'match'**. You get a list of **possible hits**, each showing how the name corresponds: the same words, the beginning of words, part of a name, or a name that sounds roughly the same (for example *Mohamed* and *Muhammad*). Each result also shows the type, the programme, the date of birth and the nationality. A person assesses whether it is the same person or organisation.

What you get

Tick a result and choose **Add to sources**. The item is a **snapshot** of one entry, with the date of the list. Each section has the breadcrumb in its heading:

- **Entry**: name, type, list, reference, programme or regime, listed since and last changed.
- **Names**: all names and aliases, with how reliable the source considers them.
- **Personal details** or **Details**: date and place of birth, nationality, position; for organisations and vessels registration numbers such as the IMO number.
- **Addresses**: only town, region and country.

- **Remarks** and **Source**: the publisher's explanation, the date and SHA-256 of the file, and the link to the current list.

Key

No key is needed. The administrator switches the connection on under **Administration > Source connections**; the lists are then installed (about 80 MB to download). If the library is more than a day old, **Update** appears.

Privacy

Searching happens **only on this computer**: the name you search for, often that of a client, is not sent anywhere. Only retrieving the lists themselves uses the internet. The lists contain names of persons and sensitive data; items are therefore processed by the local AI only by default. TrevantAI does **not take over passport or other document numbers** of persons, nor street, house number, telephone number or e-mail address.

Licence and attribution

The three lists may be reused freely with attribution: EU under Decision 2011/833/EU, UK under the Open Government Licence v3.0, US a work of the US government (public domain). The UN Security Council list is not part of this source: it has no open licence and is a separate source, **UN sanctions list (Security Council)**, which the administrator can switch on separately. The EU and the UK take over UN listings, with the UN reference, so you will also find them here. Each item states the publisher and the date of the list.

Good to know

- A snapshot is not the current state. Check the list at the source before taking a decision.
- The same person may appear on several lists; with a UN reference, the item shows where else.
- If retrieving one list fails, the previous version remains; the item says so.
- The OFAC list does not state a listing date.
- This library is not a complete screening system. For PEPs and other lists there is the OpenSanctions connection.

UN Source connection UN sanctions list (Security Council)

The UN Security Council sanctions list on this computer: search for a name without it leaving the computer. No open licence; items stay within your own organisation.

> Key: none · Type of use: Description free, text protected by copyright (Under the UN terms: no open licence; for use within your own organisation only, do not redistribute) · Personal data: may contain · Where: local (download)

The **UN Security Council** imposes sanctions on persons and organisations, for example in connection with Al-Qaida, ISIL, the Taliban or North Korea. All names are brought together in the **Consolidated United Nations Security Council Sanctions List**. All member states must implement these measures. TrevantAI retrieves the list from the UN itself and keeps it as a **source library on this computer**, together with the lists of the **Sanctions lists (EU, UK, US)** source.

What can you use this source for?

- Business and finance: checking during customer due diligence (KYC) or compliance whether a name is on the UN list.
- Legal: recording under which UN sanctions regime someone is listed and since when.
- Journalism: checking whether a person or organisation has been sanctioned by the Security Council.
- Government and policy: comparing the UN list with the lists of the EU, the UK and the US.

Why a separate source? The EU and the UK take over UN listings, with the UN reference. For most checks the **Sanctions lists** source is therefore enough. This source is for those who want to consult the UN list itself.

Searching

- **Name** (required): the name of a person or organisation. Capitals and accents do not matter. The beginning of a word also works, as does part of a name, a name without spaces or apostrophe (*alqaida*) and a UN reference (for example *QDi.001*). TrevantAI searches all names and aliases, including the original script (for example Arabic) where the UN provides it.
- **Type** (choice): persons or organisations.

There is **no automatic 'match'**. You get **possible hits**, each showing how the name corresponds: the same words, the beginning of words, part of a name, or a name that sounds roughly the same. A person assesses whether it is the same person or organisation, using date of birth, nationality and the other details.

What you get

Tick a result and choose **Add to sources**. The item is a **snapshot** of one entry, with the date of the list:

- **Entry**: name, type, UN reference, sanctions regime, listed since and last changed.
- **Names**: all names and aliases, with how reliable the UN considers them.
- **Personal details** or **Details**: date and place of birth, nationality, position or title.
- **Addresses**: only town, region and country.
- **Remarks** and **Source**: the UN's explanation, the date and SHA-256 of the file and the link to the current list.

Key

No key is needed. The administrator switches the connection on under **Administration > Source connections** and confirms the terms. The UN list is then added to the source library. If this source is switched off, TrevantAI does not retrieve the UN list, and at the next update it disappears from the library. Items already added remain.

Privacy

Searching happens **only on this computer**: the name searched for is not sent anywhere. The list contains names of persons and sensitive data; items are therefore processed by the local AI only by default. TrevantAI does **not take over passport or other document numbers** of persons, nor street or house number.

Licence and attribution

The UN publishes the list to make compliance with the sanctions possible, but the list has **no open licence**. The general terms of use of un.org only allow personal, non-commercial use, without the right to resell or redistribute the data or to create derivative works from it. By confirming the terms, your organisation takes responsibility for the use itself. Items from this source are for **use within your own organisation only**: they are left out by default when sharing or exporting outside it.

Good to know

- A snapshot is not the current state. Check the list at the UN before taking a decision.
- With a UN reference, the item shows whether the same entry also appears with the EU or the UK.
- If retrieval fails, the previous version remains; the item says so.
- If the list is more than a day old, **Update** appears.

Source connection OpenSanctions (sanctions and PEPs)

An international collection of sanctions lists and politically exposed persons: search for a person or organisation and add names, positions, sanctions and source lists as a snapshot. With your own key and licence.

> Key: required (own key) · Type of use: Under your own agreement with the source (Under the customer's licence with OpenSanctions (data CC BY-NC 4.0)) · Personal data: may contain · Where: at the source

OpenSanctions is an international collection of more than two hundred lists: sanctions lists of countries and international organisations, lists of **politically exposed persons (PEPs)** and their associates, and other lists of interest, such as wanted persons and debarments from public procurement. OpenSanctions merges the entries per person or organisation, so that you see in one place which lists someone is on. TrevantAI uses the OpenSanctions API.

What can you use this source for?

- Business and finance: checking during customer due diligence (KYC) or compliance whether someone is a PEP or on a sanctions list, including outside the EU, the UN, the UK and the US.
- Legal: recording which sanctions apply, from which authority and since when.
- Journalism: mapping the positions of a politician or the sanctions against a company or vessel.
- Government and policy: seeing which countries have sanctioned the same person or organisation.

Searching

- **Name** (required): the name of a person, organisation, company or vessel.
- **Collection** (choice): empty means everything; or sanctions only, or PEPs only.
- **Type** (choice): persons, organisations, companies or vessels.
- **Country** (optional): a two-letter country code, for example *RU* or *NL*.

Choose **Search**. Each result shows the type, the topics (for example *on a sanctions list* or *politically exposed person*), the date of birth, the countries, a position and the number of source lists. You get 20 results at a time. There is **no automatic 'match'**: the results are **possible hits**, which you assess yourself using date of birth, nationality and the other details.

What you get

Tick a result and choose **Add to sources**. The item is called "Name (OpenSanctions id)" and is a **snapshot** with the date of retrieval:

- **Details**: names and aliases, type, topics, date and place of birth, nationality; for organisations and vessels also registration numbers, such as the IMO number or the LEI.
- **Positions**: for PEPs, the positions with country, start and end.
- **Sanctions**: for each sanction the authority, the programme, the period, the measures, the reason and the source list.
- **Remarks** and **Source**: the explanation, the lists from which the entry was compiled, and the link to OpenSanctions.

Key

Your organisation needs **its own key**. How to request and enter it is described in **Key for OpenSanctions**. The administrator confirms the terms when switching the connection on.

Privacy

The search term is sent **to OpenSanctions**, with your organisation's key. If you prefer to search without a name leaving the computer, use the **Sanctions lists** connection, which searches locally. The items contain names of persons and sensitive data; they are therefore processed by the local AI only by default. TrevantAI does not take over **addresses, identity numbers** (such as passport numbers) or contact details of persons.

Licence and attribution

OpenSanctions data are licensed under **CC BY-NC 4.0**: free for non-commercial use, with attribution. For business use your organisation needs **its own licence** from OpenSanctions (or the paid API). Your organisation arranges this itself; TrevantAI does not supply any data or licence. With each item TrevantAI keeps the attribution "Source: OpenSanctions".

Good to know

- A snapshot is not the current state. Check the entry at OpenSanctions and in the official source list before taking a decision.
- OpenSanctions merges entries from different lists. A mistake in that merging is possible; if in doubt, check the source list.
- Every search and every item retrieved counts towards the use of your key.
- If OpenSanctions asks you to wait, try again a little later.

Science

🎓 OpenAlex (scholarly publications)

An open catalogue of over 250 million scholarly publications; you add the description and the abstract, with a link to the full text.

> Key: optional · Type of use: Description free, text protected by copyright (Descriptions free to use; abstracts may be protected by copyright) · Personal data: no · Where: at the source

OpenAlex is an open catalogue of scholarly publications from all over the world, maintained by the non-profit organisation OurResearch. It contains over 250 million articles, books and theses, with authors, journal, year, DOI, subjects, the number of citations and often the abstract.

What can you use this source for?

- Science and policy: mapping the state of research on a subject, for example PFAS in estuaries.
- Journalism: checking whether a claim in a press release is backed by published research, and how often it has been cited.
- Legal: finding the literature that an expert cites in a report.
- Business: collecting research on a technology or market.

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **OpenAlex (scholarly publications)**.

- **Search words or identifier** (required): for example *PFAS estuary*. Any language is allowed, but most publications are in English; English technical terms usually give more results. An OpenAlex identifier (W followed by digits, such as W2741809807) or a link containing one (openalex.org/W…) gives exactly that one publication. This source does not recognise a DOI; search Crossref for that.
- **From year** and **Up to year**: the year of publication, in four digits. One of the two is also fine.
- **Access**: **Open access only** shows only publications that are freely readable.
- **Language of the publication**: one of TrevantAI's seven languages.

Choose **Search**. OpenAlex searches titles and abstracts, among other things. You get 20 results at a time, with the first authors, journal, year, open access and DOI; use **More results** to browse further, up to at most 10,000. Too many results? Add a word or narrow the period. Tick what you want (at most 20 at a time) and choose **Add to sources**.

What you get

Each publication becomes one source with these sections (the headings in the item are in Dutch):

- **Description** (*Beschrijving*): title, authors, published in, volume and pages, date, type, DOI, identifier, number of citations and the open access link;
- **Abstract** (*Samenvatting*), if OpenAlex has one;
- **Subjects** (*Onderwerpen*): topics and keywords from OpenAlex;
- **Source** (*Bron*): the attribution.

Each section starts with a breadcrumb trail (*title › Samenvatting*). TrevantAI does not retrieve the full text; it is at the publisher or behind the open access link. Under **View** you will find the identifier, the link (the DOI), the date of retrieval, the search query and the licence. The item is a **snapshot**: the number of citations may be higher later.

Key

Not needed. Without a key a daily limit applies; with your organisation's own free key it is ten times higher. TrevantAI is supplied without keys. The administrator enters the key under **Administration > Source connections**; how to request it is explained in the Help Centre under *Key for OpenAlex (optional)*.

Privacy

The search query is sent to OpenAlex (OurResearch), with TrevantAI's contact address, never that of a user. Please do not enter client names or other confidential information. Author names are public bibliographic data; by default, therefore, an external AI may also process the items if the AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The OpenAlex data are free to use (CC0). The abstracts come from the publishers and may be protected by copyright: quoting is allowed, but reproducing one in full in a publication may require permission. TrevantAI keeps **Bron: OpenAlex** (Source: OpenAlex) with the document. In your work, also cite the publication itself.

Good to know

- Not every publication has an abstract; that is how the source is.
- The number of citations is OpenAlex's; other databases sometimes give a different number.
- If the limit has been reached, TrevantAI reports that the source asks to wait.
- Already added? Then you will see **Already in the sources**, with **View**.
- Through the identifier W… and the DOI it remains traceable which publication was used.

Unpaywall enrichment: if the **Unpaywall** enrichment is switched on in **Administration > Source connections**, TrevantAI when adding looks for a legal open version via the DOI. If its licence also allows commercial reuse, the full text is added, page by page; otherwise **Source** says why not. See Unpaywall.

📚 Crossref (publications with a DOI)

The descriptions that publishers register with Crossref for publications with a DOI; useful for finding an article, book or report exactly by its DOI.

> Key: none · Type of use: Description free, text protected by copyright (Descriptions free to use; abstracts may be protected by copyright) · Personal data: no · Where: at the source

Crossref is the organisation with which publishers register the DOIs of their publications. Each DOI comes with a description supplied by the publisher itself: title, authors, journal or book, year, publisher and sometimes the abstract. It covers journal articles, books and chapters, conference papers, reports and theses from publishers all over the world.

What can you use this source for?

- Science: going through a reference list and recording each reference as a source with the correct details.
- Legal: finding exactly the publication cited with a DOI in a report or court document.
- Government and policy: checking with which publisher and in which year a study appeared.
- Journalism: checking whether a cited article really exists and who the authors are.

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **Crossref (publications with a DOI)**.

- **Search words or DOI** (required): words from the title, names of authors or the journal, for example *PFAS estuary sediment*. Or a DOI, such as 10.1038/nature12373, also as a link (https://doi.org/10....) or in the middle of a text. With a DOI you get exactly that one publication.
- **From year** and **Up to year**: the year of publication, in four digits. One of the two is also fine.
- **Type of publication**: Journal article, Book chapter, Book, Conference paper, Report or Thesis.

Choose **Search**. Crossref searches the descriptions (title, authors, journal), not the full text. You get 20 results at a time, with the first authors, the journal, the date and the DOI; use **More results** to browse further, up to at most 10,000. TrevantAI leaves out supplementary material, peer reviews, grants and datasets; a page may therefore show fewer than 20 results. Tick what you want (at most 20 at a time) and choose **Add to sources**.

What you get

Each publication becomes one source with these sections (the headings in the item are in Dutch):

- **Description** (*Beschrijving*): title and subtitle, authors, published in, volume, issue and pages, date, publisher, type, DOI, number of citations according to Crossref, subjects and the licence of the publication, if the publisher states one;
- **Abstract** (*Samenvatting*), if the publisher supplies one;
- **Source** (*Bron*): the attribution.

Each section starts with a breadcrumb trail (*title › Beschrijving*). TrevantAI does not retrieve the full text; it is at the publisher, via the DOI. Under **View** you will find the identifier (the DOI), the link, the date of retrieval, the search query and the licence. The item is a **snapshot** of the description.

Key

No key needed.

Privacy

The search query is sent to Crossref, with TrevantAI's contact address, never that of a user. Please do not enter client names or other confidential information. Crossref contains descriptions of publications; author names are public bibliographic data. By default, therefore, an external AI may also process the items if the AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The Crossref descriptions are free to use. Abstracts may be protected by copyright: quoting is allowed, but reproducing one in full in a publication may require the publisher's permission. If the publisher states a licence for the publication, it appears in the description. TrevantAI keeps **Bron: Crossref** (Source: Crossref) with the document. In your work, cite the publication itself, with its DOI.

Good to know

- How complete a description is depends on the publisher; by no means every publication has an abstract.
- The number of citations only counts references known to Crossref.
- Already added? Then you will see **Already in the sources**, with **View**.
- If you mainly search by subject, OpenAlex often gives more abstracts; Crossref is strong at finding a publication exactly by DOI.
- The DOI is a permanent identifier: it always remains traceable which publication was used.

Unpaywall enrichment: if the **Unpaywall** enrichment is switched on in **Administration > Source connections**, TrevantAI when adding looks for a legal open version via the DOI. If its licence also allows commercial reuse, the full text is added, page by page; otherwise **Source** says why not. See Unpaywall.

arXiv (preprints)

Cornell University's open preprint server for physics, mathematics, computer science and related fields; you add the description and the abstract of an article.

> Key: none · Type of use: Free to use (Free to use (CC0)) · Personal data: no · Where: at the source

arXiv is Cornell University's open preprint server. Researchers post their articles there, often before they appear in a journal. It covers physics, mathematics, computer science, quantitative biology and finance, statistics, economics and electrical engineering. TrevantAI adds the description and the abstract, with a link to the full text.

What can you use this source for?

- Science: following the latest articles in your field, for example on graph neural networks.
- Business and technology: checking what has recently been published about a technique, such as language models or battery storage.
- Journalism: finding an article reported in the news, and seeing whether it has already appeared in a journal.
- Finance: collecting research on risk models or market behaviour (quantitative finance).

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **arXiv (preprints)**.

- **Search words or identifier** (required): English search words, for example *graph neural networks*. All words must occur; "words in quotation marks" are searched as a phrase. If you know arXiv's query language, you can use it too (e.g. *ti:*, *au:*, *abs:* or AND, OR, ANDNOT). You can also enter an identifier, such as 2401.01234 or hep-th/9901001, with or without *arXiv:* in front, or paste a link (arxiv.org/abs/..., /pdf/... or /html/...). You then get exactly that one article.

- **Field**: Computer science, Mathematics, Physics, Quantitative biology, Quantitative finance, Statistics, Economics or Electrical engineering and systems science. Physics also includes astrophysics, condensed matter, particle and nuclear physics and quantum physics, among others.
- **Order**: **Most relevant** (default) or **Newest first**.

Choose **Search**. You get 20 results at a time, with the first authors, the primary category, the year and the identifier; use **More results** to browse further, up to at most 30,000. arXiv asks for no more than one request every three seconds; TrevantAI therefore waits briefly each time, and searching or adding sometimes takes a few seconds. Tick what you want (at most 20 at a time) and choose **Add to sources**.

What you get

Each article becomes one source with these sections (the headings in the item are in Dutch):

- **Description** (*Beschrijving*): title, authors, identifier, version, primary category and categories, date of the first and of this version, DOI, journal reference and the authors' comment (e.g. number of pages);
- **Abstract** (*Samenvatting*);
- **Links** (*Links*): the abstract page and the DOI;
- **Source** (*Bron*): the attribution.

Each section starts with a breadcrumb trail (*title › Samenvatting*). TrevantAI always retrieves the latest version; the version number is in the description. The PDF is not retrieved, because the licence of the full text differs per article. Authors' affiliations are not included. The item is a **snapshot**: the authors may post a new version later.

Key

No key needed.

Privacy

The search query is sent to arXiv (Cornell University). Please do not enter client names or other confidential information. arXiv contains scholarly articles; author names are public bibliographic data. By default, therefore, an external AI may also process the items if the AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The arXiv descriptions and abstracts are free to use (CC0). The full text is covered by the licence the authors chose; it is shown on the abstract page. TrevantAI keeps **Bron: arXiv.org** (Source: arXiv.org) with the document. In your work, cite the article itself, with its identifier and version.

Good to know

- arXiv asks for three seconds between two requests; searching and retrieving therefore each take a few seconds.
- arXiv contains preprints: not every article has been peer-reviewed. Check under *Verschenen in* (published in) whether there is a journal version.
- Searching works best in English.
- Already added? Then you will see **Already in the sources**, with **View**.

- If arXiv rejects the query, check the spelling of prefixes and AND/OR.
- With the identifier and version it remains traceable which text was used.

EU OpenAIRE (European open research)

The European open research knowledge graph: publications, datasets and software with their projects and funders; you add the description and the abstract, and can follow new work.

> Key: none · Type of use: Description free, text protected by copyright (Free to use with attribution (CC BY 4.0)) · Personal data: no · Where: at the source

OpenAIRE is the European open research knowledge graph, set up with support from the European Commission. It contains publications, datasets, software and other research products from thousands of repositories, journals and registries, together with their projects and funders. TrevantAI adds the description and the abstract.

What do you use this source for?

- Science: collect publications and datasets on a topic, including from repositories that other databases miss.
- Government and policy: check what a European project (for example from Horizon 2020) has produced.
- Grants and funds: follow the work a funder made possible, such as the European Commission or NWO.
- Business: find open research and software within a field.

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **OpenAIRE (European open research)**.

- **Search words or identifier** (required): search words, usually in English, for example *open science policy*. You can also enter a DOI (10.…), an OpenAIRE identifier or a link from explore.openaire.eu; you then get exactly that item.
- **Type**: publication, dataset, software or other research products.
- **Access**: **Open access only** shows only items that are freely readable.
- **Country (code)**: the country of the organisations involved, as a two-letter code (NL, BE, DE).
- **Funder**: the short name of a funder, for example EC (European Commission) or NWO.
- **Published from** and **Published up to**: the publication date.

Choose **Search**. You get 20 results at a time, with the first authors, the journal or publisher, the year, whether there is open access and the DOI; use **More results** to browse further, up to 10,000. Tick what you want (at most 20 at a time) and choose **Add to sources**.

What you get

Each item becomes one source with:

- **Description**: title, authors, published in, publisher, date, type, DOI, identifier, language, countries, open access, licence, the projects with their funders and the locations (repository or publisher, with licence and link);
- **Abstract**, if there is one;
- **Source**: the attribution, the item's licence and what has and has not been included.

OpenAIRE does not provide the full text. If the **Unpaywall** enrichment is switched on, TrevantAI looks for a legal open version and includes it if the licence allows (see Unpaywall). If an item has very many authors, projects or locations, the **Description** says the list was cut off. The item is a **snapshot**.

Following

With **Follow…** next to **Search** you create a subscription: TrevantAI then asks at set times what has newly appeared, with the same filters (type, access, country, funder and words; all of them go to OpenAIRE). New here means: with a **publication date** from the previous round onwards. OpenAIRE has no way of asking what has been *added* since a given moment: what OpenAIRE only takes in later with an older publication date therefore does not come in. What is yet to appear comes in on its publication day. For one publication date TrevantAI can go through at most 10,000 items; if there are more (for publications of which only the year is known, OpenAIRE sets the date to 1 January), TrevantAI skips the rest of that day and continues with the next day. In that case, make the subscription narrower, for example with words or a country. More in Following sources.

Key

No key needed. Without a key, OpenAIRE allows about 60 requests per hour; TrevantAI therefore keeps search results for a short while, so adding after searching costs no extra request. If the limit is reached, TrevantAI reports that the source asks you to wait.

Privacy

The search query is sent to OpenAIRE. Please do not enter names of clients or other confidential information. Author names are public bibliographic data; by default, therefore, an external AI may also process the items if the AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

OpenAIRE's data may be reused freely with attribution (CC BY): name OpenAIRE as the source. The abstracts come from the authors or the publisher and may be protected by copyright. The licence of each item is shown in the **Description** and under **Source**. TrevantAI keeps **Bron: OpenAIRE Graph** with the document. In your work, also cite the item itself.

Good to know

- OpenAIRE merges duplicate records; one item may therefore have several locations.
- Countries and funders are sometimes inferred by OpenAIRE and not always complete.
- Already added? Then it says **Already in the sources**, with **View**.
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.

📖 Unpaywall (open versions of articles)

An enrichment: for articles from OpenAlex, Crossref, Europe PMC and OpenAIRE, TrevantAI looks for a legal open version via Unpaywall and includes the full text if the licence allows it.

> Key: required (own key) · Type of use: Free to use (Free to use (CC0)) · Personal data: no · Where: at the source

Unpaywall is a service from OurResearch, the makers of OpenAlex, that knows for each article whether a legal open version exists: at the publisher or in a repository of, for example, a university, and which version it is. In TrevantAI, Unpaywall is an **enrichment**: not a source to search in itself, but an addition to other sources.

What do you use this enrichment for?

- Science and education: have the full text of an article from OpenAlex or Crossref searchable straight away, if it is legally available in open access.
- Government and policy: query articles from European projects (via OpenAIRE) in full.
- Medical: still find an open manuscript for an article from Europe PMC without open text.

How it works

Unpaywall does not appear under **Public sources** or in the source profiles. If the enrichment is switched on in **Administration > Source connections**, then when an article is added from **OpenAlex**, **Crossref**, **Europe PMC** (only if no full text was included there) or **OpenAIRE**, TrevantAI asks Unpaywall, using the DOI, whether an open version exists. This also applies to items that a subscription adds straight away.

What you get

- If the licence of the open version also allows commercial reuse (CC0, CC BY, CC BY-SA, CC BY-ND or public domain), TrevantAI fetches the PDF and adds the text page by page (*title › Full text › Page 3*), up to 60 pages; if there are more, the **Description** says so.
- The **Source** section shows the location, the version (published, accepted manuscript or submitted) and the licence.
- Otherwise the item stays as it was, with one sentence in **Source** saying why, for example: no open version, or a licence that does not allow it.

A licence that only allows non-commercial use (NC) only counts if the installation is used non-commercially (**Settings > General**); the item then gets that label. With a no-derivatives licence (ND), a translation or summary by AI may not be shared outside the organisation.

Email address

Unpaywall asks for the email address of whoever asks the question with every request. The administrator enters the organisation's email address (not TrevantAI's) under **Administration > Source connections**, at **Unpaywall (open versions of articles)**, and switches the enrichment on. With **Test**, TrevantAI checks the address with one trial request. Without an email address, the enrichment does nothing.

Privacy

Only the article's DOI and the organisation's email address are sent to Unpaywall: no search query, no title and no user name. The PDF comes from the site that Unpaywall names, with the same protection as for web pages. An item keeps the AI setting of the source it comes from.

Licence and attribution

Unpaywall's data are free to use (CC0). The fetched text is covered by the licence of the open version; it is shown with each item in **Source**. In your work, cite the article itself, preferably the published version.

Good to know

- Unpaywall only knows Crossref DOIs; for other DOIs (for example from DataCite, for datasets) it finds nothing.
- If TrevantAI includes the full text, it stores with the document the licence of that open version (not that of the description) and an attribution with the location. With CC BY-SA the document is: free, with attribution and share-alike.
- If there are several open versions, a version under a free licence takes precedence over one under an NC licence, even if the latter is the published version.
- If Unpaywall refuses the email address, the Source of the document says so and the administrator sees it among the source failures. The document is not enriched again automatically afterwards.
- An accepted manuscript may differ in details from the published version.
- A PDF without a text layer (a scan) is not read; **Source** then says so.
- This connection has not yet been checked against the real source.

Health and medicine

🧪 ClinicalTrials.gov (clinical studies)

The register of clinical studies of the U.S. National Library of Medicine; you add the design, eligibility criteria and outcome measures of a study.

> Key: none · Type of use: Free to use (Public government information) · Personal data: no · Where: at the source

ClinicalTrials.gov is the register of clinical studies of the U.S. National Library of Medicine (NLM). It contains studies from all over the world: what is being studied, for which condition, with which drug or other intervention, who may take part, what is measured and how far the study has progressed. TrevantAI adds this description of a study as a source.

What can you use this source for?

- Medical: checking which studies are under way for a condition, for example asthma, and which are still recruiting.
- Business and finance: following a pharmaceutical company's research pipeline (which products in which section).
- Journalism: checking whether a study has been registered, who the sponsor is and which outcome measures were stated.
- Legal and policy: recording the design and eligibility criteria of a study for an opinion or case file.

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **ClinicalTrials.gov (clinical studies)**. Fill in at least one of the three text fields, in English:

- **Search words or identifier**: free search words, for example *semaglutide obesity*. Or only an identifier (NCT followed by eight digits, such as NCT01234567) or a link to the study; you then get exactly that one study. If the number is among other words, a normal search is done.
- **Condition**: for example *asthma*.
- **Intervention (e.g. drug)**: for example *metformin*.
- **Status of the study**: Recruiting, Not yet recruiting, Active, not recruiting, Completed, Terminated, Withdrawn or Suspended.

If you fill in several fields, the studies must meet all of them. Choose **Search**. You get 20 results at a time, with identifier, status, type of study, section, conditions and sponsor; use **More results** to browse further. Tick what you want (at most 20 at a time) and choose **Add to sources**.

What you get

Each study becomes one source with these sections (the headings in the item are in Dutch):

- **Summary** (*Samenvatting*): title and official title, identifier, status, conditions, sponsor, start, end (expected or actual), last update at the source and the brief summary;
- **Description** (*Beschrijving*): the detailed description;
- **Design and interventions** (*Opzet en interventies*): type of study, section and interventions;

- **Eligibility criteria** (*Deelnamecriteria*): sex, minimum and maximum age and the inclusion and exclusion criteria;
- **Outcome measures** (*Uitkomstmaten*): primary and secondary outcome measures, with time frame;
- **Results** (*Resultaten*): only whether there are results, with the titles of the outcome measures; not the figures themselves;
- **Source** (*Bron*): attribution, link and the date of processing.

Each section starts with a breadcrumb trail (*title › Deelnamecriteria*). Contact persons, investigators and locations are deliberately **not** included. The item is a **snapshot**; the date at the source is the start date.

Key

No key needed.

Privacy

The search query is sent to ClinicalTrials.gov (NLM). Please do not enter names of clients or patients, diagnoses of individuals or other confidential information. The register contains contact details of investigators, but TrevantAI does not include them; what remains contains no personal data. By default, therefore, an external AI may also process the items if the AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The data come from the US federal government and are free of copyright in the United States. The ClinicalTrials.gov terms ask for an attribution and the date on which the data were processed; both are in the Source section. TrevantAI keeps **Bron: ClinicalTrials.gov (U.S. National Library of Medicine)** with the document.

Good to know

- Values such as status, type of study and section appear in the item as the source gives them, in English (e.g. RECRUITING, PHASE2).
- If you want the results themselves, open the study with **View at the source**.
- A registration says nothing about the quality or outcome of a study.
- If TrevantAI reports that a page cannot be requested directly, search again and browse from the beginning.
- With the NCT number and the date of processing it remains traceable which version was used.

Europe PMC (biomedical literature)

The European database of biomedical and life sciences literature; you add the description and abstract, and the full text for open access articles with a free licence.

> Key: none · Type of use: Description free, text protected by copyright (Descriptions free to use; abstracts may be protected by copyright) · Personal data: no · Where: at the source

Europe PMC is the European database of biomedical and life sciences literature, run by EMBL-EBI (the European Bioinformatics Institute). It contains the references from PubMed/MEDLINE, the articles from PubMed Central and preprints. For many articles the full text is available in open access. TrevantAI adds the

description and the abstract, and the full text if the licence allows it.

What can you use this source for?

- Medicine and science: collecting the literature on a subject, for example antimicrobial resistance in wastewater, and searching open articles in full.
- Government and policy: finding support for advice on public health or the environment.
- Legal: finding medical literature for a personal injury case or a disciplinary case.
- Business: following research on a drug, medical device or nutrient.

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **Europe PMC (biomedical literature)**.

- **Search words or identifier** (required): search words, usually in English, for example *antimicrobial resistance*. The words are sent to Europe PMC unchanged, so Europe PMC's query language (such as AND, OR and "phrases") also works. You can also enter a single identifier: a PMID (such as 12345678), a PMCID (PMC1234567), a DOI (10.… or DOI:10.…), an identifier with its source (MED:12345678, PMC:PMC1234567, PPR:… for preprints) or a link from Europe PMC, PubMed or PMC. You then get exactly that article.
- **Access**: **Open access only** shows only articles that are freely readable.
- **From year** and **Up to year**: the year of publication, in four digits. One of the two is also fine.

Choose **Search**. You get 20 results at a time, with the first authors, the journal, the year, whether there is open access and the DOI; use **More results** to browse further. Tick what you want (at most 20 at a time) and choose **Add to sources**.

What you get

Each article becomes one source with (the fixed headings in the item are in Dutch):

- **Description** (*Beschrijving*): title, authors, published in, year, first publication, DOI, PMID, PMCID, link, licence and whether there is open access;
- **Abstract** (*Samenvatting*), if there is one;
- for open access with a licence that also allows commercial use (CC0, CC BY, CC BY-SA or CC BY-ND): the **full text**, per heading of the article, with a breadcrumb trail (*title › Methods › Statistical analysis*). Tables become lines with " | " between the cells; the reference list and footnotes are left out;
- **Source** (*Bron*): the licence of the article and why the full text was or was not included.

With a non-commercial licence (NC) or without open access, TrevantAI includes only the description and the abstract; you can then read the full text at the source. The item is a **snapshot**.

Key

No key needed.

Privacy

The search query is sent to Europe PMC (EMBL-EBI). Please do not enter names of clients or patients or other confidential information. Europe PMC contains publications, no data on patients or contact persons; author names are public bibliographic data. By default, therefore, an external AI may also process the items if the AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The descriptions are free to use. Abstracts and full texts are covered by the copyright of the authors or the publisher, unless the article's licence says otherwise. Under CC BY you may reuse the text with attribution; the licence is shown with each item. TrevantAI keeps **Bron: Europe PMC** (Source: Europe PMC) with the document. In your work, cite the article itself.

Good to know

- Whether the full text is included depends on the licence of the article, not on the choice under **Access**.
- If an article has neither an abstract nor open text, TrevantAI reports that it has no readable text.
- If TrevantAI reports that a page cannot be requested directly, search again and browse from the beginning.
- Already added? Then you will see **Already in the sources**, with **View**.
- With the PMID, PMCID or DOI it remains traceable which article was used.

Unpaywall enrichment: if the **Unpaywall** enrichment is switched on in **Administration > Source connections**, TrevantAI for an article for which Europe PMC gives no full text (or only under a licence that does not allow it) looks for a legal open version via the DOI. If its licence also allows commercial reuse, the full text is added, page by page; otherwise **Source** says why not. See Unpaywall.

📖 PubMed (medical literature)

The biomedical literature database of the U.S. National Library of Medicine; you add the description, abstract and MeSH terms of a publication.

> Key: optional · Type of use: Description free, text protected by copyright (Descriptions free to use; abstracts may be protected by copyright) · Personal data: no · Where: at the source

PubMed is the biomedical literature database of the U.S. National Library of Medicine (NLM), with over 37 million references to articles from medical and life sciences journals (MEDLINE and more). Each publication has a description, often an abstract, and MeSH terms: the controlled keywords with which NLM indexes each article. TrevantAI retrieves the data through NCBI's E-utilities and adds them as a source.

What can you use this source for?

- Medicine and care: collecting the literature on a treatment, for example *asthma budesonide*, for a guideline or a second opinion.
- Science: searching with MeSH terms for a systematic review.
- Legal: finding medical publications for a personal injury case or a dispute about a treatment.
- Government and journalism: checking what is known about the health effects of a substance.

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **PubMed (medical literature)**.

- **Search words or identifier** (required): search words in English, for example *asthma budesonide*. PubMed's search syntax is also allowed, for example *asthma[mh]* for a MeSH term, *smith j[au]* for an author, and AND, OR and NOT. You can also enter a PMID (a number of at least five digits, optionally preceded by *PMID:*) or paste a link to pubmed.ncbi.nlm.nih.gov; you then get exactly that one publication.
- **From year** and **Up to year**: the year of publication, in four digits. One of the two is also fine.

Choose **Search**. You get 20 results at a time, with the first authors, the journal, the date, the DOI and the PMID; use **More results** to browse further, up to the first 10,000. Tick what you want (at most 20 at a time) and choose **Add to sources**.

What you get

Each publication becomes one source with these sections (the fixed headings in the item are in Dutch):

- **Description** (*Beschrijving*): title, authors, published in, volume, issue and pages, date, publication type, PMID, DOI, PMCID and the link to PubMed;
- **Abstract** (*Samenvatting*): for a structured abstract one section per part, with a breadcrumb trail (*title › Samenvatting › Methods*);
- **Subjects** (*Onderwerpen*): the MeSH terms (with subheadings) and the authors' keywords;
- **Source** (*Bron*): the attribution and where the full text can be found.

TrevantAI does not retrieve the full text here; it is at the publisher or, for open access, in PubMed Central (the PMCID is included). The item is a **snapshot**.

Key

Not needed. Without a key NCBI allows at most 3 requests per second; with your organisation's own free key, 10. TrevantAI keeps to that limit. TrevantAI is supplied without keys. The administrator enters the key under **Administration > Source connections**; how to request it is explained in the Help Centre under *Key for PubMed (optional)*.

Privacy

The search query is sent to NCBI (NLM), with the name TrevantAI and TrevantAI's contact address, never that of a user. Please do not enter names of clients or patients or other confidential information. PubMed contains descriptions of publications, no patient data; author names are public bibliographic data. By default, therefore, an external AI may also process the items if the AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The NLM descriptions are free to use. Abstracts are covered by the copyright of the authors or the publisher: quoting is allowed, but reproducing one in full in a publication may require permission. NLM asks for the attribution that TrevantAI keeps with the document: **Bron: PubMed, courtesy of the U.S. National Library of Medicine**.

Good to know

- Not every publication has an abstract; the description then says so.
- Search for a short series of digits (up to four) as a word instead; TrevantAI does not treat it as a PMID.
- If PubMed rejects the query, check the search syntax (brackets, quotation marks).
- If the limit has been reached, TrevantAI reports that PubMed asks you to wait a moment. If TrevantAI reports that PubMed rejects the key, the administrator should check the key.
- With the PMID it remains traceable which publication was used.

🌍 WHO Global Health Observatory

WHO health statistics by country and region, such as life expectancy, mortality and immunisation: find an indicator, choose country and years, and add the figures as a snapshot. Use under the WHO terms.

> Key: none · Type of use: Free to use (Under the WHO terms for datasets: for public health purposes, with attribution; not to be sold or used for commercial promotion) · Personal data: no · Where: at the source

The **Global Health Observatory** (GHO) is the health statistics database of the **World Health Organization (WHO)**. It contains more than two thousand **indicators**, such as life expectancy, child mortality, causes of death, immunisation coverage, smoking, overweight, air pollution and the number of physicians, by country, by WHO region and for the world. TrevantAI uses the GHO's official OData service.

What can you use this source for?

- Medical and clinical: compare figures on a disease or risk factor in the Netherlands with those of other countries.
- Government and policy: support health targets with internationally comparable figures.
- Research: record a series over years and countries as a source, with the lower and upper bounds of the estimate.
- Journalism: check a claim about health in the world.

Searching

- **Search words or indicator code** (required): words from the name, for example *life expectancy* or *measles immunization*. Every word must appear in the name. You can also enter a code such as *WHOSIS_000001*.
- **Country or area**: for example the Netherlands, Belgium or *World (total)*. Leave it empty for all countries and regions.
- **From year** and **Up to year**: for example *2010* and *2024*. Leave them empty for all years.

Country and years apply to the item you add. Each result shows the code and your selection. Short names are listed first. You get 20 results at a time; use **More results** to continue.

What you get

Tick a result and choose **Add to sources**. The item is a **snapshot** with the retrieval date. It is called *indicator name (code)* and has these parts, each with the breadcrumb in its heading:

- **Beschrijving** (description): the query (area and years), the date of the WHO's last update and the number of values.

- **Gegevens (deel 1, 2, ...)** (data, part 1, 2, ...): a table with area, period, the breakdown of the indicator (for example sex or age group), the value and, where present, a WHO comment. A value such as *81.4 [81-81.8]* is the estimate with the lower and upper bounds in brackets. Names of countries and groups are written out in full.
- **Bron** (source): the attribution and the conditions.

An item contains at most **20,000 values**. If the selection is larger, the description states that the item has been truncated. Then choose a country or fewer years.

Key

No key is needed. Conditions do apply, see below. The administrator switches the connection on under **Administration > Source connections** and confirms the conditions when doing so.

Privacy

The search query is sent to the WHO. Do not use client names or other confidential information in the query. The GHO contains only aggregated figures by country or region, no personal data. That is why, by default, an external AI may also process these items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

The WHO makes its datasets available under its own **terms for datasets**. You may use, reproduce and adapt the data free of charge for **public health purposes**. Not allowed: using the data for or in connection with the promotion of a commercial enterprise, and selling or transferring them outside that use. Attribute the source: WHO, the name of the dataset, the year, the date of access and the countries that provided the data. Do not suggest that the WHO endorses your use. The administrator confirms these terms when switching the connection on. Whether your organisation's use falls within them is for the organisation itself to assess; TrevantAI does not block the source for commercial use. With the document, TrevantAI keeps the attribution "Bron: WHO, Global Health Observatory" (Source: WHO, Global Health Observatory).

Good to know

- The WHO has announced that this service will be replaced by a new one. Until that is publicly available, TrevantAI uses the current service. If searching suddenly stops working, please report it; the connection will then need to be adapted.
- Indicator names are in English. Search with English words.
- Not every indicator has figures for every country and year. If you get no values, choose another country or leave the years empty.
- Since late October 2026 the WHO service only returns errors (seen on 4 November 2026). The address of its successor is not yet known; until the connection has been adapted, this source does not work.
- Under **View** you see the source, the code, the link, the retrieval date and the search query (**Found with**).

💊 EMA (European Medicines Agency)

Medicines, shortages, DHPCs and referrals of the European Medicines Agency, from the official data files; for a medicine also the product information (SmPC and package leaflet) in the screen language. Can also be followed.

> Key: none · Type of use: Free to use (Reuse permitted, including commercially, with the EMA acknowledged as the source in each copy (not for third-party content)) · Personal data: no · Where: at the source

The **European Medicines Agency (EMA)** evaluates and monitors medicines in the European Union. TrevantAI uses the official **data files** of the EMA website, updated twice a day: centrally authorised medicines (EPARs), shortages, letters to healthcare professionals (DHPCs) and referral procedures. For a medicine, TrevantAI also fetches the **product information**: the summary of product characteristics (SmPC), the labelling and the package leaflet.

What can you use this source for?

- Medical and pharmaceutical: read the indications, the dosage or the warnings of a medicine in the official European product information.
- Pharmacy and procurement: check whether there is a shortage and whether alternatives are available.
- Government and policy: follow safety procedures (referrals) and letters to healthcare professionals.
- Science and journalism: record when a medicine was authorised, suspended or withdrawn.

Searching

- **Name, active substance or identifier** (required, unless you enter a therapeutic area): the name of a medicine, an active substance or an ATC code, in English as the EMA writes them, for example *inclisiran*. All words must occur; capitals and accents do not matter. You can also enter a product number (*EMEA/H/C/005333*) or a link to an EMA page.
- **Type**: medicine (EPAR), shortage, letter to healthcare professionals (DHPC) or referral procedure. Empty: all types.
- **Status**: for example authorised, withdrawn or lapsed, suspended, ongoing or concluded. Not every status fits every type: a shortage is ongoing, resolved, or the medicine has been discontinued.
- **Therapeutic area**: words from the area (MeSH), for example *Hypercholesterolemia*.

TrevantAI searches the files itself; the query is not sent to the EMA. Names that start with your search words come first, then what was updated most recently. You get 20 results at a time; use **More results** to see more.

What you get

Tick a result and choose **Add to sources**. The item has these parts, each with the breadcrumb in its heading:

- **Beschrijving** (description): the data from the file, such as status, active substance, therapeutic area, ATC code, marketing authorisation holder, particulars (for example orphan medicine or additional monitoring) and the key dates.
- **Therapeutische indicatie** (for a medicine): what the medicine is authorised for.
- **Productinformatie** (for a medicine): the PDF with SmPC, labelling and package leaflet, page by page. In the screen language if the EMA has it, otherwise in English; the description states which language was included.
- **Bron** (source): the attribution and what has and has not been included.

The product information is limited to **200 pages** and 600,000 characters. If it is larger, the description states that the item has been truncated; the full PDF is available from the EMA. For a shortage, a DHPC or a referral, the item consists of the data from the file; the letter and the documents of the procedure are on the EMA page.

Following

Follow… creates a subscription: TrevantAI fetches the files at set times and puts in *New arrivals* what has been newly published or updated since the previous time (according to the dates in the files, by day). Filters: **Type**, **Status**, **Therapeutic area**, **Only newly published** (no updates of existing pages) and **Words** (in the name and the description). TrevantAI checks the filters itself; nothing is sent to the EMA. If the page of an item you already have has been updated, TrevantAI reports this with that item. See Following sources with a subscription.

Key

Not needed. The administrator switches the connection on under **Administration > Source connections**.

Privacy

The query and the filters stay in TrevantAI; only the requests for the files and the PDF go to the EMA. The files contain medicines, companies and procedures, no personal data. That is why, by default, an external AI may also process these items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

Under the EMA reuse policy (*Legal notice*), documents and data may be reproduced and distributed in whole or in part, **including commercially**, provided the EMA is acknowledged as the source in each copy. This does not apply to third-party content or to the EMA logo. With the document, TrevantAI keeps the attribution "Bron: European Medicines Agency (EMA)" and the link to the page.

Good to know

- The data are in English, as the EMA publishes them; only the product information is in the screen language.
- The medicines file contains only **centrally authorised** medicines (through the EMA). Nationally authorised medicines are listed by the national medicines authority, in the Netherlands the CBG.
- TrevantAI keeps the files for six hours, so a change at the EMA becomes visible at most six hours later.
- Do not use the product information as the only basis for treatment decisions; consult the current version at the EMA.
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.
- The first search retrieves the EMA lists (a few megabytes); after that, searching is fast.
- Under **View** you see the source, the identifier, the link, the date of retrieval and the search (**Found with**).

🧪 openFDA (U.S. FDA)

Drug labels, recalls of drugs, devices and food, and overviews of reported adverse events (counts only) from the U.S. FDA. Recalls and labels can also be followed.

> Key: optional · Type of use: Free to use (Free to use (CC0)) · Personal data: no · Where: at the source

openFDA is the open data service of the U.S. **Food and Drug Administration (FDA)**. TrevantAI uses three parts: the **drug labels** of medicines in the USA, the **recalls** of drugs, medical devices and food (the weekly Enforcement Report), and the reports of **adverse events** (FAERS), only as an overview of counts.

What can you use this source for?

- Medical and pharmaceutical: read the U.S. label of a drug, for example the boxed warning or the interactions.
- Quality and procurement: check whether a product or a batch has been recalled, and why.
- Pharmacovigilance: see which adverse events have been reported most often for a drug.
- Journalism and research: follow the recalls of a company or a product group.

Searching

- **What**: drug label, drug recall, medical device recall, food recall or adverse events (overview of counts). Empty: drug labels.
- **Search words or identifier**: English words, for example *metformin*; all words must occur. You can also enter an identifier: the set ID of a label, a recall number such as *D-0123-2026* or a DailyMed link. For adverse events: the generic or brand name of one drug.
- **Class (recalls)**: class I (serious hazard), II or III (hazard unlikely).
- **From** and **Up to**: for labels the effective date, for recalls the date in the Enforcement Report, for adverse events the date on which the FDA received the report. For adverse events the period is at least 30 days.

You get 20 results at a time; use **More results** to see more. For adverse events there is one result: the overview for the drug and the period.

What you get

Tick a result and choose **Add to sources**. Each item has a **Beschrijving** (description) and a **Bron** (source), with in between:

- For a **label**: a section per part, such as *Indications and usage*, *Warnings and cautions*, *Adverse reactions* and *Drug interactions*, in the order of the label.
- For a **recall**: the reason, the product (batches, quantity, distribution) and, in the description, class, status, company, location and dates. TrevantAI leaves out the street address of the company.
- For **adverse events**: a **snapshot** with the date of retrieval: the number of reports, how many of them serious, and a table of the 100 most frequently reported reactions (MedDRA terms) with the number of reports. Individual reports are never fetched. If there are fewer than 10 reports, there is no overview; a reaction with fewer than 5 reports is not in the table, and a number below 5 is shown as "fewer than 5". This way no report can be traced back to one patient.

A report of an adverse event **does not prove a causal relationship**: according to the FDA there is no certainty that the drug caused the reaction, and the reports say nothing about how often a reaction occurs. This warning is included in every overview. A label is limited to 400,000 characters; if it is larger, the description states that it has been truncated.

Following

Follow… creates a subscription to recalls or labels: TrevantAI asks at set times what has been added since the previous time (recalls by the date in the Enforcement Report, labels by the effective date) and puts it in *New arrivals*. Filters: **What** (required; after a search it shows what you searched for, and without a choice or for adverse events: drug labels, as in searching), **Class (recalls)** and **Words**; they are included in the request to

openFDA. Adverse events cannot be followed: an overview of counts is not a new item. A date in the future at the source does not move following ahead. If there are more than 25,000 items with the same date, TrevantAI skips the rest of that day. See Following sources with a subscription.

Key

Not needed. Without a key, openFDA allows 240 requests per minute and 1,000 per day (per IP address); with a free key 240 per minute and 120,000 per day. TrevantAI respects the limit per minute. TrevantAI does not ship with keys; the administrator enters the organisation's own key under **Administration > Source connections**. A key can be requested free of charge at open.fda.gov/apis/authentication.

Privacy

The search query is sent to openFDA (api.fda.gov). Do not enter names of patients or other confidential information. Labels and recalls are about products and companies; for adverse events TrevantAI includes only counts. That is why, by default, an external AI may also process these items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

The openFDA data are in the **public domain** (CC0): you may copy, adapt and distribute them, including commercially. An exception is the GMDN data for medical devices; this connection does not include them. The FDA asks for attribution; with the document, TrevantAI keeps "Data provided by the U.S. Food and Drug Administration (openFDA)". Do not suggest that the FDA endorses your use.

Good to know

- The data are in English.
- openFDA is updated roughly weekly. The effective date of a label is the date given by the manufacturer; a label may appear at openFDA only later and is then missed by Following. To follow labels, DailyMed is more reliable.
- Do not rely on openFDA alone for decisions about medical care; the FDA says so itself.
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.
- Under **View** you see the source, the identifier, the link, the date of retrieval and the search (**Found with**).

📋 DailyMed (U.S. product labeling)

The current U.S. product labeling (SPL labels) of medicines via DailyMed of the NLM: search by name or manufacturer, one section per part of the label. New labels and versions can also be followed.

> Key: none · Type of use: Free to use (Public government information) · Personal data: no · Where: at the source

DailyMed is the service of the **U.S. National Library of Medicine (NLM)** with the current product labeling of medicines in the United States: the labels (Structured Product Labeling, SPL) that manufacturers submit to the FDA. TrevantAI uses the official DailyMed web services.

What can you use this source for?

- Medical and pharmaceutical: read the full U.S. product labeling of a medicine, with indications, dosage, warnings and adverse reactions.

- Comparing: put the U.S. information next to the European information (EMA).
- Quality and registration: check which version of a label applies and since when.

Searching

- **Medicine name or set ID**: the generic or brand name, for example *metformin*. You can also enter a set ID or a DailyMed link.
- **Manufacturer**: the name of the manufacturer or labeler, for example *Pfizer*.
- **Type**: prescription or over-the-counter.

Enter a name, a manufacturer or both. You get 20 results at a time; use **More results** to see more. Each result shows the manufacturer, the version and the date of publication.

What you get

Tick a result and choose **Add to sources**. The item is the current version of the label, with these parts, each with the breadcrumb in its heading:

- **Beschrijving** (description): the title, the document type, the manufacturer, the version, the date of this version and the set ID.
- A section per part of the label, with the title from the label, for example *1 INDICATIONS AND USAGE* or *2.1 Recommended Dosage*; subsections appear in the breadcrumb. Tables and lists are included as text.
- **Bron** (source): the attribution and what has and has not been included.

Not included: the product data (ingredients and packaging as data), the Highlights (a summary of the same text) and images. An item is limited to 500 sections and 800,000 characters; if the label is larger, the description states that it has been truncated.

Following

Follow… creates a subscription: TrevantAI asks at set times, day by day, which labels DailyMed has published since the previous time (a new label or a new version) and puts them in *New arrivals*. Filters: **Medicine name or set ID**, **Manufacturer** and **Type** (included in the request to DailyMed) and **Words** (in the title and the description). Without filters there are hundreds per day. TrevantAI uses the date in the US, where DailyMed publishes: a day is only closed once it is over there, so that what still appears in the evening in the US is not missed. See Following sources with a subscription.

Key

Not needed. The administrator switches the connection on under **Administration > Source connections**. TrevantAI waits a quarter of a second between two requests.

Privacy

The search query is sent to DailyMed (NLM). Do not enter names of patients or other confidential information. Labels contain product information, no patient data. That is why, by default, an external AI may also process these items, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

DailyMed is a service of the U.S. government; government information is free of copyright there. The labels themselves were submitted to the FDA by the manufacturers; the FDA publishes the same labels through openFDA as public domain. You may therefore **use the items freely**, including commercially, with attribution. With the document, TrevantAI keeps the attribution "Bron: DailyMed, U.S. National Library of Medicine".

Good to know

- The labels are in English.
- According to DailyMed, the labeling in use has not always been verified by the FDA, and the NLM does not review the content beforehand. The most recently approved version is available at Drugs@FDA.
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.
- Under **View** you see the source, the identifier, the link, the date of retrieval and the search (**Found with**).

Heritage and knowledge

🖼️ Europeana (European cultural heritage)

Descriptions of millions of heritage objects from European museums, archives and libraries; you add the description with maker, date, subjects and rights.

> Key: required (own key) · Type of use: Free to use (Free to use (CC0)) · Personal data: no · Where: at the source

Europeana is the European platform for cultural heritage. Museums, archives and libraries from all over Europe supply the descriptions of their collections: millions of artworks, photographs, prints, books, newspapers, sound and film recordings and 3D objects. Each object has, among other things, a title, maker, date, subjects, the institution and a rights statement. TrevantAI adds the description of an object as a source; images and files are not added.

What can you use this source for?

- Heritage and history: collecting objects about a person, place or event from collections in different countries.
- Journalism and communication: finding images and checking whether they are freely reusable.
- Legal: recording the rights status of an image or recording, with the institution that manages it.
- Education and science: bringing together sources for a lesson, exhibition or research project.

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **Europeana (European cultural heritage)**. Without a key no search is possible; TrevantAI then reports that the administrator must enter a key (see **Key**).

- **Search words or link** (required): search words, for example a maker, place or subject, such as *Vermeer*. Descriptions are in many languages; if needed, also search in English or in the language of the country of origin. You can also paste a link to a single object, such as https://www.europeana.eu/nl/item/2021672/resource_document_mauritshuis_670, or a Europeana identifier (/2021672/...). That object then appears as the only result; if Europeana does not know the object, the list stays empty.
- **Type of object**: Image, Text, Sound, Video or 3D.
- **Reuse**: Freely reusable, Reuse with restrictions or Only with permission.

Choose **Search**. You get 20 results at a time, with maker, year, type, institution and country; use **More results** to browse further, up to at most the first 1,000. Too many? Choose a type of object or add a word. Tick what you want (at most 20 at a time) and choose **Add to sources**.

What you get

Each object becomes one source with these sections (the headings in the item are in Dutch):

- **Description** (*Beschrijving*): title, other titles, type of object, material and format, language and the institution's description;
- **Maker and date** (*Maker en datering*): maker, contributors, publisher, date and period;

- **Subjects** (*Onderwerpen*): subjects, place and coverage;
- **Rights and provenance** (*Rechten en herkomst*): the rights statement with a plain-language explanation, the rights according to the institution, the institution, the country, the provenance, the identifier at the institution and the link to the object there;
- **Source** (*Bron*): the attribution.

Each section starts with a breadcrumb trail (*title › Maker en datering*). Where texts exist in several languages, TrevantAI prefers Dutch, then English. The item is a **snapshot** with the date of retrieval.

Key

Required: your organisation's own key, one per installation. TrevantAI is supplied without keys. At Europeana, choose a project key (Project API key); a personal key is meant for trying out and has lower limits. Europeana does not mention any costs. The administrator enters the key under **Administration > Source connections**; how to request it is explained in the Help Centre under *Key for the Europeana API*.

Privacy

The search query is sent to Europeana; the key is sent in the request header, not in the link. Please do not enter client names or other confidential information. Europeana contains descriptions of heritage, published by the institutions themselves; names of makers are public information about the work. By default, therefore, an external AI may also process the items if the AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The descriptions (metadata) are free to use (CC0). The object itself is covered by the institution's rights statement, for example Public Domain Mark, CC BY or In copyright; it is shown with each item, with an explanation. If you want to use an image, follow that statement. TrevantAI keeps **Bron: Europeana** (Source: Europeana) with the document, with the name of the institution.

Good to know

- How detailed a description is varies greatly between institutions.
- TrevantAI does not retrieve images, sound or video; open the object with **View at the source**.
- If an object has no title or description, TrevantAI cannot add it.
- Already added? Then you will see **Already in the sources**, with **View**.
- With the Europeana identifier (/dataset/object) it remains traceable which object was used.

📚 Wikipedia (encyclopaedia)

Articles from the free encyclopaedia in seven languages; you add the text per heading, with the version of the article, so it remains traceable which text was used.

> Key: none · Type of use: Free, with attribution and share-alike (Free to use with attribution and share-alike (CC BY-SA 4.0)) · Personal data: may contain · Where: at the source

Wikipedia is the free encyclopaedia of the Wikimedia community, written by volunteers. TrevantAI can retrieve articles from the Wikipedia in seven languages: Dutch, English, German, French, Spanish, Italian and Portuguese. The article is added as a source per heading, with the number of the version you retrieved.

What can you use this source for?

- Government and policy: quickly gathering background on a subject, organisation or area, for example the Afsluitdijk.
- Journalism: a first overview of an event or person, and searching further through the article's sources.
- Heritage and history: comparing articles about a place or period in several languages.
- Business and legal: an explanation of a concept, technique or institution as a starting point.

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **Wikipedia (encyclopaedia)**.

- **Search words or link** (required): search words, for example *Afsluitdijk*. You can also paste a link to an article, such as https://nl.wikipedia.org/wiki/Afsluitdijk (also the mobile version or a link with ?curid=). You then get exactly that one article; a redirect is followed.
- **Language of Wikipedia**: Dutch, English, German, French, Spanish, Italian or Portuguese. If you choose no language, TrevantAI searches the Dutch Wikipedia. For a link, the language of the link applies.

Choose **Search**. Wikipedia searches the text of the articles. You get 20 results at a time, with a snippet of text, the number of words and the date of the last change; use **More results** to browse further. Only articles are found, no talk or user pages. Tick what you want (at most 20 at a time) and choose **Add to sources**.

What you get

Each article becomes one source: the introduction and then each heading of the article as a separate section, with a breadcrumb trail (*Afsluitdijk › History › Construction*). Finally a section **Source** (*Bron*, the fixed headings are in Dutch) with the language, the version number and the date of that version, a permanent link to exactly that version (?oldid=) and a link to the article's history, which shows who contributed to it. Tables, infoboxes and lists consisting only of links are not included. The item is a **snapshot**: the article may change later, but the version used remains readable through the permanent link.

Key

No key needed. TrevantAI identifies itself to Wikipedia with a recognisable name and TrevantAI's contact address, as the Wikimedia Foundation asks.

Privacy

The search query is sent to Wikipedia (Wikimedia Foundation). Please do not enter client names or other confidential information. Wikipedia also contains biographies of living people. By default, therefore, **only the local AI** processes items from this source. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The text is covered by Creative Commons Attribution-ShareAlike 4.0 (CC BY-SA 4.0). You may reuse it, also commercially, provided you credit the source and the contributors and share any adaptation under the same licence. TrevantAI keeps **Bron: Wikipedia, de vrije encyclopedie (bijdragers aan het artikel)** (Source: Wikipedia, the free encyclopaedia, contributors to the article) with the document, with the link to the version.

Good to know

- Wikipedia is written by volunteers; check important facts in the sources that the article cites.
- The same title may be a different, longer or shorter article in another language.
- For a very long article, check whether the text is complete; the permanent link lets you read the version used at the source.
- Already added? Then you will see **Already in the sources**, with **View**.
- With the language and the version number it remains exactly traceable which text was used.

Wikisource (historical texts)

Wikimedia's free library: Church Fathers, Reformers, confessions, old Bibles and more, in nine languages; a work comes with its subpages.

> Key: none · Type of use: Free, with attribution and share-alike (Free to use with attribution and share-alike (CC BY-SA 4.0)) · Personal data: no · Where: at the source

Wikisource is Wikimedia's free library: volunteers publish texts that are free to use, often with a scan of the original. For theology these include Church Fathers, Reformers, confessions and catechisms in their original language and old Bible translations; in addition laws, letters, speeches and literature.

Languages

Dutch, English, German, French, Spanish, Italian, Portuguese, Latin and Greek (including Ancient Greek). Each language is a separate Wikisource with its own texts: German confessional writings are on the German one, Latin Church Fathers on the Latin one. Therefore choose the language of the original.

Searching

Type search words and choose the language, or paste a link to a page on Wikisource. The search query is sent to Wikisource, so please do not enter confidential information.

A work with subpages

Many works on Wikisource are split into subpages, for example the Lord's Days of a catechism or the chapters of a book. If you choose the main page, the subpages come with it, in the right order (at most 40, or about 500,000 characters); anything left over can be added separately. A subpage can also be chosen on its own.

What happens to the text

TrevantAI takes over the text without Wikisource's headers, navigation, page numbers and footnote markers. The version (revision number) of every page is recorded; the *Source* section links to exactly that version.

Licence and privacy

Per item: public domain if the page says so; a translation by volunteers, or a page without a licence statement, as CC BY-SA 4.0 (free to use with attribution and share-alike). If the page only has a US public-domain template (PD-US, *not renewed*), the item gets the label *Public domain in the US only*: in Europe the

text may still be protected (70 years after the death of the author or translator). Use such an item within the organisation for your own research, and do not share it, or a translation or summary of it, outside. For a work with subpages the strictest licence applies. No personal data: these are historical texts.

See also: Creeds and catechisms · Theological sources: overview

🏛️ Dutch National Monuments Register (RCE)

The register of national monuments of the Cultural Heritage Agency of the Netherlands: search by number, name, place or street and add the register entry, with the statement of significance. Free to use with attribution (CC BY 4.0).

> Key: none · Type of use: Free to use (Free to use with attribution (CC BY 4.0)) · Personal data: no · Where: at the source

The **National Monuments Register** contains all national monuments of the Netherlands: buildings, churches, mills, bridges, gardens and sites protected by the State. The **Cultural Heritage Agency of the Netherlands (RCE)** maintains the register. TrevantAI uses the official **Linked Data** of the RCE.

What can you use this source for?

- Spatial planning and permits: check whether a building is a national monument and why.
- Heritage and history: record the description of a monument as a source.
- Real estate and management: list the monuments in a street or municipality.

Searching

- **Monument number**: the national monument number, for example *518611*. A link from the monument register also works.
- **Search words**: a word from the name or the description, for example *molen* (mill) or *boerderij* (farm).
- **Place or municipality**: the town or municipality, exactly as it is named, for example *Utrecht*.
- **Street**: (part of) the street name, for example *Oudegracht*.

Fill in at least one field; more fields narrow the search. Results are sorted by monument number.
You get 20 results at a time; use **More results** to continue.

What you get

The item is called *name (national monument number)* and has these parts:

- **Description**: number, name, addresses, town, municipality, province, cadastral reference (if the RCE lists it), legal status, kind of monument, function, type, building period, date of registration and the link to the monument register.
- **Statement**: the text from the register; for a national monument this is the *statement of significance*, which explains why the monument is protected.
- **Source**: the attribution and the licence.

A very long statement is cut off after 100,000 characters; the Description then says so.

Key

No key is needed. The administrator turns the connection on under **Administration > Source connections**.

Privacy

The search goes to the RCE. The register describes buildings and sites; owners and contact details are not in the open data. A cloud AI may therefore also process these items by default, if your organisation's AI policy allows cloud AI.

Licence and attribution

The data are under **CC BY 4.0**: free to use, also commercially, with attribution to the Cultural Heritage Agency of the Netherlands. TrevantAI does not include images; they have their own rights. With the document, TrevantAI stores the attribution "Bron: Rijksdienst voor het Cultureel Erfgoed (RCE), Rijksmonumentenregister".

Good to know

- In case of doubt, the monument register itself applies (monumentenregister.cultureelerfgoed.nl).
- This source cannot be followed: the open data have no modification date per monument.
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.
- Search with a **place**. A monument has no name of its own; the title is its function (for example *house*) with the address.
- Under **View** you see the source, the identifier, the link, the retrieval date and the search (**Found with**).

🖼️ Rijksmuseum (collection)

Descriptions of artworks and objects in the Rijksmuseum collection: search by title, maker, type, dating or object number and add the description, with a link to the image. Free to use (CC0); can also be followed.

> Key: none · Type of use: Free to use (Free to use (CC0)) · Personal data: no · Where: at the source

The **Rijksmuseum** in Amsterdam makes its collection data openly available through the **Rijksmuseum Data Services** (data.rijksmuseum.nl). For each object they include the title, object number, maker, dating, type, material, technique, dimensions and the museum's texts, among other things. TrevantAI uses the official search service (Search API) and retrieves the description of each object as Linked Art.

What can you use this source for?

- Heritage and history: bringing together objects by one maker, from one period or on one subject.
- Education and communication: recording an artwork with the museum's description as a source.
- Journalism: checking what the museum says about a work, its maker and its provenance.
- Legal: recording the provenance of an object as the museum describes it.

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **Rijksmuseum (collection)**. Fill in at least one field; all fields filled in must match:

- **Title**, **Maker**, **Type of object**, **Material**, **Technique** and **Words in the description**: one or more words, for example *Nachtwacht*, *Rembrandt van Rijn*, *schilderij* or *doek*. Dutch words usually give more results than English ones.
- **Dating**: a year, such as *1642*, or years with question marks, such as *16??* for the seventeenth century.

- **Object number or link**: an object number, such as *SK-C-5*, or a link from id.rijksmuseum.nl or from the collection website. That object then appears as the only result.
- **Set**: a part of the museum's collection, for example Dutch paintings of the seventeenth century.
- **Image**: **Only with an image** shows only objects with an available image.

You get 20 results at a time, with maker, dating and object number; use **More results** to browse further. Tick what you want (at most 20 at a time) and choose **Add to sources**.

What you get

Each object becomes one source with the sections:

- **Description**: title, object number, maker(s), dating, type of object, material, technique, dimensions, the link to the image (IIIF) and the museum's description;
- **Inscriptions and marks** and **Provenance**, if the museum records them;
- **Source**: the attribution, the licence and what is and is not included.

Each section starts with a breadcrumb (*title › Provenance*). If the texts are in Dutch and English, TrevantAI chooses the language of your screen; for any other language, English. The image itself is not retrieved; the link is in the Description. The item is a **snapshot** with the date of retrieval.

Following

With *Follow…* you create a subscription to new and changed objects: at set times TrevantAI asks the Rijksmuseum what has been added or updated since the previous time (via OAI-PMH). Filters: *Set* (a part of the museum's collection) and *Words* (in the title and the summary). Objects that the museum removes are reported. See Following sources with a subscription.

Key

No key is needed. The administrator switches the connection on under **Administration > Source connections**.

Privacy

The search query goes to the Rijksmuseum. Do not enter names of clients or other confidential information. The descriptions are about artworks and objects and are published by the museum itself; names of makers are public information about the work. Therefore an external AI may also process these items by default, if your organisation's AI policy allows cloud AI. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

The Rijksmuseum's collection data are open data and free to use (CC0). The museum does ask you to credit the Rijksmuseum as the source. The images have their own statement: most are in the public domain, some fall under CC BY 4.0 or are still protected. TrevantAI includes only the link to the image. With the document, TrevantAI keeps the attribution **Bron: Rijksmuseum, Amsterdam**.

Good to know

- The search service only finds museum objects, not the books and archival records of the Rijksmuseum library.

- How detailed a description is varies per object.
- Already added? Then **Already in the sources** is shown, with **View**.
- The identifier of the object (the number in id.rijksmuseum.nl) makes it possible to trace which object was used.
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.

🎨 RKD (artists and works of art)

Artists from RKDartists and works of art from RKDimages of the RKD – Netherlands Institute for Art History: search by name, title or subject and add the description. Free to use with attribution (ODC-BY).

> Key: none · Type of use: Free to use (Free to use with attribution (Open Data Commons Attribution License, ODC-BY 1.0)) · Personal data: may contain · Where: at the source

The **RKD – Netherlands Institute for Art History** holds the largest art-historical documentation on Western art. **RKDartists** describes more than two hundred thousand artists, workshops and institutions; **RKDimages** describes works of art, mainly Dutch paintings, drawings and prints. TrevantAI uses the official **Linked Open Data** of the RKD.

What can you use this source for?

- Art history and museums: record data on an artist or a work as a source.
- Provenance research and trade: check the attribution, date and collection of a work.
- Education and journalism: look up the life and places of work of an artist. Teachers and pupils are not in the RKD's open data, only on rkd.nl.

Searching

- **Name, title or subject**: (part of) the name of an artist, for example *Judith Leyster*; for works of art, words from the title or subject. Every word must appear. A link or identifier of the RKD (*artists/…*, *images/…*) also works.
- **Search in**: empty or *Artists* for RKDartists, *Works of art* for RKDimages.
- **Maker (for works of art)**: (part of) the name of the maker, also without a title.

You get 20 results at a time; use **More results** to continue.

What you get

An **artist** is called *name (RKDartists number)* and has the parts **Description** (names, kind, gender, nationality, birth and death with place, qualifications and places of work with period, the link to the record), **Description by the RKD**, **Links** (to other sources) and **Source**.
A **work of art** is called *title (RKDimages number)* and has the parts **Description** (title, numbers, kind of object, maker, date, technique, material, dimensions, collection), **Iconography**, **Notes** and **Source**. TrevantAI does not include images or addresses.

Key

No key is needed. The administrator turns the connection on under **Administration > Source connections**.

Privacy

The search goes to the RKD. RKDartists also contains **living artists**, with date and place of birth.
Items from this source therefore go to the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

The RKD makes the data available under the **Open Data Commons Attribution License (ODC-BY) 1.0**: free to use, also commercially, with attribution to the RKD and preferably a link to the record (*https://data.rkd.nl/artists/...*). Images are not covered. With the document, TrevantAI stores the attribution "Bron: RKD – Nederlands Instituut voor Kunstgeschiedenis".

Good to know

- Attributions and dates may change; the RKD updates them. An added item is a snapshot.
- This source cannot be followed: the RKD offers no official way to ask what is new or changed.
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.
- Artworks are searched via the **maker**: the name of the artist, preferably in full (*Johannes Vermeer*). Words then search the Dutch and English titles of that artist's works.
- A qualification such as *after Rembrandt* means a copy or imitation, not a work by Rembrandt himself.
- A new search can take a few seconds; the same search afterwards is faster.
- Under **View** you see the source, the identifier, the link, the retrieval date and the search (**Found with**).

🗄️ Nationaal Archief (archival inventories)

The inventories of the archives in the Nationaal Archief: search by access number or by words in the name of the archive and add the inventory, with the introduction and the inventory numbers per section. Free to use (CC0); can also be followed.

> Key: none · Type of use: Free to use (Free to use (CC0)) · Personal data: may contain · Where: at the source

The **Nationaal Archief** (National Archives of the Netherlands) in The Hague holds the archives of central government and of organisations and persons of national importance. Each archive has an **inventory**: a description of the archive and a list of the records (the inventory numbers), arranged in sections. Each inventory has a fixed **access number**, for example 2.05.170. TrevantAI retrieves the inventories through the official open data service of the Nationaal Archief (OAI-PMH, in EAD).

What can you use this source for?

- History and heritage: checking which records an archive contains before consulting it.
- Government and policy: getting an overview of the archives of a ministry or agency over a period.
- Journalism: finding files on a subject and recording the inventory numbers in order to consult them.
- Legal: recording in which archive and under which inventory number a record is kept.

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **Nationaal Archief (archival inventories)**. The fields:

- **Access number or link**: the access number, such as *2.05.170*, or a link from nationaalarchief.nl (/onderzoeken/archief/2.05.170). That inventory then appears as the only result.
- **Words in the name of the archive**: one or more words, for example *Hoge Raad van Adel*. TrevantAI searches the names of all inventories; each word must occur.

The Nationaal Archief has no open search service for the content of the inventories. To search by subject, first look up the archive on nationaalarchief.nl and enter the access number here. You get 20 results at a time; use **More results** to browse further.

What you get

Each inventory becomes one source with the sections:

- **Description**: title, access number, creator(s), period, extent, language, the number of inventory numbers, the link and a summary;
- **Introduction**: the parts of the introduction, such as the history of the creator, the custodial history, access conditions and how to cite;
- **Inventory**: the inventory numbers with title, dates and extent, per section, with the breadcrumb (*title › Inventory › section › subsection*). A record with restricted access is marked *[beperkt openbaar]*;
- **Source**: the attribution and the licence.

An item contains at most **400 sections** and **600,000 characters**. If the inventory is larger, the Description says that it has been truncated and which part is the first to be missing; the full inventory is on nationaalarchief.nl. TrevantAI does not retrieve scans or the records themselves. The item is a **snapshot** with the date of retrieval.

Following

With *Follow…* you create a subscription to new and changed inventories: at set times TrevantAI asks the Nationaal Archief which inventories have been added or updated since the previous time (via OAI-PMH). Filters: *Access number* (follow one inventory) and *Words* (in the name of the archive). An inventory that the Nationaal Archief removes is reported. See Following sources with a subscription.

Key

No key is needed. The administrator switches the connection on under **Administration > Source connections**.

Privacy

The search query goes to the Nationaal Archief. Do not enter names of clients or other confidential information. Inventories, especially of twentieth-century archives, may contain names of persons, for example in the titles of files. Therefore items from this source are processed by the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The archival inventories of the Nationaal Archief are open data and free to use (CC0); attribution is appreciated. A small part of the inventories was made by others and has not been released. If an inventory carries a statement other than CC0, TrevantAI includes only the description of the archive and says so in the Source section. The GDPR continues to apply to personal data in an inventory. With the document, TrevantAI keeps the attribution **Bron: Nationaal Archief**.

Good to know

- An inventory describes the records, not their content. To consult a record, use the access number and the inventory number at the Nationaal Archief.
- Records with restricted access do appear in the inventory, but cannot always be consulted.
- Already added? Then **Already in the sources** is shown, with **View**.
- The access number and the inventory number make it possible to trace which record is meant.
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.
- Search with the **access number** (for example *2.05.170*): the open service of the Nationaal Archief gives no names of archives, and not every inventory is in it. For the same reason **Words** does not work when following; follow by access number.

🎞️ Sound & Vision (Open Images)

The descriptions of the videos on Open Images from Sound & Vision, such as Polygoon newsreels: search with words within a collection or by link and add the description, with the licence of the video. Descriptions free to use (CC0); can also be followed.

> Key: none · Type of use: Free to use (Free to use (CC0)) · Personal data: may contain · Where: at the source

Open Images (Open Beelden, openbeelden.nl) is the open platform of the **Netherlands Institute for Sound & Vision**. It offers thousands of freely reusable videos, such as Polygoon newsreels, information films and amateur films, also from other institutions such as Eye Filmmuseum. For each video it gives the title, date, maker, broadcaster or publisher, summary, keywords and the licence, among other things. TrevantAI uses the official OAI-PMH service of Open Images.

What can you use this source for?

- Heritage and history: finding and recording film footage of an event, place or period.
- Education and communication: finding a video that may be freely reused, with its licence and attribution.
- Journalism: checking when and by whom historical footage was made.
- Legal: recording the licence of a video before it is used.

Searching

Under **Sources > Explore & add sources > Public sources**, choose the source **Sound & Vision (Open Images)**. The fields:

- **Search words or link**: one or more words, for example *Afsluitdijk*; each word must occur in the title, description, keywords, maker or broadcaster. You can also paste a link from openbeelden.nl (/media/47540) or the number of the video; that video then appears as the only result.

- **Collection**: the collection on Open Images, for example that of Sound & Vision. Without search words, TrevantAI shows the videos of the collection.

Open Images has no search service. TrevantAI therefore browses the collection and checks the words itself; this may take a moment. Choosing a collection is therefore recommended. You get 20 results at a time; use **More results** to browse further. Tick what you want (at most 20 at a time) and choose **Add to sources**.

What you get

Each video becomes one source with the sections:

- **Description**: title, date, maker(s), broadcaster or publisher, type, duration, language, place, keywords, the licence of the video, the attribution, the link to the video and the summary;
- **Content**: the detailed description, if there is one;
- **Source**: the attribution and the licence.

Each section starts with a breadcrumb (*title › Content*). If the texts are in several languages, TrevantAI chooses the language of your screen, otherwise Dutch. The video itself is not retrieved; the link is in the Description. The item is a **snapshot** with the date of retrieval.

Following

With *Follow…* you create a subscription to new and changed videos: at set times TrevantAI asks Open Images what has been added or updated since the previous time (via OAI-PMH). Filters: *Collection* and *Words* (in the title and the summary). A video that disappears from Open Images is reported. See Following sources with a subscription.

Key

No key is needed. The administrator switches the connection on under **Administration > Source connections**.

Privacy

The query goes to Open Images; search words are not sent, because TrevantAI checks them itself. Do not enter names of clients or other confidential information. Descriptions of newsreels and programmes may concern people who are still alive. Therefore items from this source are processed by the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**. Questions about the items always go to the local AI.

Licence and attribution

The descriptions (metadata) of Open Images are free to use (CC0). Each video has its own licence, usually a Creative Commons licence such as CC BY-SA, or the public domain. That licence and the requested attribution are given with each item. If you want to use the video, follow that licence. With the document, TrevantAI keeps the attribution **Bron: Open Beelden, Beeld & Geluid**.

Good to know

- Open Images is a selection: most programmes held by Sound & Vision are not on it.
- How detailed a description is varies per video.
- Already added? Then **Already in the sources** is shown, with **View**.

- The number of the video on Open Images makes it possible to trace which video was used.
- On 3 and 4 November 2026, Open Images could not be reached. If TrevantAI reports that the source cannot be reached, try again later.

🕰️ Wayback Machine (web archive)

Old versions of web pages from the web archive of the Internet Archive: search by address and period, add the text of a capture or follow a page for new versions. The pages belong to their makers.

> Key: none · Type of use: Description free, text protected by copyright (Copyright held by the original source (©); not under CC BY 4.0, reuse subject to that source's conditions) · Personal data: may contain · Where: at the source

The **Wayback Machine** is the web archive of the **Internet Archive**, an American non-profit organisation. Since 1996 it has kept captures of web pages, including pages that no longer exist. TrevantAI uses the official **CDX server** (the list of captures) and retrieves the capture itself without the archive's toolbar.

What can you use this source for?

- Law and supervision: record what was on a website at a given moment.
- Research and journalism: find a page that has disappeared or changed.
- Policy and communication: see how the text of a page changed over time.

Searching

- **Address** (required): for example *https://www.rijksoverheid.nl/* or *rijksoverheid.nl*. A link from web.archive.org also works.
- **Scope**: *Exactly this address* (default), *Everything below this address* or *The whole website (with subdomains)*.
- **From** and **Up to**: a period, for example a year.
- **Versions**: empty gives only the captures in which the page changed; *Every capture* gives them all.

Only captures of web pages, plain text and PDFs with status 200 count; the oldest are at the top.
You get 20 results at a time; use **More results** to continue.

What you get

The item is the **text of one capture**: a web page or a PDF as it was at that moment. The **Description** gives the original address, the time of the capture (UTC), the status code, the type and the link to the archive. A very large capture is cut off after 300,000 characters; the Description then says so. The **Source** is at the bottom.

Following

With *Follow…* a subscription to an address is created: at set times TrevantAI asks the archive for new captures in which the page changed. Filters: **Address** (always required), **Scope** and **Words** (TrevantAI checks them in the text of the capture). A new capture of an unchanged page does not come in. See Following sources.

Key

No key is needed. The administrator turns the connection on under **Administration > Source connections**.

Privacy

The address goes to the Internet Archive in the United States. A web page can contain anything, including personal data.
Items from this source therefore go to the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

The pages belong to their **makers** and fall under the copyright of that source; reuse according to that source's terms. The **terms of the Internet Archive** (archive.org/about/terms.php): access for study and research, use within copyright, and not disrupting the archive. With the document, TrevantAI stores the attribution "Bron: Internet Archive, Wayback Machine".

Good to know

- Not every page has been archived, and not every capture is complete: images, layout and pages that only work with JavaScript are sometimes missing.
- With a very long address the identifier is shortened, followed by a check code; the full address is in the Description. An address with a space is not recognised; a space is written as %20.
- TrevantAI waits one second between two requests to the archive, so as not to burden it.
- Checked against the real source on 4 November 2026: searching and retrieving work as described here.
- Retrieving the first item can take a few seconds.
- Under **View** you see the source, the identifier, the link, the retrieval date and the search (**Found with**).

News and media

Source connection Feeds (RSS, Atom, JSON Feed)

Follow news and publications from any website with a feed: paste the address of the feed or of the website. Items from a feed go to the local AI only by default.

> Key: none · Type of use: Description free, text protected by copyright (Copyright of the publisher of the feed; reuse according to the terms of the website) · Personal data: may contain · Where: at the source

What a feed is

Many websites also publish their latest posts in a *feed*: a list that programs can read (RSS, Atom or JSON Feed). News sites, governments, municipalities, blogs and magazines often have one. With the *Feeds* connection, TrevantAI reads such a list.

Adding a feed

Paste the address of the feed itself, or simply of the website, under *Address of the feed or website* and choose *Search*. If the page names a feed, TrevantAI reads it; if it names several, a drop-down list *Other feed on this page* appears. The latest items are shown with title, date and summary. Tick what you want to keep and choose *Add*.

What becomes a source

For each chosen item, TrevantAI fetches the page it points to (also if that is a PDF), as with a web page. If that fails, or the page has hardly any text (for example behind a paywall), the text from the feed itself is used. The source shows where the text comes from. The same article from two feeds does not become a duplicate source.

Podcasts and YouTube channels

A podcast also has a feed: paste the address of the podcast feed, or choose *Follow…* for a podcast in the *Podcasts & Streaming* tab. An episode then does not become a text source but an episode with the audio, just as with *Import & Transcribe*: with the creator's transcript if the feed names one, and otherwise transcribed by speech recognition if the administrator has switched it on. A YouTube channel or playlist works the same way: paste the address of the channel (youtube.com/@name), of a playlist or of a video from the channel. A new video comes in with the subtitles, like a YouTube video added by its address. With *Only the text from the feed*, only the description comes in, without audio. An episode or video that is already among the sources does not become a second source.

Following

Follow… creates a subscription to the feed: TrevantAI reads the feed at set times (every hour to weekly, daily by default) and only asks for the feed if something has changed, so a feed without news costs almost nothing. Filters: *Category* (the categories the feed itself names), *Words* (in title and summary) and *Only the text from the feed* (faster and less storage; the page is then not fetched). See Following sources with a subscription.

Privacy and AI

A feed can come from anyone and contain anything. Items from a feed therefore go to the local AI only by default; the administrator can change this for Feeds. The address goes to the website itself; do not enter confidential data. TrevantAI does not fetch addresses within the own network.

Licence

The content of a feed belongs to the publisher. Use items for your own research and observe the terms of the website when sharing.

Status

The status (working, problem, the website asks to wait) is shown per website, as with web sources that TrevantAI checks.

🌐 GDELT (worldwide news)

News articles from all over the world from the last three months, via the GDELT Project: search by words, language, country and website, add the article or follow a topic. The articles are protected by copyright.

> Key: none · Type of use: Description free, text protected by copyright (Copyright held by the original source (©); not under CC BY 4.0, reuse subject to that source's conditions) · Personal data: may contain · Where: at the source

The **GDELT Project** (Global Database of Events, Language and Tone) monitors news sites all over the world in more than a hundred languages and records which articles it sees. TrevantAI uses GDELT's official **DOC 2.0 API**, which searches the articles of the **last three months**.

What can you use this source for?

- Journalism and communication: see how a topic appears in the news at home and abroad.
- Policy and research: collect coverage of an organisation, law or event and keep it traceable.
- Business: follow news about customers, suppliers or your own sector.

Searching

- **Search words** (required): for example *nitrogen farmers*. Put an exact phrase *"in quotation marks"*; *(a OR b)* gives either; *-word* excludes a word. GDELT refuses very short or very common words.
- **Language**, **Country of the source** and **Domain (website)**: for example Dutch, Belgium or *nos.nl*.
- **Period**: the last 24 hours, week, month or three months. Empty: three months.

GDELT returns at most 250 articles per search, newest first.
You get 20 results at a time; use **More results** to continue.

What you get

The item is the **article itself**: TrevantAI retrieves the page at the publisher. The **Description** gives the title, link, domain, when GDELT saw the article, and language and country according to GDELT. If the page cannot be read (a paywall, a page that only works with JavaScript, or gone), the item contains only the Description, with the reason. A very long page is cut off after 200,000 characters. The **Source** is at the bottom.

Following

With *Follow…* a subscription is created: at set times TrevantAI asks GDELT which articles have appeared since the previous time. Filters: **Words** (required; GDELT searches them itself), **Language**, **Country of the source** and **Domain (website)**. New items appear in *Newly arrived*; GDELT looks back at most three months. See Following sources.

Key

No key is needed. The administrator turns the connection on under **Administration > Source connections**.

Privacy

The search words and filters go to GDELT in the United States; TrevantAI retrieves the page of an article from the website itself. Do not use confidential data in the search. News is often about people.
Items from this source therefore go to the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

GDELT's data are free to use, also commercially, with attribution to and a link to the GDELT Project. The **articles themselves** belong to their publishers and are protected by copyright: use them for your own research and observe the website's terms when sharing. With the document, TrevantAI stores the attribution "Gevonden via het GDELT Project (www.gdeltproject.org)".

Good to know

- GDELT allows at most one request every five seconds. TrevantAI therefore waits between two requests itself; a search may take a moment.
- If you change the words of a subscription, the new words apply from the last round onwards; earlier articles do not come in again.
- GDELT records for each article when it saw it, often per quarter of an hour. If more than 250 articles share one time, TrevantAI asks for that time again per chosen language or country (with two or more languages or countries). If that is not possible, articles from that time may be missing; more specific words help in that case.
- Language and country are what GDELT recognises; this is not always correct.
- On 3 and 4 November 2026, GDELT only returned *too many requests* (429) from the test network. This happens on networks that share one internet address with many others; then try again later or from another network.
- Under **View** you see the source, the identifier, the link, the retrieval date and the search (**Found with**).

🎬 Video sites (Vimeo, Dailymotion, TED)

Individual public videos from Vimeo, Dailymotion and TED by link, without an account: the title, the creator, the date, the duration, the description and the spoken text from the subtitles. TrevantAI does not store the video itself. This requires yt-dlp on the TrevantAI computer.

> Key: none · Type of use: Description free, text protected by copyright (Copyright held by the original source (©); not under CC BY 4.0, reuse subject to that source's conditions) · Personal data: no · Where: at the source

With this source, an **individual public video** from **Vimeo**, **Dailymotion** or **TED** becomes a source in TrevantAI. An item is **one video**. **No account and no key** are needed; TrevantAI retrieves the details and the subtitles with *yt-dlp*, a tool on the TrevantAI computer (the same one as for YouTube videos).

What is this source used for?

- Keeping a lecture, a TED talk or a report in a traceable way, together with what is said in it, to search it and chat with it.

Adding a video

- Paste the **link to one video** into the field *Link to a video*, for example *https://vimeo.com/76979871*, *https://www.dailymotion.com/video/…*, *https://dai.ly/…* or *https://www.ted.com/talks/…*. A non-public Vimeo link (with the extra code after the number) also works, as long as you have that link.
- TrevantAI retrieves the details of that video; with **Add** it becomes a source.
- A channel, a playlist or a video from another site is not possible. A **YouTube video** is added as a web page, see YouTube videos as a source; a **TikTok video** with the source TikTok.

What you get

The item contains the **Description** (the title, the creator, the site, the date, the duration, the link and where the spoken text comes from), the **Summary** the creator wrote for the video, the **Spoken text** and, at the bottom, the **Source**. TrevantAI does not store the **video itself**; watch the video on the site via the link. TrevantAI does not include counters (views, likes).

The spoken text

Subtitles by the creator take precedence over automatic ones; among them first Dutch, then English, then the language of the video. For automatic subtitles TrevantAI chooses the original spoken language (an automatic translation is less good); names and technical terms may then be wrong. If a video has no subtitles, only the title and the description are searchable. The spoken text is shown in blocks with the time at which they start, such as *[1:23]*.

If yt-dlp is missing

Without yt-dlp on the TrevantAI computer this source does not work; TrevantAI then says so when searching. How to install yt-dlp is explained in the guide *Installing TrevantAI* in the Help centre. If TrevantAI says that yt-dlp could not read a video, updating yt-dlp usually helps.

Licence and attribution

The videos, the description and the spoken text belong to their creators and are protected by their copyright; use them for your own research and credit the creator. **TED talks** are licensed under **CC BY-NC-ND 4.0**: non-commercial only, with attribution and without modification. If the administrator has set *Use of this installation* to commercial, TrevantAI does not add a TED talk.

Good to know

- For this source TrevantAI does not (yet) check whether a video has been deleted on the site; the item remains until you remove it yourself.

- Under **View** you see the source, the identifier (site:number), the link, the date of retrieval and the search query (**Found with**).

Social media

Social media and marketplaces as sources

Three ways

1. **Importing your own export**: you make an export in the platform's app and add it to TrevantAI. Now: WhatsApp and Instagram, Facebook and Threads (your own data from Meta). Also: Telegram (the export from Telegram Desktop).
2. **A single public post by link**: TrevantAI fetches that one post and checks whether it is still there. Now: X, TikTok and other video sites (Vimeo, Dailymotion, TED; for videos the spoken text from the subtitles, never the video itself). And a single marketplace listing as a web source, with a summary card and periodic checks (eBay too).
3. **Searching and following via an official connection**: now Bluesky, Mastodon and X; for eBay no connection (see Marketplaces).

Where to find them

In *Explore & add sources*, Bluesky, Mastodon, X and TikTok are listed under **Social media** (with a reference under *News and media*); the other video sites are listed under **News and media**; under **Marketplaces** there are two information tiles, *Marktplaats, 2dehands and others* and *eBay*. In the source overview, *Type* has the group *Social media*, with a separate type per platform (*WhatsApp chat* ...).

Rules

Each platform has an (i) with what its terms say and a link to them. TrevantAI does not put up fences but points to the rules; what is prohibited by law, TrevantAI does not build (such as systematically copying search results from a marketplace). Whatever contains data of others, such as a chat, is local only by default.

Marketplace listings as a source

A **listing** on a marketplace can become a **web source** in TrevantAI: for example a car sold in a dispute, or an offer a regulator wants to follow. TrevantAI recognises the well-known marketplaces and adds a **summary card**; with **Check automatically for changes** you see what changes.

Why no search function?

- **Marktplaats and its foreign counterparts** (2dehands, Kleinanzeigen, leboncoin, Wallapop, Subito, OLX, Vinted, Etsy, willhaben, Facebook Marketplace) prohibit the systematic copying of listings in their terms, and their collections are protected by database rights. TrevantAI therefore builds no search or follow function.
- **eBay** does have an official connection (the API), but it does not allow listings to be stored or kept once they are gone. That is exactly what matters in TrevantAI. Hence no eBay connection: eBay works here just like Marktplaats.

Capturing a listing

1. Copy the address of the listing in the browser.
2. In **Sources** choose the tab **Web & Articles** > **Enter direct URL** and paste the address.

3. Tick **Check automatically for changes** if TrevantAI should check the listing periodically.

For a recognised listing the screen already says so: *A listing on …*. **Explore underlying links** is then off: the links on a listing page are other listings.

The summary card

A recognised listing opens in the **Listing** tab: the link to the site, a summary card with **title, price, status, location, seller** and the **posting date**, the **features** and the seller's **description**, one line about the check and the screenshot of the page. Everything else (description field, labels, AI use, *Check automatically*, the text, the links and the versions) is under **Source information**. The data come from what the page itself provides for search engines and, for Marktplaats and 2dehands, from fixed places on the page. TrevantAI does not piece anything together from the running text: if the page provides no data, the card says so and the check compares only the text. Every version is kept, with the snapshot (archive and screenshot).

The text as a source

For searching, chatting and comparing TrevantAI only uses the listing itself: title, price, location, features and description. The menus, *popular search terms* and the surrounding listings do not count, so they cause no false changes. The full page text is collapsed under **Source information > Text**; the snapshot is the evidence of what the whole page looked like. A listing that already existed gets a new version with the cleaned text once, at the next check, without a notification.

What the check reports

In the **Notifications** bell (and by e-mail if set up), in plain words:

- *Price changed: €250 → €200*;
- *Status: reserved*, *Status: sold*, *No longer available* or *Listing no longer found*;
- *Title changed* or *Description changed*.

If only numbers change, such as the view count or *3 hours ago*, no notification is sent (a new version is kept). **View differences** shows exactly what is different.

How often, and when does it stop?

- By default **daily**; not more often than **every 6 hours**, and for **eBay** at most **once a day** (eBay's terms prohibit automated access; TrevantAI only retrieves this one page, responsibility lies with the user).
- The check **stops automatically** if the listing is **gone three times in a row**, or after **90 days without change**. A notification follows; **Resume** continues the check.

Search results pages

A page of search results (for example *all folding bikes*) can be captured **once**, but **not periodically** and **without underlying pages**: that would be the systematic copying the terms prohibit. For such a page *Check automatically* is therefore off, with the explanation next to it.

If the marketplace refuses

Some marketplaces block programs, and **Facebook Marketplace** always requires signing in. If the site refuses the page, TrevantAI says so. Save the page in the browser as a **PDF** (*Print* > *Save as PDF*) and add the file; periodic checking is then not possible, and there is no summary card. TrevantAI does not try to get

around a block.

Privacy

A listing usually names the seller. Therefore by default only the **local AI** may process a recognised listing; this can be changed per source under AI use.

Good to know

- Recognised: Marktplaats, 2dehands/2ememain, Kleinanzeigen, leboncoin, Wallapop, Subito, OLX, Vinted, Facebook Marketplace, eBay (all country sites), Etsy and willhaben.
- See also Check automatically, Notify on change and Social media and marketplaces as a source.

🦋 Bluesky (social media)

Public posts from Bluesky: search by words, author, language and period, add a post or follow a topic or account. Deleted posts, or posts that are no longer public, lose their text.

> Key: optional · Type of use: Description free, text protected by copyright (Copyright held by the original source (©); not under CC BY 4.0, reuse subject to that source's conditions) · Personal data: may contain · Where: at the source

Bluesky is a social network on the open **AT Protocol**. TrevantAI searches and retrieves public posts through Bluesky's official API. An item is **one post**.

What is this source used for?

- Journalism and communication: seeing what is said about a topic, organisation or person, and keeping a statement traceable.
- Policy and research: collecting reactions to a decision, law or event.
- Following an account, for example of a ministry, regulator or newsroom.

Searching

- **Search words or link** (required): for example *nitrogen farmers*; an exact phrase goes "*in quotation marks*". The link of a post (*https://bsky.app/profile/.../post/...*) or an AT URI (*at://...*) returns exactly that post.
- **Author (handle)**: for example *bsky.app* or *name.bsky.social*, with or without @.
- **Language**, **From** and **Up to and including**, and the **Order**: newest or most popular first.

Results come 25 at a time; **More results** shows the next ones.

Without or with an account

Searching, retrieving and following work without an account. Bluesky's public service sometimes refuses searches to limit the load; a message then says that searching without an account was refused. The administrator can then enter, under **Administration > Source connections**, the handle of the organisation's own Bluesky account and an **app password**, in one field, separated by a space (see **Bluesky account for searching**). Only searches then go through that account; retrieval and the check for deleted posts always work without it.

What you get

The item contains the **Description** (the author with display name and @handle, the date, the language, the link, what the post replies to or quotes, as a link only, a link card and the links in the text), the **Post** itself, the **description of images and video** (the alt texts the author provided) and, at the bottom, the **Source**. TrevantAI does not include counters (likes, reposts) and does not fetch pages behind links. A post that is replied to or quoted can be added as a separate item if needed.

Not meant for everyone

An author can indicate that people without an account should not see their posts (the label *!no-unauthenticated*). TrevantAI does not show such posts in the results and does not add them.

Deleted at the source

Every 12 hours TrevantAI checks whether the stored posts still exist, even if the connection has since been switched off. If a post has been deleted or is no longer public (the account was deactivated, suspended or deleted, or the author set the label), its text disappears from TrevantAI. The item remains as a card "Post by @name" with the link and the date, so that earlier answers stay traceable. Verbatim quotations in answers and fields are replaced; the project owner receives a notification. TrevantAI cannot recall an export that has already been made.

Following

Follow… creates a subscription to **Words** and/or an **Author**, optionally with **Language**. With words, TrevantAI asks Bluesky at set times which posts have been added since the previous time; with only an author it follows that account's posts (without reposts). New items appear in *New arrivals*. See Following sources.

Privacy

The search words and filters are sent to Bluesky in the United States. Do not use confidential information in the query. Posts are written by people; items from this source are therefore processed by the local AI only by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

The posts belong to their authors and are protected by their copyright (Bluesky Terms of Service). Use them for your own research and name the author; TrevantAI stores the attribution with the author's name and @handle with each item.

Good to know

- Following by words looks at the moment Bluesky received a post; a post that appears later with an old date (for example an imported post) is not picked up.
- On 4 November 2026 Bluesky refused searching without an account; adding a post by its link does work. Searching requires an app password in Administration (see above).
- **View** shows the source, the identifier (the AT URI), the link, the retrieval date and the search query (**Found with**).

Mastodon (social media)

Public posts from Mastodon servers: search by hashtag or account, add a post by link, follow a hashtag or account. Deleted or no longer public posts lose their text; people who do not want to be found are not shown.

> Key: optional · Type of use: Description free, text protected by copyright (Copyright held by the original source (©); not under CC BY 4.0, reuse subject to that source's conditions) · Personal data: may contain · Where: at the source

Mastodon is a social network made up of many independent **servers** (together the *fediverse*): for example *mastodon.social*, *fosstodon.org* or a server run by a government body or university. Each server knows the posts of its own users and the posts that reach it from other servers. TrevantAI therefore always asks one server, through the official Mastodon API. An item is **one post**.

What is this source used for?

- Journalism and research: collecting public posts on a topic (a hashtag) and keeping them traceable.
- Government and policy: following reactions to a decision, or the account of an organisation.
- Science: discussions in the own research community, which are often on Mastodon.

Searching

- **Server**: the server TrevantAI searches. The administrator chooses the default server (pre-filled *mastodon.social*) and whether another server may be chosen. If so, the default server is filled in here and TrevantAI remembers the server used last. Enter only the name, such as *fosstodon.org*.
- **Hashtag**: for example *climate* or *#climate*. TrevantAI shows the public posts with that hashtag that this server knows; another server may know other posts.
- **Account**: for example *@Mastodon@mastodon.social*, *name@server* or the link of the profile. Together with a hashtag: only that account's posts with that hashtag.
- **Search words or link**: the link of a post (such as *https://mastodon.social/@name/113…*) gives exactly that post, from any server. Searching by words only works with an access token on the default server (see below).
- **Language**: only posts in that language.

Results come 20 at a time; **More results** shows the next ones.

Without or with an access token

Searching by hashtag and account, adding a post by link and following work without an account. Mastodon searches by **words** only for signed-in users, and then only in posts of people who want to be searchable. For this, the administrator can enter an access token of an own account on the default server (see **Mastodon: default server and access token**). Some servers also show hashtags or accounts only to signed-in users; TrevantAI then says: *This server shows this only to signed-in users. Choose another server.*

People who do not want to be found

A Mastodon user can indicate in the profile not wanting to be indexed by search engines, or not wanting to be discoverable in the discovery features. TrevantAI respects this: such posts do not appear in the search results and such an account cannot be followed. Below the results, TrevantAI says how many posts were not shown for this reason. A pasted link to a single post does work (a deliberate choice); TrevantAI then reports *This user has indicated not wanting to be indexed* and also notes this in the source.

Which posts

In the results and when following, only **public** posts. A post that is public but not shown in the timelines can only be added by link. Posts for followers and direct messages never. Replies are included, passed-on posts (*boosts*) are not: add the original post instead.

What is retrieved

The item contains the **Description** (a content warning at the top if there is one, the author with display name and @name@server, the date, whether the post was edited, the language, the link, the server it was read from, what it replies to, a link card, the hashtags and the options of a poll), the **Post** itself, the **description of images and video** (the alt texts given by the author) and at the bottom the **Source**. Counters (favourites, boosts) are not included.

Deleted at the source

Every 12 hours TrevantAI checks with the server where the post was found whether it still exists, also when the connection has been switched off in the meantime. If a post was deleted or is no longer public, or the account was suspended, the text disappears from TrevantAI. The item remains as a card "Post by @name@server" with the link and the date. Verbatim quotes in answers and fields are replaced; the project owner receives a notification. If a server does not respond for a while, TrevantAI waits and deletes nothing.

Following

Follow… creates a subscription to a **Hashtag** and/or an **Account** on a server, optionally with a **Language**; with an access token also to **words**. At set times TrevantAI asks the server which posts have been added since the last time. What comes in appears in *Newly received*. See Following sources. Prefer an RSS feed? Every Mastodon server offers one per account (*https://server/@name.rss*) and per hashtag (*https://server/tags/name.rss*); these can be followed as a feed, but then without the check for deleted posts and without respecting people who do not want to be found.

Privacy

The hashtag, the account and the search words are sent to the chosen Mastodon server. Do not use confidential data in the search. Posts are from people; that is why items from this source go only to the local AI by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

The posts belong to their authors and are protected by their copyright; each server has its own terms (for example *mastodon.social/terms*). Use them for own research and credit the author; TrevantAI keeps the attribution with the name and @name@server with every item.

Good to know

- The same post, found via two servers, becomes two items (a post has its own number on each server).
- A link to a post on other fediverse software (such as Pixelfed) only works if that server supports the Mastodon API.
- Under **View** the source, the identifier (server/number), the link, the retrieval date and the search (**Found with**) are shown.

𝕏 X (formerly Twitter)

Individual public posts from X, added by their link, without an account, key or costs. Searching and following are not possible with this connection; deleted or no longer public posts lose their text.

> Key: none · Type of use: Description free, text protected by copyright (Copyright held by the original source (©); not under CC BY 4.0, reuse subject to that source's conditions) · Personal data: may contain · Where: at the source

X (formerly Twitter) is a social network. TrevantAI retrieves individual public posts through the public **embed service of X** (oEmbed): the same service websites use to display a post. This requires **no account, no key and no costs**. An item is **one post**.

Adding a post

- Paste the **link to a post** into the field *Link to a post*, for example *https://x.com/BarackObama/status/266031293945503744*. A twitter.com link or just the number of the post also works.
- TrevantAI retrieves that one post; with **Add** it becomes a source.

What this connection cannot do

- **Searching** by words, author or language: the embed service only knows individual posts.
- **Following** a topic or an account.
- **Images and video**: TrevantAI only stores the text, with the links contained in the post.

What you get

At the top of the item is a **header** with the name and @username of the author (with a link to the profile on X), the date (with a link to the post), the X logo and **View on X**. Below it follow the **Description** (author, date, language, link and the links in the post), the **Post** itself, unchanged, and at the bottom the **Source**. For very old posts (from before late 2010) TrevantAI only knows the day, not the time. The embed service does not say what a post replies to or quotes; TrevantAI does not store counters (likes, reposts).

Deleted at the source

Every 12 hours TrevantAI checks whether the stored posts are still public, even if the connection has been switched off in the meantime. This check is free. If X no longer returns a post (deleted, protected, account suspended or deactivated), the text disappears from TrevantAI; the item remains as a card "Post on X" with the link and the date. Literal quotations in answers and fields are replaced; the project owner receives a notification. An export that has already been made cannot be recalled by TrevantAI.

Terms

The administrator switches X on under **Administration > Source connections** and confirms the terms there: the terms of service of X and the display rules of X (Display Requirements). TrevantAI shows the text unchanged, with the author, the date and a link to X, without a profile picture and without the buttons of X. TrevantAI does not train any AI model with content from X, and only the link or the number of a post is passed on to third parties.

Privacy

When retrieving and checking, only the link of the post goes to X in the United States, without an account. Posts belong to persons; that is why items from this source go only to the local AI by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

The posts belong to their authors. TrevantAI keeps the attribution with the name and @username of the author with each item.

Need searching or following?

Searching X and following are only possible through the paid X API, with your organisation's own developer account at X. If your organisation has such an account, it can ask TrevantAI to investigate **for a fee** what more is possible.

Good to know

- Checked against the real source on 4 November 2026: retrieving, the check for deleted posts and a post that does not exist work as described here.
- TrevantAI requests at most one post per second from X.
- Under **View** you see the source, the identifier (the number of the post), the link, the retrieval date and the search (**Found with**).

🎵 TikTok (video by link)

Individual public TikTok videos by link, without an account: the caption, the creator, the date, the sound, the hashtags and the spoken text from the subtitles. TrevantAI does not store the video itself; deleted or private videos lose their text.

> Key: none · Type of use: Description free, text protected by copyright (Copyright held by the original source (©); not under CC BY 4.0, reuse subject to that source's conditions) · Personal data: may contain · Where: at the source

TikTok is a social network for short videos. TrevantAI retrieves individual public videos through TikTok's public **embed service** (oEmbed): the same service websites use to show a video. This requires **no account and no key**. An item is **one video**.

What is this source used for?

- Journalism and research: keeping a video that is in the news in a traceable way, together with what is said in it.
- Government and policy: checking what is circulating about a decision or an organisation.

Adding a video

- Paste the **link to a video** into the field *Link to a video*, for example *https://www.tiktok.com/@scout2015/video/6718335390845095173*. A short link (*vm.tiktok.com/…*, as the Share button in the app gives it) also works.
- TrevantAI retrieves that one video; with **Add** it becomes a source.

- A profile, a hashtag or a photo post is not possible. Searching TikTok and following are not possible either: TikTok has no open service for that.

What you get

The item contains the **Description** (the creator with name and @name, the time of posting, the link, the sound, the hashtags and where the spoken text comes from), the **Post** (the creator's caption), the **Spoken text** and, at the bottom, the **Source**. TrevantAI does not store the **video itself**, nor a thumbnail; watch the video on TikTok via the link. TrevantAI does not include counters (likes, views).

The spoken text

TrevantAI retrieves the video's **subtitles** with *yt-dlp* on the TrevantAI computer (the same tool as for YouTube videos). Subtitles by the creator take precedence over automatic ones; in automatic subtitles, names and technical terms may be wrong. Not every TikTok video has subtitles. The Description says what happened: *from the creator's subtitles*, *from TikTok's automatic subtitles*, *this video has no subtitles*, or *not retrieved (yt-dlp is not installed on the TrevantAI computer)*. The spoken text is shown in blocks with the time at which they start, such as *[0:12]*.

Deleted at the source

Every 12 hours TrevantAI checks with TikTok whether the video is still public, even if the connection has been switched off in the meantime. If the video has been deleted or made private, or the account no longer exists, the text disappears from TrevantAI. The item remains as a card "Video by @name on TikTok" with the link and the date. Literal quotes (from the caption and from the spoken text) in answers and fields are replaced; the project owner receives a notification. If TikTok does not respond for a while, TrevantAI waits and deletes nothing.

Taking it slowly

TikTok quickly refuses when many requests come from one address. TrevantAI therefore requests one video at a time, with two seconds in between. If TikTok refuses for a moment anyway, the source shows *asks to wait a moment*; TrevantAI tries again later.

Privacy

The link of the video is sent to TikTok. Videos are made by people; that is why items from this source go only to the local AI by default. The administrator can change this under **Who may process items from this source?** in **Administration > Source connections**.

Licence and attribution

The caption and the spoken text belong to the creator and are protected by the creator's copyright; TikTok's terms of service also apply, which the administrator confirms when switching the source on. Use them for your own research and credit the creator; TrevantAI stores the attribution with the name and @name for every item.

Good to know

- If a short link does not work (TikTok does not always redirect), open it in the browser and paste the full link (*tiktok.com/@name/video/…*).

- Other video sites (Vimeo, Dailymotion, TED) are in the source Video sites; YouTube videos are added as a web page, see YouTube videos as a source.
- Under **View** you see the source, the identifier (name/number of the video), the link, the date of retrieval and the search query (**Found with**).

Religion and theology

Theological sources: overview

TrevantAI has seven sources for Bible study, theology and exegesis. Five of them are stored **on the TrevantAI computer itself**: searching and adding work without internet, and no search query leaves the computer. *Wikisource* and *Sefaria* search live on the internet.

Source	Contents
Bibles	19 free translations in 10 languages, from the Dutch Statenvertaling to the Vulgate and the Hebrew and Greek text
Marginal notes	the almost 60,000 marginal notes to the Statenvertaling (1637)
Original text and lexicon	Hebrew OT and Greek NT word by word, with Strong's number, lemma, grammar and lexicon
Cross-references, places and people	related verses, biblical people and places with their likely location
Creeds and catechisms	24 historical creeds, confessions and catechisms (English), from Nicaea to Westminster, partly with proof texts
Wikisource	Church Fathers, Reformers and confessions in their original language (such as the Heidelberg Catechism), live
Sefaria	Jewish texts: Tanakh, Mishnah, Talmud and commentaries (such as Rashi), Hebrew or Aramaic with a translation alongside, live

How it works

1. Choose the source on the *Public sources* tab.
2. Type a Bible reference (*John 3:16*, *Ps 23; 121*, *Gen 1:1-2:3*), words, a name or a Strong's number. References may be written in seven languages.
3. Tick the results and choose **Add**: they become sources of the project.
4. Ask questions about them in the chat: the answer names the items used. See Asking questions about theological sources.

Stronger together

- *Original text* places the Statenvertaling alongside the Hebrew or Greek.
- *Cross-references* uses the Bible text from *Bibles* and the Dutch names from *Original text*.
- Not all Bibles number verses the same way; TrevantAI converts. See Verse numbering.
- If what you are looking for is (also) in another source, for example a name in *Cross-references* or a confession in *Creeds and catechisms* or *Wikisource*, the results show a tip with a button that searches there straight away.
- If you search *Bibles*, *Original text* or *Marginal notes* for an Old Testament passage and *Sefaria* is switched on, a tip points to the same passage in Sefaria, with Jewish translations and commentaries.

Eight ways of working

- *Sermon or Bible study*: the passage with a second translation, the marginal notes and the cross-references; ask the chat for the main line and discussion questions.

- *A key word*: the original text of the verse and the lexicon item of the word (meanings, how the Statenvertaling renders it, occurrences).
- *A person or place*: the item from *Cross-references, places and people*, with the main passages.
- *Comparing translations*: the same passage in several translations, or with *Alongside*, plus the original text.
- *Textual criticism in outline*: in the Greek original text, *only in: TR/KJV* marks words that the Statenvertaling follows and modern translations do not.
- *A theme*: search for words in Bibles or Marginal notes, and add the cross-references of the strongest hits.
- *A doctrine*: search *Creeds and catechisms* for a Bible reference or a word, add the articles or questions with the proof texts written out, and ask the chat how the confessions put it; the original text can be found through *Wikisource*.
- *Jewish interpretation*: search *Sefaria* for the passage, add the Hebrew text with a translation, and through *Related texts* the commentaries (such as Rashi); ask the chat how they read the verse.

The administrator installs the sources under Administration > Source connections.

Verse numbering: why a verse sometimes has another number

Not all Bibles number verses the same way. The best-known differences:

Passage	Statenvertaling / Hebrew	English (KJV) numbering
Psalm titles	the title is verse 1 (*Psalm 3:1: A Psalm of David*)	the title is verse 0
Malachi 4	Hebrew: Malachi 3:19-24	Malachi 4:1-6 (the Statenvertaling follows the English numbering here)
Joel 3	Hebrew: Joel 3:1-5	Joel 2:28-32

TrevantAI converts using the STEPBible table (TVTMS):

- **Bibles** reads a reference in the numbering of the chosen translation. With *Alongside*, the matching verse always appears next to it, even when the number differs.
- **Original text and lexicon** reads the OT in the Hebrew numbering. An English number is converted, and the title says so: *Malachi 3:19 (Hebrew) = Malachi 4:1 in the standard numbering*.
- **Cross-references, places and people** reads in the numbering of the chosen translation; the referenced verses use that numbering too.

Every item states which numbering it uses. *(standaard 3:0)* after a reference is the number in the English numbering.

Asking questions about theological sources

The chat works with the sources of the project. So first add what is to be discussed; a passage, the lexicon item of a key word and the marginal notes together give the richest answer.

Example questions

- *Compare Romans 3:28 in the Statenvertaling, the Luther 1912 and the Vulgate. Where do they differ?* (Bibles)
- *What do the marginal notes say about "image and likeness" in Genesis 1:26?* (Marginal notes)

- *Which verb is used in Genesis 1:1 and how is it used elsewhere in the OT?* (with the lexicon item of H1254; Original text and lexicon)
- *Where does the Statenvertaling follow a different Greek text in John 3 than the critical editions?*
- *Which verses are most strongly linked to Isaiah 53:5, and what do they have in common?* (Cross-references, places and people)
- *Give an overview of the places Abraham visited, with their likely location.*

Tips

- **Ask for quotations with references.** Every line in an item starts with the reference, so the chat can quote precisely.
- **Lexicon and descriptions are in English.** The chat translates on request.
- **Keep checking.** The AI can make mistakes; marginal notes and descriptions are interpretations. The Bible text and the lexicon are the touchstone.
- On numbering: see Verse numbering.

Back to the overview of the theological sources.

📖 Bibles (public domain)

Nineteen free Bible translations on your own computer: search by Bible reference or by words, and add a passage, optionally with a second translation alongside.

> Key: none · Type of use: Free to use (Public domain) · Personal data: no · Where: local (download)

Bibles is a source library of free Bible translations (public domain) from eBible.org: Dutch (Statenvertaling, Petrus Canisius 1939), English (King James Version, World English Bible, Berean Standard Bible, Douay-Rheims 1899), German (Lutherbibel 1912, Elberfelder 1905), French (Louis Segond 1910, Ostervald), Spanish (Reina-Valera 1909), Italian (Diodati 1885, Riveduta 1927), Portuguese (Bíblia Portuguesa Mundial), Latin (Vulgata Clementina), Greek (Septuagint, Robinson-Pierpont, Textus Receptus) and Hebrew (Westminster Leningrad Codex).

The Dutch Statenvertaling (1637) does not come from eBible.org but from the transcription by Isidore Guild (CC0, text of the 1888 edition), with the Statenvertaling's own verse numbering: Psalm 3:1 is the title, Isaiah 8:23 and John 1:52 exist.

Installing

Unlike the other public sources, the translations are stored on the TrevantAI computer itself. When the administrator switches the source on in **Administration > Source connections**, TrevantAI retrieves them (about 90 MB to download, about 390 MB on disk); this runs as a background task. After that, searching and adding also work without internet. With *Update* the administrator retrieves new versions; with *Free up space* the translations are removed from the computer. Passages already added remain.

Searching

- A Bible reference in one of the seven languages, for example *John 3:16*, *Jn 3,16*, *1 Cor 13*, *First Corinthians 13:4-7*, *Ps 23; 121* or *Rev 21:1-4*. The numbers are those of the chosen translation.
- Or search words: every verse containing all the words; put a phrase in quotation marks. Capitals, accents, Greek accents and Hebrew vowel points are ignored.

Two translations side by side

Choose a second translation under *Alongside*. Each verse is then followed by the verse (or verses) of the second translation with the same text. Bibles are not all numbered the same way: in the Hebrew Bible Malachi 4 is chapter 3:19-24, in the Vulgate Psalm 23 is Psalm 22, and a psalm title is sometimes verse 1. TrevantAI converts this using the verse numbering of STEPBible (TVTMS).

What an item contains

One part per chapter; each line starts with the reference and the abbreviation of the translation, for example *John 3:16 (KJV)*, so that the chat can quote precisely. Footnotes, headings and introductions are left out. The origin is included: the translation, the version (date and checksum of the package) and the link to the translation at eBible.org.

Licence and privacy

All translations are in the public domain: free to use, also commercially. The verse numbering of STEPBible.org is under CC BY 4.0. The source contains no personal data.

Not included

Translations under copyright (such as the NBV21 or the HSV) are not included.

See also: Theological sources: overview · Verse numbering · Asking questions

📝 Marginal notes to the Statenvertaling

Almost 60,000 marginal notes by the translators of the Dutch Statenvertaling (1637), to be looked up per verse or searched by words.

> Key: none · Type of use: Free to use (Public domain) · Personal data: no · Where: local (download)

The **marginal notes** (kanttekeningen) are the annotations that the translators of the Dutch Statenvertaling (1637), commissioned by the States General, wrote alongside the text: explanations of words and phrases, alternative renderings ("Anders, …"), background, and references to other Bible passages. There are almost 60,000 of them.

Where they come from

From the same file as the Statenvertaling in the *Bibles* source: the transcription by Isidore Guild (CC0), text of the 1888 edition. They are stored on the TrevantAI computer itself, in the same source library as the Bibles: switching on either Marginal notes or Bibles installs both. If *Bibles* was installed before the Statenvertaling was added, Administration > Source connections shows *Update*.

Searching

- A Bible reference, for example *Gen 1:1*, *John 3*, *Ps 23; 121*: the notes on those verses. The verse numbering is that of the Statenvertaling (Psalm 3:1 is the title).
- Or search words: in the keywords and text of the notes; every verse with a match is a result.

What an item contains

One part per chapter. For each verse, first the text of the Statenvertaling, then the notes, for example: *Kanttekening 1 bij Genesis 1:1 «beginne»: Van den tijd der schepping* …. This way the chat can quote both the Bible text and the note precisely. The references remain in the note as in the source.

Licence and privacy

Public domain (the transcription: CC0): free to use, also commercially. No personal data.

Please note

The notes are a seventeenth-century Dutch text, in the spelling of 1888, and reflect the interpretation of the Dordrecht translators.

See also: Theological sources: overview · Verse numbering · Asking questions

🔤 Original text and lexicon (Hebrew and Greek)

The Hebrew Old and Greek New Testament word by word, with Strong's number, lemma, spelled-out grammar and a lexicon; for each original word, how the Statenvertaling translates it.

> Key: none · Type of use: Free to use (Free to use with attribution (CC BY 4.0)) · Personal data: no · Where: local (download)

With **Original text and lexicon** you work with the text that was translated from: the Hebrew Old Testament (Westminster Leningrad Codex) and the Greek New Testament, word by word.

Where it comes from

From STEPBible.org, based on work by Tyndale House Cambridge (CC BY 4.0): the word files TAHOT and TAGNT, the lexicons TBESH and TBESG and the explanation of the grammar codes. Which Dutch words the Statenvertaling uses for each original word comes from the transcription by Isidore Guild (CC0). Everything is stored on the TrevantAI computer itself (about 135 MB to download, about 85 MB on disk); searching does not go over the internet. The administrator switches the source on under Administration > Source connections; installation then starts automatically.

Searching

- A Bible reference, for example *Gen 1:1* or *John 3:16-18*: the original text of those verses.
- A Strong's number, for example *H1254* or *G26*: the lexicon.
- A Hebrew or Greek word, with or without vowel points and accents, including an inflected form such as *ἠγάπησεν*.
- A Dutch word from the Statenvertaling, for example *schiep*: the original words behind it. An English word searches the short English meanings.

What an item contains

A text item gives, for each verse, the original text, the Statenvertaling (if the *Bibles* source is installed) and for each word: the form with pronunciation, the Strong's number with the lemma, the English gloss and the grammar, spelled out in the language of your screen (for example *verb, Qal, perfect, 3rd person, masculine,*

singular). A lexicon item gives the meaning (English), how the Statenvertaling translates the word, and where it occurs. At most 180 verses per item: every chapter fits.

Numbering and editions

The Old Testament uses the Hebrew numbering, the numbering of the text itself (Psalm 3:1 is the title, as in the Statenvertaling). If you type the English numbering, TrevantAI converts it and says so in the title (for example *Malachi 3:19 (Hebrew) = Malachi 4:1 in the standard numbering*). If a reference exists in both numberings (*Ps 51:3*), the Hebrew one applies. For the Greek, the text of the critical editions (Nestle-Aland/SBL) comes first; words found only in the Textus Receptus (the original text behind the Statenvertaling and the KJV) or other editions are marked.

Licence and privacy

CC BY 4.0: free to use, also commercially, with attribution to STEPBible.org; that attribution is included in every item. No personal data.

See also: Theological sources: overview · Verse numbering · Asking questions

🗺️ Cross-references, places and people

For every Bible verse the related verses, the places with their likely location and the people; and every person or place itself, with names, family and occurrences.

> Key: none · Type of use: Free to use (Free to use with attribution (CC BY 4.0)) · Personal data: no · Where: local (download)

With **Cross-references, places and people** you can see, for a Bible passage, where else the Bible speaks about the same thing, where it happened and who was involved.

Where it comes from

- *Cross-references*: OpenBible.info, over 340,000 links between verses. Users have voted on each link; the more votes, the stronger. For each verse you see the ten strongest.
- *People and places*: STEPBible.org (TIPNR, based on work by Tyndale House Cambridge): over 3,100 people and 1,000 places, with their names in Hebrew or Greek, family, tribe or region and all occurrences.
- *Location*: OpenBible.info Bible Geocoding: for each place the modern locations proposed for it, with coordinates and a confidence of 0 to 1000 drawn from more than seventy reference works.
  Everything is stored on the TrevantAI computer itself (about 20 MB to download, about 20 MB on disk). The Bible text comes from the *Bibles* source; the Dutch names (as the Statenvertaling spells them) from the *Original text and lexicon* source, if those are installed.

Searching

- A Bible reference, for example *Gen 12:8*, *John 3:16* or *Ps 23*: for each verse the text, the strongest cross-references (with their text), the places and the people. Optionally choose a translation; the numbering is that of the translation.
- A name, for example *Abraham*, *Bethel*, *Jacob* or *Jerusalem*: people and places. Dutch, English and other spellings find each other (Betel, Bethel, Beth-El).

A person or place

Overview (description, names in the original text with Strong's number, family, tribe or region), for a place the location (the most likely first, with confidence and coordinates), all occurrences (the first sixty with the Bible text) and a short English description. STEPBible wrote that description with the help of AI; check the facts against the Bible text.

Licence and privacy

CC BY 4.0: free to use, also commercially, with attribution to OpenBible.info and STEPBible.org; that attribution is included in every item. No personal data (only biblical persons).

See also: Theological sources: overview · Verse numbering · Asking questions

📜 Creeds and catechisms

24 historical creeds, confessions and catechisms (English), by article or question and answer, several with the proof texts; search by name, Bible reference or words.

> Key: none · Type of use: Free to use (Public domain) · Personal data: no · Where: local (download)

With **Creeds and catechisms** you add historical creeds, confessions and catechisms as a source, by article, section or question and answer, with the Bible texts they appeal to (the *proof texts*).

Contents

24 documents in English, from the Creeds.json collection:

- *Early church*: the Nicene Creed, the Athanasian Creed and the rules of faith of Tertullian and Gregory Thaumaturgus.
- *Middle Ages*: the Waldensian Confession.
- *Reformation*: Zwingli's Sixty-seven Articles and Fidei Ratio, the Ten Theses of Berne, the First Helvetic Confession, the Consensus Tigurinus, the French Confession and the Canons of Dort.
- *Westminster*: the Westminster Confession, the Larger and Shorter Catechisms, and the expositions by Matthew Henry, John Flavel and Erskine and Fisher.
- *Baptists*: the 1689 Confession, the Baptist Catechism (1695), the catechisms of Keach and Spurgeon and the Abstract of Principles.
  Everything is stored on the TrevantAI computer itself (about 4 MB to download, about 9 MB on disk).

Why not everything

Only documents whose English text is in the public domain. For the Heidelberg Catechism, the Belgic Confession, the Apostles' Creed, the Chalcedonian Definition and a few others, the collection only has modern, copyrighted English translations; they are therefore not included. The original texts (Dutch, German, French, Latin) can be found through the Wikisource source. Anyone searching for such a confession also gets this tip with the search results.

Searching

- A name, for example *Dort*, *Westminster* or *Nicene*: the whole document (a large document in parts).

- A Bible reference, for example *Rom 8:28* or *John 3:16*: the articles and questions that cite it as a proof text. Proof texts are included with the Westminster Confession, the Westminster Larger Catechism and three Baptist catechisms.
- Words, for example *justification* or *election*.
- Optionally choose a *Document* to search only that one, and under *Write out proof texts in* a translation from the Bibles source: each article or question then also shows the text of its proof texts.

An item

First *About this document* (year, place, original language, origin), then one section per chapter or per ten questions. Every question and article keeps its number, so that the chat can refer precisely (*Westminster Shorter Catechism, Question 33*). The numbers in [ ] point to the proof texts below; the references follow the English (KJV) verse numbering, see Verse numbering.

Licence and privacy

Public domain: free to use, also commercially. Every item states where the text comes from. No personal data.

Sefaria (Jewish texts)

The free online library of Jewish texts: Tanakh, Mishnah, Talmud, commentaries and more, with the Hebrew or Aramaic text and a translation side by side; licence per version.

> Key: none · Type of use: Free, with attribution and share-alike (Licence per text version: public domain, CC0, CC BY or CC BY-SA; versions under CC BY-NC for non-commercial use only (labelled)) · Personal data: no · Where: at the source

Sefaria is a free online library of Jewish texts. It includes the Tanakh (the Hebrew Bible), the Mishnah, the Talmud, Midrash, Halakhah, Kabbalah, liturgy and the classical commentaries (such as Rashi), in the original language and in translation. The connection searches Sefaria live.

Searching

- A **reference**: *Genesis 1:1-5*, *Psalms 23*, *Berakhot 2a* (a folio of the Talmud), *Mishnah Avot 1*, *Rashi on Genesis 1:1*, or in Hebrew (*בראשית א*). You get one result per translation in the chosen language (without a choice: the screen language and English), and one result with the original text only. A whole book (*Genesis*) gives its first part.
- **Search words** in English or Hebrew (*love your neighbor*): the results are the passages that contain them, one per reference. With *Part of the library* you can narrow the search, for example to the Talmud. Put a word that is also a title (*Shabbat*) in quotation marks.
- A **link** to a text on sefaria.org; if the link names a translation, that one comes first.

Write references the way Sefaria does, with English book names (*I Samuel*, *Song of Songs*). If you search *Bibles*, *Original text* or *Marginal notes* for an Old Testament passage and Sefaria is switched on, a tip with a button points to the same passage in Sefaria. If *Original text and lexicon* is installed, the reference is converted to the Hebrew verse numbering (for example Malachi 4:5 → Malachi 3:23); otherwise the numbering may differ.

An item

An item is a reference with the original text (Hebrew or Aramaic) and one translation, side by side by verse, line or paragraph, marked with the language (*[HE]*, *[EN]*). Choose the translation with *Language of the translation*; without a choice it is the screen language, or English if there is none. Each version contributes only its own text, without footnotes and without filling gaps from other versions. A reference spanning several chapters or folios is retrieved part by part (at most 30 parts, or about 500,000 characters).

Related texts

At the end of an item you can see which commentaries and other texts on Sefaria refer to that text (for example Rashi or a passage in the Talmud). They are listed as references only: search for them to add them as items of their own.

Licence

On Sefaria every text version has its own licence; it is shown with every result and in the *Bron* section of the item.

- Public domain, CC0 and CC BY: free to use, with attribution.
- CC BY-SA: free to use, with attribution and share-alike.
- CC BY-NC: non-commercial use only. Such versions are labelled; if the installation is used commercially, they cannot be added (the administrator sets this under Administration > Settings > General).
- No derivatives (ND): anything is allowed within the organisation; a translation or summary by AI is not shared outside it.
- TrevantAI does not offer versions without licence information or under copyright.

An item gets the strictest licence of its versions; the default choice prefers a free version to a non-commercial one, unless the installation is used non-commercially: then Sefaria's own preference applies (for the Talmud, for example, the edition with vowel points). Attribution: *Powered by Sefaria*, with the titles of the versions.

Privacy

References and search words are sent to Sefaria, so please do not enter confidential information. The texts themselves contain no personal data.

See also: Theological sources: overview · Original text and lexicon