PAICON CLOUD GMBH

PaiX Navigator Privacy Notice

Version 2.0 · 20.10.2026

Contents

1. Who is responsible
2. Data collected automatically
3. Purposes, legal bases and retention
4. Registration and user account
5. Conversations: inputs, outputs and chat logs
6. AI-based processing and tools
7. Improving the Platform
8. Data requests and data products
9. Payment, support, service messages and analytics
10. Recipients and processors
11. International transfers
12. Cookies, similar technologies and the pages before you log in
13. The datasets the Platform works with
14. Storage periods
15. Your rights
16. Data security
17. Changes to this notice and version
18. Contact

1. Who is responsible

The controller for the PaiX Navigator is PAICON Cloud GmbH, Am Taubenfeld 21/2, 69123 Heidelberg, Germany, registered with the Amtsgericht Mannheim under HRB 740848. General business enquiries: info@paicon.com. This mailbox is not a channel for data protection requests. For those, see below.

Data protection officer. You can reach our data protection officer at privacy@paicon.com, or by post at the address above, marked for the attention of the data protection officer.

Where we ask for data you must provide, for example to create an account, we say so at the point of collection and tell you what happens if you do not provide it.

2. Data collected automatically

Server log files. Each time the Platform is accessed, our servers record information the accessing system sends: browser type and version, operating system, the IP address, and the date and time of access. These log files are stored on Amazon Web Services in Frankfurt, Germany, for 14 days and are then deleted. Where a statutory obligation requires us to retain specific log records for longer, we keep those records for the period the law prescribes, and the legal basis for that retention is Art. 6(1)(c) GDPR. Otherwise the legal basis is Art. 6(1)(f) GDPR, and our legitimate interest is the functionality, stability and security of the Platform.

Activity records in your account. Where you are logged in, we also record actions taken in your account, so that we can trace misuse and support you. These records contain your user ID, the action taken, its date and time, and your IP address. They are kept for 30 days and are accessible to our DevOps and administration staff. The legal basis is Art. 6(1)(f) GDPR, our interest being the security and integrity of the Platform and the traceability of use.

3. Purposes, legal bases and retention

We process personal data for the purposes below. The legal basis for each purpose is stated here and in the section dealing with it. Where we process special categories of personal data, a condition under Art. 9(2) GDPR applies in addition and is stated in the same place.

Purpose	Data	Legal basis	Retention
Operating and securing the Platform	Server log data (section 2)	Art. 6(1)(f), our interest in a functional, stable and secure Platform	14 days
Account activity records (section 2)	Actions taken in your account (section 2)	Art. 6(1)(f), our interest in the security and integrity of the Platform and the traceability of use	30 days
Providing your user account (section 4)	Registration data, date and time of registration	Art. 6(1)(b) where you contract with us directly; Art. 6(1)(f) where your organisation contracts with us and you use the Platform as one of its members	Duration of the contract; afterwards see section 14
Providing the conversation service (section 5)	Inputs, outputs, chat logs	Art. 6(1)(b) where you contract with us directly; Art. 6(1)(f) where your organisation contracts with us, our interest being to operate the service for the organisation that engaged us	Until your account is deleted; deleted within 30 days of your deletion request (section 14)
Metering your use (section 5)	Number of prompts, number of tokens processed	Art. 6(1)(b) where you contract with us directly; Art. 6(1)(f) where your organisation contracts with us, our interest being to meter the use of the service we provide to it	Until your account is deleted (section 14)
Security, prevention of misuse, correction of faults (section 5)	Inputs, outputs and chat logs, with sampled and indication-based review of conversation content by staff	Art. 6(1)(f), our interest in the security and integrity of the Platform and in keeping the service working	Until your account is deleted (section 14). Log data and activity records are covered by their own rows above and are deleted after 14 and 30 days
Improving the Platform (section 7)	Aggregated usage data (number and type of requests, features used, error rates, response times)	Art. 6(1)(f), our interest in understanding how the Platform is used and improving it	Underlying usage data until your account is deleted; the aggregated statistics do not identify you
Data requests and data products (section 8)	Questionnaire answers and request report, contact details, order and contract data	Art. 6(1)(b) for steps you ask us to take before a contract and for the contract itself where you contract with us directly; Art. 6(1)(f) where your organisation contracts with us, our interest being to handle requests from its members; Art. 6(1)(a) where you consent to your name and contact details being passed to a data partner; Art. 6(1)(c) for accounting records	Request records until your account is deleted. Accounting records for six, eight or ten years by record class (section 14)
Support (section 9)	Your message, contact details, account identifier	Art. 6(1)(b) where you contract with us directly; Art. 6(1)(f), our interest in answering members of a customer organisation	Until your account is deleted
Service messages and notifications (section 9)	E-mail address, message content	Art. 6(1)(b) for messages needed to run your account; Art. 6(1)(f) for security notices	Until your account is deleted

Purpose	Data	Legal basis	Retention
Error tracking and performance monitoring (section 9)	Error reports and performance metrics from our servers, including IP address	Art. 6(1)(f), our interest in detecting and fixing faults	30 days
Working with the datasets the Platform serves (section 13)	Dataset metadata and query results: pseudonymised health data, both from public cohorts and from partner institutions	Art. 6(1)(f) with Art. 9(2)(j) GDPR, § 27(1) BDSG and the safeguards of Art. 89(1) GDPR	Public cohorts: as long as the research purpose requires, reviewed annually. Partner datasets: the term of the licence agreement, reviewed on a change request or a legal obligation
Direct marketing (section 8)	Contact details, the subject of your earlier request	Art. 6(1)(f), our interest in offering related products, within the limits of § 7 UWG; consent under Art. 6(1)(a) where § 7 UWG requires it	Until you object or withdraw your consent
Disclosure on a sale or reorganisation of the business (section 10)	Account and usage data affected by the transaction; not the datasets described in section 13	Art. 6(1)(f), our interest in being able to transfer or reorganise the business	Until the transaction completes or is abandoned
Keeping you logged in and remembering your interface settings (section 12)	Session cookie; local storage entries for interface settings	§ 25(2) no. 2 TDDDG for the cookie and the local storage entries themselves; Art. 6(1)(f) for the data processed afterwards	Session cookie until you close your browser; local storage entries until you clear your browser storage
Statutory retention of commercial and tax records	Invoices, accounting vouchers, books, annual accounts, commercial correspondence	Art. 6(1)(c) GDPR with § 257 HGB and § 147 AO	Six, eight or ten years by record class (section 14)
Defending a legal claim	The records needed for the claim	Art. 6(1)(f), our interest in defending ourselves	Until the claim is resolved or time-barred (section 14)

Where we intend to process your personal data for a purpose other than the one it was collected for, we will tell you about that purpose, and give you the information in this notice that relates to it, before we begin (Art. 13(3) GDPR).

4. Registration and user account

Use of the Platform requires a user account. We process the data you enter in the registration form: e-mail address, first and last name, organisation or institute, and your role in that organisation. The fields marked as mandatory are required to create the account, and without them we cannot provide access. The date and time of registration are also stored. We do not offer sign-in through external identity providers.

Where you contract with us directly, the legal basis is Art. 6(1)(b) GDPR, performance of that contract. Where your organisation contracts with us and you use the Platform as one of its members, our legal basis for your account data is Art. 6(1)(f) GDPR, and our legitimate interest is operating the service for the organisation that engaged us. In both cases we are the controller for your account data and for the content you create on the Platform.

5. Conversations: inputs, outputs and chat logs

When you use the Platform, we process your inputs (prompts), the generated outputs, and the associated chat logs. This data is assigned to your user account.

We process it to provide the contractually owed service, to meter your use against the quota of your plan (in particular the number of prompts and the number of tokens processed), to ensure the security and integrity of the Platform, to prevent misuse and correct faults, and to keep the accounting records listed in section 14. The legal bases are Art. 6(1)(b) GDPR for the service and for metering where you contract with us directly, Art. 6(1)(f) GDPR for the service and for metering where your organisation contracts with us, Art. 6(1)(f) GDPR for security, the prevention of misuse and keeping the service working, and Art. 6(1)(c) GDPR for the retention of accounting records. Improving the Platform is dealt with separately in section 7.

Access by our staff. Our DevOps and administration staff can read identifiable conversations in two cases: for fault analysis and support, where a support or incident ticket requires it, and to detect and trace misuse, on the basis of random samples and where there is a specific indication of abuse. Each access must be requested from and granted by general management before it takes place, and every access is logged. This review serves the security purposes above and rests on Art. 6(1)(f) GDPR. It is not used to improve the Platform (section 7).

What should not be entered. Conversations are not intended to receive personal data about other people, or health data or other special categories of personal data, whether about you or about anyone else. This concerns what you type into a conversation. It is separate from the datasets the Platform works with, described in section 13. What may be entered is governed by our terms of use.

We do not ask for such data, the Platform is not designed to receive it, and we have no condition under Art. 9(2) GDPR that would allow us to process it for our own purposes. We therefore do not use it to improve the Platform, we do not include it in usage statistics, and we do not make it the subject of the sampling review described above. Where we become aware that such content has been entered, we remove it. You can delete a conversation yourself at any time (section 14).

How such content is blocked. A content filter (a guardrail model) screens every request before a model answers it. It is configured to detect health data and other special categories of personal data, as well as attempts at misuse, and where it detects them it blocks the request, so that no answer is generated and the content is not processed further. The filter is itself a model, so screening a request means sending the content of that request to it. The filter runs inside the European Economic Area. The screening is a security and data-minimisation measure under Art. 32 and Art. 25 GDPR. It is not a legal basis for processing such data, and it does not change the position stated above.

6. AI-based processing and tools

The Platform processes your inputs using large language models (LLMs). Inference runs through specific routes: Amazon Web Services Bedrock in the European Economic Area, Google Vertex in the European Economic Area, and the routing gateway Requesty Ltd in the United Kingdom, which forwards requests to the model providers on our fixed list, which are located in the European Economic Area and, under the configuration we have chosen, do not retain the content of requests or answers beyond the short-lived caching described below. Your prompts and the results of dataset queries are sent to these providers to generate the response. The route for a given request is chosen at run time according to the model selected and provider availability, so a conversation may be processed by any of the providers listed in section 10. Additionally, in order to optimise costs of LLM usage, the inference providers named above cache the content of requests for at most 60 minutes (prompt caching) and then delete it. The Platform itself keeps no separate cache of prompts or query results.

Within the scope of its agentic functionality the assistant may perform multi-step tasks largely automatically, invoke tools and retrieve and analyse data. Some of these tools reach services outside the Platform: scientific literature databases and clinical-trial registries. Code the assistant writes is executed on our own infrastructure and reaches no outside service. To carry out a task the assistant may pass parts of your conversation to one of these tools — for example a search query or the identifiers for a database lookup — and it chooses which to use from your request as it works, so the services involved differ from one conversation to the next. The tools, the services they reach and their countries are listed in section 10. The results are based on probabilistic methods and may be erroneous. How outputs may be used, and who bears the consequences of acting on them, is governed by our terms of use.

Automated decisions. The Platform does not take decisions about you based solely on automated processing that produce legal effects concerning you or similarly significantly affect you (Art. 22 GDPR).

7. Improving the Platform

We use aggregated usage data — such as the number and type of requests, the features used, error rates and response times — to produce usage statistics and to improve the assistant's instructions. We do not read the content of your conversations for this purpose, and we do not use your inputs, outputs or chat logs to train or fine-tune machine-learning models.

The legal basis for compiling these statistics from your usage data is Art. 6(1)(f) GDPR, and our legitimate interest is understanding how the Platform is used and improving it. The statistics themselves are aggregated and do not identify you, and we do not attempt to re-identify anyone from them.

You can object to the use of your usage data for these purposes at any time on grounds relating to your particular situation, and we will stop unless we can demonstrate compelling legitimate grounds that override your interests, rights and freedoms (Art. 21(1) GDPR). Write to privacy@paicon.com.

Review of conversation content to detect misuse is a security measure and is described in section 5.

8. Data requests and data products

If you request data or a data product through the Platform, the assistant guides you through a questionnaire about what you need. Your answers are stored with your account and compiled into a request report, which is sent by e-mail to our commercial team. A member of that team then contacts you personally, and no automated offer or delivery takes place. Any delivery of data follows a separate contract, and we process the order, contract and billing data required for it.

Our data partners are laboratories, pathology institutes, biobanks, hospitals and research institutions worldwide. Where your request concerns data held by one of them, we pass the content of your request to that partner. We do not pass your name or your contact details, and we do not pass your conversations. Where it becomes necessary for a partner to know who you are, for example to perform a contract you have concluded, we ask for your explicit consent first and pass only what you have agreed to.

The legal basis is Art. 6(1)(b) GDPR for the steps you ask us to take before any contract and for the contract itself, and Art. 6(1)(c) GDPR for the accounting records listed in section 14.

Contact about related data products. Where you are an existing customer of ours, we may contact you by e-mail after your request has been handled about data products related to the one you asked about. We do this only within the limits of § 7(3) UWG, which permits such e-mails to existing customers without their consent, and we do not contact users of the free pilot or enquirers who have not concluded a contract with

us in this way unless they have given their consent. The legal basis is Art. 6(1)(f) GDPR, our interest in offering you products related to the one you asked about. We send such messages only where § 7 UWG permits them without your consent, and where it does not, we ask for your consent first.

Your right to object to direct marketing. You can object to the use of your data for direct marketing at any time. You do not have to give a reason, and your objection costs you nothing. Once you object, we will stop using your data for direct marketing immediately (Art. 21(2) and (3) GDPR). Write to privacy@paicon.com, or use the unsubscribe link in any message we send you.

9. Payment, support, service messages and analytics

Payment. The Platform currently takes no payments: access is free of charge during the pilot phase, subject to a usage quota. If we introduce paid plans, we will update this notice before any payment data is collected.

Support. You reach our support by e-mail at support@paicon.com. Your messages, your contact details and, where relevant, your account identifier are held in our Microsoft 365 mailbox (Microsoft Ireland Operations Ltd., data stored in EU data centres) and are seen by the staff handling your request. Where fault analysis requires it, support staff may look at your conversations as described in section 5. The legal basis is Art. 6(1)(b) GDPR where you contract with us directly and otherwise Art. 6(1)(f) GDPR, our interest in answering the members of a customer organisation. Support correspondence is kept until your account is deleted.

Service messages and notifications. We send the messages needed to run your account — registration confirmation, password reset, and security and change notices — by e-mail through Amazon Simple Email Service in Frankfurt, Germany. The legal basis is Art. 6(1)(b) GDPR for messages needed to perform the contract and Art. 6(1)(f) GDPR for security notices. Sent messages are retained until your account is deleted.

Error tracking and performance monitoring. We use AppSignal (AppSignal B.V., Netherlands) to monitor our servers for errors and performance problems. AppSignal receives error reports and performance metrics from our systems, which may include technical request data including your IP address. The data is stored in the European Union. The legal basis is Art. 6(1)(f) GDPR, our interest in detecting and fixing faults. We do not use product analytics or session recording. Error data is retained for 30 days.

10. Recipients and processors

We disclose personal data to service providers that process it on our behalf and on our instructions. The data processing terms required by Art. 28 GDPR form part of our contracts with them. We also send search queries, which contain no personal data, to the operators of the public research databases the assistant queries (below), and we disclose personal data to public authorities where the law requires it.

Function	Provider	Country	What it receives
Hosting, infrastructure and code execution	Amazon Web Services EMEA SARL	Germany (Frankfurt)	All Platform data
Cloud inference	Amazon Web Services EMEA SARL (Amazon Bedrock)	European Economic Area	Prompts, dataset query results, outputs
Cloud inference	Google Cloud EMEA Limited (Vertex AI)	European Economic Area	Prompts, dataset query results, outputs
Routing gateway	Requesty Ltd	United Kingdom; data stored in Frankfurt, Germany	Prompts, dataset query results, outputs, forwarded to the model deployments we have configured for the Navigator, inside the EEA (section 6)

Function	Provider	Country	What it receives
Queries on TCGA data	Google Cloud EMEA Limited (BigQuery)	United States	Database queries the assistant generates
Support	Microsoft Ireland Operations Ltd. (Microsoft 365)	Ireland, EU data centres, support access from the United States possible	Support e-mails (section 9)
E-mail and notifications	Amazon Web Services EMEA SARL (Simple Email Service)	Germany (Frankfurt)	E-mail address, message content (section 9)
Error tracking and performance monitoring	AppSignal B.V.	Netherlands	Error reports and performance metrics from our servers, including IP addresses (section 9)

Research databases the assistant may query. To answer questions that need context beyond the dataset you are exploring, the assistant can query the public services below (section 6). It sends them only search terms and identifiers derived from your request. It does not send your name, your e-mail address or your account data, and it does not send data from the datasets the Platform serves. Personal data contained in your question does not reach these services either: the content filter described in section 5 blocks a request that contains such data before any query is sent. The queries are sent from our servers, so these services do not receive your IP address. The operators process the queries as independent controllers under their own privacy policies.

Service	Operator	Country
PubMed, PubMed Central, MeSH	US National Library of Medicine (NCBI)	United States
ClinicalTrials.gov	US National Library of Medicine	United States
Europe PMC	European Bioinformatics Institute (EMBL-EBI)	United Kingdom
bioRxiv, medRxiv	openRxiv	United States
arXiv	Cornell University	United States

Data requests forwarded to data partners. Data partners (laboratories, pathology institutes, biobanks, hospitals and research institutions) receive the content of data requests as independent controllers, in order to prepare offers and perform contracts (section 8). They do not receive your name or your contact details unless you have given your explicit consent.

If our business or part of it is sold or reorganised, personal data may transfer to the acquirer. This concerns your account data, your conversations and the other data described in sections 2 to 9. The basis for that disclosure is Art. 6(1)(f) GDPR, our interest in being able to transfer or reorganise the business, and the acquirer becomes responsible for the data as controller from that point, including for informing you under Art. 14 GDPR. The datasets described in section 13 are not covered by this paragraph: they would pass to an acquirer only after we have established, for that step, a condition under Art. 9(2) GDPR and any safeguard a transfer requires.

11. International transfers

To provide the service we use models, computing capacity and other infrastructure operated in part by service providers outside the European Economic Area, and the assistant queries public research databases outside the EEA. The recipients, the countries concerned and the safeguard relied on for each are listed

below. Where an adequacy decision exists, we base the transfer on Art. 45 GDPR. Where none exists, we rely on appropriate safeguards under Art. 46 GDPR, in particular the EU standard contractual clauses, with supplementary measures where the transfer assessment for that route requires them.

Recipient	Country	What is transferred	Safeguard
Requesty Ltd	United Kingdom; data stored in Frankfurt, Germany	Prompts, dataset query results, outputs	European Commission adequacy decision for the United Kingdom (Art. 45 GDPR). We configure the Navigator to use the EU endpoint and model deployments inside the EEA, so no onward transfer to a model provider outside the EEA takes place for this service.
Amazon Web Services (hosting, code execution, Bedrock inference, Simple Email Service)	Data stored in Frankfurt, Germany; inference in the European Economic Area; access from the United States for support possible	All Platform data stored on AWS; prompts, dataset query results and outputs for Bedrock inference; service e-mails	European Commission adequacy decision for the EU-U.S. Data Privacy Framework (Art. 45 GDPR): Amazon.com, Inc. is certified under the Framework, with Amazon Web Services, Inc. as a covered entity.
Google Cloud (Vertex AI inference)	European Economic Area, access from the United States for support possible	Prompts, dataset query results and outputs	European Commission adequacy decision for the EU-U.S. Data Privacy Framework (Art. 45 GDPR): Google LLC is certified under the Framework.
Google Cloud (BigQuery queries on TCGA data)	United States	Database queries the assistant generates	European Commission adequacy decision for the EU-U.S. Data Privacy Framework (Art. 45 GDPR): Google LLC is certified under the Framework.
Microsoft 365 (support mailbox)	Ireland, EU data centres, access from the United States for support possible	Support e-mails	European Commission adequacy decision for the EU-U.S. Data Privacy Framework (Art. 45 GDPR): Microsoft Corporation is certified under the Framework.
Europe PMC (EMBL-EBI)	United Kingdom	Search terms and identifiers the assistant sends	As for the row below: the assistant sends search terms and identifiers only, no personal data is transferred, and Chapter V GDPR does not apply to these queries.
PubMed, PubMed Central, ClinicalTrials.gov (US National Library of Medicine); bioRxiv, medRxiv (openRxiv); arXiv (Cornell University)	United States	Search terms and identifiers the assistant sends	The assistant sends search terms and identifiers only. It sends no account data, no data from the datasets the Platform serves, no personal data from your question, and not your IP address. No personal data is transferred, so Chapter V GDPR does not apply to these queries.
Data partners receiving request details (section 8)	Worldwide	Content of your data request, without your name or contact details; your name and contact details only where you have given your explicit consent	Where, with your consent, personal data is transferred: European Commission adequacy decision where one exists for the partner's country (Art. 45 GDPR), otherwise the EU standard contractual clauses for transfers between controllers (Art. 46(2)(c) GDPR)

You can obtain a copy of the standard contractual clauses we rely on by writing to privacy@paicon.com. Transfers based on an adequacy decision need no further safeguard, and the European Commission publishes these decisions on its website.

12. Cookies, similar technologies and the pages before you log in

We use a single cookie, `_paix_navigator_session`. It keeps you logged in and contains a security token that protects your requests against cross-site request forgery. The cookie is encrypted, cannot be read by scripts on the page (httpOnly), is sent only over encrypted connections (Secure), and is deleted when you close your browser. It is strictly necessary to provide the service you have requested, and it is used on the basis of § 25(2) no. 2 TDDDG and, for the personal data processed afterwards, Art. 6(1)(f) GDPR. In addition, your browser's local storage holds your interface settings, so that the Platform looks the way you configured it when you return; these entries are strictly necessary for the service in the form you have requested it and are used on the same basis. Apart from these, we store or read nothing else on your device: no analytics, advertising or tracking storage. Because nothing we store requires consent, the Platform shows no cookie banner.

What we store or read on your device:

Category	What is stored or read	Purpose	Set by	Lifetime	Third-party access
Necessary	Session cookie `_paix_navigator_session` (encrypted; httpOnly, Secure, SameSite=Lax)	Keeping you logged in, and protecting your requests against cross-site request forgery	PAICON, paix-navigator.paicon.com	Until you close your browser	None
Necessary	Local storage entries for interface settings	Remembering your interface preferences	PAICON, paix-navigator.paicon.com	Until you clear your browser storage	None

The pages before you log in. The login, registration and information pages collect the server log data described in section 2 and set only the storage described above.

13. The datasets the Platform works with

What you enter into conversations. The conversation function is not for health data or other special categories of personal data, whether about you or about anyone else, and our terms of use say so. We do not ask for such data and the conversation function does not require it. This is separate from the datasets described below, which the Platform works with. Section 5 explains what happens if such data reaches us anyway.

The datasets. The Platform lets you query biomedical research datasets across modalities — histopathology, radiology, molecular and clinical data — drawn from public cohorts and from partner institutions. The catalogue currently centres on pathology cohorts, is not limited to oncology, and grows as datasets are added. Users query the datasets. They do not upload data and they do not connect their own sources.

These data are personal data. The datasets we work with are pseudonymised health data. This applies to both kinds of source.

Datasets from partner institutions are pseudonymised by the partner before they reach us, with the direct identifiers removed. Public research cohorts are published by their publishers in de-identified form. We treat the data from both sources as personal data within the meaning of Art. 4(1) GDPR and, as health data, as special categories of personal data under Art. 9(1) GDPR.

We process them to make the datasets searchable and to answer your queries. The legal basis is Art. 6(1)(f) GDPR — our legitimate interest, and that of our users, in enabling biomedical research on existing data — together with Art. 9(2)(j) GDPR and § 27(1) BDSG. § 27(1) BDSG permits the processing of special categories of personal data for scientific research purposes where that processing is necessary and our interest in it

substantially outweighs the interest of the person concerned in not having it processed. The safeguards are those in Art. 89(1) GDPR and §§ 27(3), 22(2) BDSG: we process the data in pseudonymised form; the characteristics that could link a record to a person are not held by us at all but remain with the partner institution or, for public cohorts, with the publisher; we restrict access to the staff who need it; and we review at least once a year whether a dataset can be anonymised without defeating the research purpose, and anonymise it where it can.

The assistant answers from these datasets and is designed to return derived information rather than individual records. Query results are sent to the inference providers named in section 10 to generate the answer.

Where this information comes from. This paragraph informs the people whose data is in the datasets, who did not give it to us (Art. 14 GDPR). We obtain datasets from two kinds of source: publicly accessible research cohorts, released by their publishers in de-identified form, which we hold as our own copy or, in the case of The Cancer Genome Atlas, query where the publisher makes it available (section 10), and partner institutions — hospitals, pathology institutes and laboratories — under data licence agreements, which pseudonymise the data before transfer. The data comprise slide and case metadata, diagnoses, tumour characteristics, limited demographics and molecular markers. We process them for the purposes and on the basis stated above. They are disclosed to the inference providers named in section 10 when a query is answered, under the safeguards in section 11. We keep public cohorts for as long as the research purpose requires, reviewed annually, and partner datasets for the term of the licence agreement. No decisions producing legal effects on the people concerned are taken.

We do not inform each person individually, because we hold no contact details for them and obtaining them would involve disproportionate effort. We therefore make this information publicly available in this notice, which is the measure Art. 14(5)(b) GDPR provides for.

Rights in respect of dataset records. The rights in section 15 apply to this data. We hold the datasets under pseudonymous identifiers only. For partner datasets, the key that would link a record to a person stays with the partner institution; for public cohorts, any such link is held by the publisher. We cannot identify a person from a record ourselves. We can therefore act on a request about dataset data only where you give us information, such as the identifier of the record, that lets us find the record concerned, and we will act on it as soon as you do (Art. 11(2) GDPR).

14. Storage periods

We store personal data only for as long as is necessary for the respective purpose. Account data is stored for the duration of the contractual relationship, unless statutory retention obligations apply to particular records, or legitimate interests, such as the defence of legal claims, require longer storage. The periods are:

What	How long
Account data	Deleted within 30 days after you ask us to delete your account
Inputs, chat logs and outputs	Deleted together with your account, within 30 days of your deletion request
Results of dataset queries	Kept as part of the conversation they belong to, and deleted with it. There is no separate store of query results beyond the cache described in section 6
Cached prompts and query results	At most 60 minutes, at the inference provider that handled the request (section 6)
The datasets themselves	Public cohorts: as long as the research purpose requires, reviewed annually. Partner datasets: the term of the licence agreement (section 13)
Server log files; account activity records	14 days; 30 days (section 2)

What	How long
Request records from the data-request questionnaire	Until your account is deleted
Support e-mails and service messages; error data	Until your account is deleted; error data 30 days
Backups	Daily backups, kept for 30 days. Data you have deleted leaves the backups within 30 days at the latest
Inactive accounts	We do not delete an account for inactivity. Where your organisation's contract with us ends, the accounts of its members are deleted 30 days after the end of the contract
Accounting vouchers and invoices	Eight years (§ 257 HGB, § 147 AO)
Books, inventories and annual accounts	Ten years (§ 257 HGB, § 147 AO)
Commercial correspondence	Six years (§ 257 HGB)
Records needed to defend a legal claim	Until the claim is resolved or time-barred (§§ 195, 199 BGB)

The statutory periods run from the end of the calendar year in which the relevant event happened. That event differs by record class: the last entry in the books, the preparation of the annual accounts, or the receipt or despatch of the correspondence (§ 257(5) HGB, § 147(4) AO). The periods are extended for records that matter to a tax whose assessment period has not yet expired (§ 147(3) AO).

To have your account and your conversations deleted, write to privacy@paicon.com or use the deletion function in your account settings.

15. Your rights

You have the following rights in respect of your personal data:

- **Access (Art. 15 GDPR):** confirmation of whether and which personal data we process about you, including the purposes, categories, recipients, storage period and information on third-country transfers and the safeguards under Art. 46 GDPR, and a copy of that data.
- **Rectification (Art. 16 GDPR):** correction of inaccurate data and completion of incomplete data.
- **Erasure (Art. 17 GDPR):** deletion of your data where one of the grounds in Art. 17(1) applies and no retention obligation stands in the way.
- **Restriction of processing (Art. 18 GDPR):** where one of the grounds in Art. 18(1) applies, for example while the accuracy of the data is contested.
- **Notification of recipients (Art. 19 GDPR):** where we rectify, erase or restrict your data, we inform each recipient it was disclosed to, unless that proves impossible or involves disproportionate effort, and we tell you who they are if you ask.
- **Withdrawal of consent (Art. 7(3) GDPR):** where we have asked for your consent, for example to send you advertising by e-mail, you may withdraw it at any time with effect for the future. Withdrawing it does not affect the lawfulness of what we did before.
- **Data portability (Art. 20 GDPR):** where processing is based on consent or a contract and is carried out by automated means, you can receive the data you provided to us, including data about your use of the Platform that we observed, in a structured, commonly used and machine-readable format. It does not extend to conclusions we drew or content the Platform generated.

Your right to object (Art. 21(1) GDPR). Where we rely on our legitimate interests, you have the right to object at any time on grounds relating to your particular situation. We will then stop unless we can demonstrate compelling legitimate grounds that override your interests, rights and freedoms, or the processing serves the establishment, exercise or defence of legal claims. Write to privacy@paicon.com.

Your right to object to direct marketing (Art. 21(2) and (3) GDPR). This right is separate from the one above and is not subject to any condition. Where we process your data for direct marketing, you can object at any time, without giving a reason. Once you object, we stop processing your data for direct marketing immediately. Write to privacy@paicon.com, or use the unsubscribe link in any message we send you.

Your right to object to research processing (Art. 21(6) GDPR). Where we process personal data for scientific research purposes with the safeguards of Art. 89(1) GDPR, which is how we process the datasets described in section 13, you have the right to object on grounds relating to your particular situation. We will then stop unless the processing is necessary to perform a task carried out for reasons of public interest. Write to privacy@paicon.com. Section 13 explains what information we need from you in order to find the record concerned.

How we handle a request. We answer within one month of receiving it. Where a request is complex, or where we receive many, we may extend this by two further months and will tell you why within the first month (Art. 12(3) GDPR). Exercising your rights costs you nothing (Art. 12(5) GDPR). If we do not act on a request, we tell you the reasons, and that you may complain to a supervisory authority and seek a judicial remedy (Art. 12(4) GDPR).

Complaint. You have the right to lodge a complaint with a supervisory authority, in particular in the Member State where you live, where you work, or where the matter you are complaining about happened (Art. 77 GDPR). The authority competent for us is: Der Landesbeauftragte für den Datenschutz und die Informationsfreiheit Baden-Württemberg, Lautenschlagerstraße 20, 70173 Stuttgart, Germany, poststelle@lfdi.bwl.de, www.baden-wuerttemberg.datenschutz.de.

16. Data security

We take appropriate technical and organisational measures pursuant to Art. 32 GDPR to protect your data against accidental or unlawful destruction, loss, alteration, and unauthorised disclosure or access. Our measures are continuously adapted in line with technological developments.

17. Changes to this notice and version

We may update this notice so that it reflects current legal requirements and changes to our services. Where we make a material change, we will inform you by e-mail to the address in your account and by a notice in the Platform before it takes effect. Previous versions are available from us on request.

This is version 2.0 of the notice, dated 20.10.2026.

18. Contact

For questions about data protection and the exercise of your rights, write to privacy@paicon.com, which reaches our data protection officer, or by post to PAICON Cloud GmbH, Am Taubenfeld 21/2, 69123 Heidelberg, Germany, marked for the attention of the data protection officer. General business enquiries that do not concern your data: info@paicon.com.