



WHITE PAPER | AI GOVERNANCE | OCTOBER 2026

Operating AI Tools in Life Sciences

A practical guide to approving, controlling and documenting third-party AI, drawn from our experience with pharmaceutical, biotech and healthtech companies.

For CEOs, General Counsel, DPOs and Quality, IT Security and Clinical Operations leaders

Before We Begin

For two years, almost every conversation we have had with life-sciences companies about AI has started with the same questions. **Can we use this assistant? Who decides? What may my team put into it, and where does that data go afterwards? Does the AI Act apply to us yet?** And, more and more often: **what are companies like ours actually doing?**

These questions come from General Counsel, DPOs, heads of Quality, IT Security leads and Clinical Operations directors. They rarely come from people who want to block AI. They come from people who are being asked to say yes quickly, without a process that lets them say yes safely.

This guide is our answer. It is not a commentary on the AI Act. It is the operating model we have built from our experience of these problems across the sector. AI raises two distinct questions: a data governance question, about how a company's own data is shared and reused once AI tools are involved, and a compliance question, about the rules that apply to AI systems, in particular high-risk ones. We address both. The situations in this guide are ones we deal with regularly: an AI assistant blocked by customer contracts, an AI meeting companion that turns out to be a recording tool, a roll-out date set before any process exists and an AI assistant asked to review clinical-trial data. **The pain points are real, and we see them across the sector.**

You will find the regulatory picture, but mostly you will find what to do with it: who decides, on what evidence, within what time and what data may go into which tool. Every chapter ends with something you can use on Monday morning: a matrix, a checklist, a set of questions to ask a vendor.

Our thanks go to the iliomad team, who deal with these questions day to day and whose experience this paper distils.

Enjoy the read.

Seamus Larroque

Founder and CEO, iliomad Health Data

ABOUT THIS PAPER

This white paper draws on iliomad's practical experience of AI governance and AI compliance in the life-sciences sector. The situations it describes are typical of the problems we deal with. They are generalised and simplified, and no organisation, vendor or product is named.

It reflects the law and guidance available on 6 October 2026. It provides general information for life-sciences organisations and is not legal advice on any specific situation.

How to Read This Guide

This guide is written for organisations at every stage. It reads in one sitting, but it is also built to stay on your desk: practical sheets, checklists and decision tools are placed where you will need them during a project. Start where you are.

PROFILE 1 Just starting	AI tools are arriving and there is no framework yet. You need to understand where the pressure comes from and which rules already bind you. Read Chapters I to III in order, then the practical sheet ‘Five questions to situate your organisation’.
PROFILE 2 A tool is on your desk	A business team wants an assistant, a meeting companion or an AI feature switched on, and you need a decision you can defend. Go to the four situations in Chapter IV, then to triage, assessment and the permitted-use matrix (V.c to V.e).
PROFILE 3 Framework in place	You have a policy and an inventory. You want the model to run at speed, generate evidence and cover the AI you build. Go to V.f to V.h and Chapter VI: register, literacy, provider track, 90-day plan and indicators.

Along the way, **practical sheets** give you checklists to use as they are, **situations from our experience** show how recurring problems are resolved, and **key takeaways** sum up each situation or chapter in one sentence you can bring to your executive team.

PRACTICAL SHEET

Five questions to situate your organisation

Answer yes or no. Each ‘no’ points to the part of this guide that will help most.

- | | | |
|---|--|--------------------|
| 1 | Do you have a list of the AI tools and AI features actually in use, including those vendors have switched on? | See III and V.f |
| 2 | Do you know which customer or sponsor contracts restrict new sub-processors or the use of their data for model training? | See III.d and IV.a |
| 3 | Can any employee tell, in seconds, what data they may put into which tool? | See V.e |
| 4 | Is one person accountable for AI governance, with a named decision-maker for high-risk tools? | See V.b |
| 5 | Can a low-risk tool be approved within a week? | See V.c and VI |

Two or more ‘no’ answers: start with the 90-day plan in Chapter VI and do not wait for the full framework before deciding the tools already in the queue.

Contents

Before We Begin	2
How to Read This Guide	3
At a Glance.....	5
I. AI, Data Governance and Compliance	7
a. Two Questions behind Every AI Tool.....	8
b. Why AI Makes Data Governance Harder	9
c. Our Dual Approach.....	10
II. Why AI Tools Strain Your Controls.....	11
a. Four Ways AI Enters	12
b. Six Questions Due Diligence Misses	13
c. Four Beliefs That Stall Governance	14
d. Two Failure Modes	15
III. What Already Applies to Deployers	16
a. The AI Act after the Digital Omnibus	17
b. Data Protection Law	18
c. United States.....	19
d. Sector Guidance and Contracts.....	20
e. The Map at a Glance.....	21
IV. What We See in Practice	22
a. Situation 1: An Assistant Blocked by Contracts	23
b. Situation 2: An AI Meeting Companion Pilot.....	25
c. Situation 3: A Roll-Out Date without a Process.....	27
d. Situation 4: An AI Assistant on Clinical-Trial Data.....	28
e. Twelve Recurring Pain Points	30
V. The iliomad Operating Model	31
a. Two Tracks, One Foundation	32
b. Roles and Decision Rights	32
c. Intake and Triage	33
d. Four-Dimension Assessment.....	33
e. Permitted-Use Matrix.....	34
f. Decision Register and Life Cycle	36
g. AI Literacy by Role	36
h. When You Also Build AI	36
VI. A 90-Day Implementation Plan.....	38
a. The Roadmap.....	39
b. Three Pitfalls to Avoid	39
c. Measuring Whether It Works.....	39
d. Five Decisions for the Executive Team	40
VII. Conclusion.....	41
Annex: References	42
About iliomad Health Data.....	43

At a Glance

Most life-sciences companies will meet AI regulation as users of AI long before they meet it as builders.

The tools arrive through three doors: an enterprise assistant built on a large language model, an AI companion that records and summarises meetings, and AI features that vendors switch on inside software the company already uses. Each arrives faster than procurement, legal, privacy and IT security can assess it.

The difficulty is rarely the absence of rules. It is the absence of an **operating model**: who decides, on what evidence, within what time and what data may go into which tool. Without one, organisations swing between a blanket block that pushes use onto personal accounts and enablement without assessment that exposes health data, customer commitments and, for recording tools, criminal-law obligations that sit entirely outside data protection law.

Behind every tool sits a second question that compliance reviews rarely ask: where the company’s own data goes, and whether it is retained, reused or used to train models. AI governance has to answer both, which is why we treat it as a matter of **data governance** as much as of compliance (Chapter I).

4 assessment dimensions before any tool is connected	3 triage tracks, so low-risk tools are not held up	5 × 5 data tiers and tool statuses in one permitted-use matrix	2 Dec 2027 new application date for Annex III high-risk obligations
--	---	--	---

EIGHT KEY FINDINGS

- 01** **AI governance starts with data governance.**
Before classifying any system, a company must know which of its data goes into which tool, and whether it can be retained or reused to train models.
- 02** **The binding constraint is usually contractual.**
Customer data processing agreements, sub-processor notification and objection rights and sponsor confidentiality clauses decide what a tool may touch long before the AI Act does.
- 03** **Users need one answer: what may I put into which tool.**
A matrix that crosses data tiers with tool statuses answers it and makes the policy usable at the moment of use.
- 04** **AI meeting companions are recording tools first.**
They sit at the junction of recording law, biometric privacy, the GDPR and the AI Act prohibition on emotion recognition at work. Only a single, uniform consent standard is workable.
- 05** **Shuffling values or shifting dates does not anonymise a patient-level dataset.**
Anonymisation is a legal conclusion, not a technical setting.
- 06** **Governance must be fast to be followed.**
A fast track for low-risk tools, decided in days, removes the main incentive for shadow AI.
- 07** **The Digital Omnibus moved the high-risk dates, not the rest.**
Prohibitions, AI literacy, the GDPR and contractual obligations already apply to every deployer.
- 08** **Using AI on clinical-trial data is rarely prohibited.**
It adds a processor, a transfer and a risk to an existing study, and it must be cleared country by country before the first upload.


CHAPTER



I. AI, Data Governance and Compliance

AI raises two questions: where your data goes, and which rules apply to the system.

READ THIS CHAPTER IF you want to understand why AI governance starts with your own data, and how our dual approach is organised.



a. Two Questions behind Every AI Tool

Most discussions of AI in life sciences start with compliance: does the AI Act apply, is the system high-risk, which obligations follow. In our experience, **another question comes first** and is asked far less often: **where does our data go when we use AI, and what happens to it there?**

That question is one of data governance. Data governance is a company's ability to know and control how its own data is shared and used: which data leaves its environment, with whom, on what terms, for which purposes and what the recipient may do with it afterwards. In life sciences, that data is the core asset: clinical-trial results, research data, regulatory dossiers, manufacturing know-how, commercial plans and the data the company holds on behalf of customers, sponsors and patients.

The second question is one of compliance: the ability to show that an AI system meets the rules that apply to it, from the GDPR to the AI Act and, for the most sensitive uses, the requirements for high-risk systems. Both questions apply to every AI tool. They are answered by different people, with different evidence, and a company that answers only the second can meet its compliance obligations and still lose control of its most valuable data.



Figure 1: Two questions behind every AI tool

b. Why AI Makes Data Governance Harder

Data governance is not new. Life-sciences companies already control what goes to CROs, laboratories and software vendors through contracts, data transfer agreements and access rights. AI strains these controls for four reasons.

REASON 1

AI is everywhere

Assistants, meeting companions and features switched on inside existing software mean that data leaves through many small doors, often without a procurement decision.

REASON 2

AI reuses what it receives

Unless the contract excludes it, prompts, files and outputs may be used to train, fine-tune or improve models. What has shaped a model cannot, in practice, be recalled.

REASON 3

AI connects to everything

Connectors and agents reach mail, files and clinical systems, and vendors rely on model providers as sub-processors. One approval can open many data flows.

REASON 4

AI keeps traces

Inputs, outputs and logs are retained by the provider. Retention terms are not an absolute guarantee: data held by a provider can be preserved and disclosed by court order.

Figure 2: Four reasons AI strains data governance

For a life-sciences company, the consequence is concrete. Unpublished trial results, a candidate's development data or a sponsor's programme data entered into a tool on the wrong terms can be retained abroad, reused to improve a model that competitors also use or disclosed in litigation to which the company is not a party. None of this requires a security breach. An ordinary use on terms nobody checked is enough.

Fine-tuning adds a further layer. A company that adapts a model with its own data must know whether that data can later be extracted from the model, who else can use the adapted model and on what terms. In Opinion 28/2024, the EDPB confirmed that a model trained on personal data is not anonymous by default: its anonymity must be assessed case by case.

c. Our Dual Approach

We organise our AI work around these two questions. They call for different expertise and lead to different deliverables, but they share one inventory, one policy and one scoring scale.




Figure 3: Our dual approach to AI: AI Governance and AI Compliance

AI Governance applies to every company, because every company already uses AI. It is data governance applied to AI tools: deciding what data may go into which tool and on what terms, and keeping evidence of those decisions. AI Compliance applies where an AI system faces the AI Act's most demanding requirements: AI in a medical device or diagnostic that requires notified-body assessment, AI used in an Annex III area such as employment and AI that a company develops for others, which must first be qualified and classified.


The two approaches meet. The AI Act itself makes data governance a legal requirement for high-risk systems: Article 10 requires training, validation and testing data sets to be subject to data governance and management practices appropriate to the system's intended purpose. A company that already governs its data has answered many of the questions a compliance review will ask.

This paper follows the same order. Chapters II to VI address the AI a company uses; section V.h addresses the AI it builds.

KEY TAKEAWAY

Govern the data before the model. In life sciences, most AI exposure starts when company data goes where nobody decided it should go.


CHAPTER



II. Why AI Tools Strain Your Controls

AI arrives through routes that ordinary vendor due diligence was not designed to see.

READ THIS CHAPTER IF you are being asked to approve AI tools and want to understand why your usual checks are not enough.



a. Four Ways AI Enters

AI rarely arrives through a single procurement decision. In the organisations we support, it enters through four routes, each with a different risk profile and a different owner.

ROUTE 1 Enterprise assistants An LLM assistant procured for whole teams, often with connectors to mail, files and chat.	ROUTE 2 Meeting companions Transcription, summaries and action items built into the videoconferencing platform.	ROUTE 3 Features switched on AI functions that vendors activate inside software the company already uses.	ROUTE 4 Use outside any decision Personal accounts and free tools used when the approved route is slower than demand.
--	--	--	--

Figure 4: How AI tools enter a life-sciences organisation

b. Six Questions Due Diligence Misses

Conventional vendor due diligence asks where data is stored and who can access it. AI tools raise five further questions that standard questionnaires do not cover, and agentic tools add a sixth.

QUESTION 1

Inputs are unstructured

Whatever a user pastes into a prompt or says in a meeting becomes processed data, whatever its sensitivity.

QUESTION 2

Vendors depend on model providers

The tool vendor's own sub-processors, often several large model providers, process the content, sometimes in other jurisdictions.

QUESTION 3

Contract terms may permit reuse

Unless the contracted tier excludes it, inputs and outputs may be retained or used to improve models.

QUESTION 4

Connectors act with delegated authority

On-behalf-of access to mail, files and chat can expose far more than the user intends to share.

QUESTION 5

Outputs can be wrong and look right

Summaries, extracts and drafts need human verification before they enter regulated records, submissions or customer deliverables.

QUESTION 6

Agents remember and act

What the tool remembers and does on the user's behalf. The CNIL and the French AI and Digital Council identified persistent memory and delegated action as the main sources of risk in their July 2026 note on agentic AI.

c. Four Beliefs That Stall Governance

The obstacles are rarely where teams expect them. Four beliefs come up in almost every first meeting, and each one leads either to paralysis or to unmanaged exposure.

'The AI Act was postponed, so we can wait.'

The Digital Omnibus moved the high-risk dates. Prohibitions, AI literacy, the GDPR and contractual commitments already apply to every deployer (Chapter III).

'If we block it, we are safe.'

A blanket block pushes use onto personal accounts and unapproved tools, where nothing is assessed and nothing is recorded.

'We anonymised the data before using it.'

Shuffled values and shifted dates leave a patient-level dataset very likely personal. Anonymisation is a legal conclusion, not a technical setting (Situation 1).

'It is just a note-taker.'

A meeting companion is a recording tool first. Recording, criminal and biometric law apply before data protection law does (Situation 2).

d. Two Failure Modes

Organisations without an operating model tend to fall into one of two positions. Both are understandable; neither holds for long. The governed alternative sits between them.




Figure 5: Two failure modes and the governed alternative

KEY TAKEAWAY

The risk is not AI itself. It is a decision process slower than demand, or no decision process at all.


CHAPTER



III. What Already Applies to Deployers

Several regimes apply at once, and most of them already bind deployers today.

READ THIS CHAPTER IF you need to explain to your executive team which obligations are live now and which can wait.



A company that uses third-party AI tools is, in AI Act terms, a **deployer**. Several regimes apply at once and differ in nature: some are legal requirements, some are regulator guidance and some are contractual. Figure 7 at the end of this chapter maps them.

a. The AI Act after the Digital Omnibus

Regulation (EU) 2024/1689 (the **AI Act**) was amended by Regulation (EU) 2026/1744, the **Digital Omnibus on AI**, adopted on 8 July 2026, published in the Official Journal on 24 July 2026 and in force since 27 July 2026. The amendment postponed the high-risk obligations but left the earlier milestones in place.

2 Feb 2025	2 Aug 2025	27 Jul 2026	2 Dec 2027	2 Aug 2028
Prohibitions and AI literacy apply	General-purpose AI model obligations apply	Digital Omnibus in force; AI literacy becomes a duty of means	Annex III high-risk obligations apply (employment, among others)	Annex I high-risk obligations apply (AI in regulated products such as devices and IVDs)

Figure 6: AI Act milestones relevant to deployers, as amended by Regulation (EU) 2026/1744

Four points matter most for a deployer today.

ART. 5(1)(F) Emotion recognition at work is already prohibited AI systems that infer the emotions of people in the workplace are prohibited, except for medical or safety reasons. The Commission's February 2025 guidelines confirm that this covers inference from biometric data, including voice, and not sentiment analysis of written text. Check sentiment and engagement features in meeting and contact tools; we recommend switching them off by default.	ANNEX III, POINT 4 Employment uses are high-risk AI used for recruitment, promotion or termination decisions, task allocation or to monitor and evaluate performance is high-risk. A meeting summary tool is not high-risk by design; it becomes so if its outputs feed performance evaluations.
ART. 25 A deployer can become a provider A deployer that puts its name on a high-risk system, substantially modifies it or changes its intended purpose so that it becomes high-risk takes on provider obligations. Companies that fine-tune general-purpose models should also apply the Commission's July 2025 guidelines: a modification using more than one third of the original training compute makes the modifier a model provider.	ART. 4 AI literacy remains an obligation Since 27 July 2026, providers and deployers must take measures to support the AI literacy of their staff and of others operating AI on their behalf, rather than ensure a sufficient level. The duty is now one of means rather than result. It still calls for evidence that measures were taken.

b. Data Protection Law

The Digital Omnibus moved the high-risk dates. It did not move the prohibitions, the GDPR or the commitments in your contracts.

For most deployments, the **GDPR** and the **UK GDPR** create the earliest and most concrete obligations.

Topic	What it requires of a deployer
Processor terms (Art. 28)	The tool vendor is usually a processor and its model providers are sub-processors. Both must be covered by the contract and, where customer data is involved, by the customer's own sub-processor approval mechanism.
International transfers (Chapter V)	The transfer mechanism must cover the vendor's model providers, not only the vendor.
Lawful basis	Legitimate interest can support internal use, documented by a legitimate interest assessment. It is much harder to sustain for external participants who have no relationship with the company and do not expect AI processing.
DPIA (Art. 35)	New technology, large scale and systematic monitoring often come together in AI deployments, which makes a data protection impact assessment necessary.
Special category data and anonymisation	Health and genetic data require an Article 9 condition. Data is anonymous only if individuals can no longer be identified by any means reasonably likely to be used (Recital 26). In Opinion 28/2024, the EDPB confirmed that the anonymity of AI models trained on personal data must be assessed case by case.

Table 1: GDPR and UK GDPR obligations most often triggered by AI tools

c. United States

Topic	What it requires of a deployer
Recording consent	About a dozen states, including California, Florida, Illinois, Maryland, Massachusetts, Pennsylvania and Washington, require the consent of all parties to record a conversation. Where participants are in several states, apply the strictest rule.
Biometric identifiers	Illinois (BIPA), Texas and Washington regulate voiceprints. In December 2025, a putative class action filed in Illinois alleged that an AI meeting assistant captured participants' voiceprints without the notice and written consent BIPA requires. AI note-taking tools have since attracted further privacy claims.
HIPAA	Where a tool processes protected health information for a covered entity or business associate, a business associate agreement and a Security Rule risk analysis are required.
State AI laws	Colorado replaced its 2024 AI Act with SB 26-189, signed on 14 May 2026 and effective on 1 January 2027: notice, explanation and human-review duties for automated decision-making technology used in consequential decisions, including employment and healthcare. California's CCPA regulations on automated decision-making technology, in force since 1 January 2026, require compliance for significant decisions from 1 January 2027.

Table 2: United States rules most often triggered by AI tools

d. Sector Guidance and Contracts

In January 2026, the European Medicines Agency and the US Food and Drug Administration published ten common guiding principles of good AI practice in drug development, covering the medicines life cycle from research to post-authorisation. They are not binding, but they set expectations on human-centric design, context of use, data governance, risk-based performance assessment and life cycle management that inspectors and partners will increasingly apply. Where AI output feeds GxP records, the ordinary requirements for computerised systems and data integrity apply.

Contracts often bind first. Customer data processing agreements give customers information and objection rights over new sub-processors. Sponsor master services agreements restrict the disclosure of programme data, and broad clauses allowing data to be used to improve products and services rarely address model training explicitly.

e. The Map at a Glance

LEGAL REQUIREMENTS Binding today	REGULATOR GUIDANCE Sets expectations	CONTRACTS Often bind first
■ AI Act: prohibitions (Art. 5), AI literacy (Art. 4), role shift (Art. 25), high-risk from 2027 ■ GDPR and UK GDPR: Art. 28, Chapter V, lawful basis, DPIA, Art. 9 ■ Recording, criminal and biometric law; HIPAA; Colorado and California ADMT rules	■ Commission guidelines on prohibited practices and on the AI system definition (Feb 2025) ■ Commission GPAI guidelines (Jul 2025); EDPB Opinion 28/2024 ■ EMA and FDA guiding principles (Jan 2026); CNIL note on agentic AI (Jul 2026)	■ Customer DPAs: sub-processor notification and objection rights ■ Sponsor MSAs: confidentiality of programme data ■ 'Improve our services' clauses silent on model training

Figure 7: What applies to a deployer of third-party AI tools, by nature of obligation

KEY TAKEAWAY

Before asking whether the AI Act applies, ask what your contracts, data protection law and recording law already require. That is where the first constraints come from.


CHAPTER

IV

IV. What We See in Practice

Four situations we deal with regularly, and the pain points they share.

READ THIS CHAPTER IF a tool request is on your desk and you want to see how these situations are resolved.



The four situations below are ones we deal with regularly. Each follows the same structure: the situation, the pain points we see and how we address them.

a. Situation 1: An Assistant Blocked by Contracts

AT A GLANCE

TYPICAL PROFILE	Health-data analytics company processing sensitive clinical data for hospital and industry customers; it both uses third-party AI and develops AI within regulated software
TRIGGER	Business teams want to use a leading enterprise LLM assistant with company data
OUR ROLE	Two workstreams: a deployer framework for the AI the company uses, and a qualification and classification method for the AI it builds

i. Pain points observed

- **A blocked tool with no route to approval.** The assistant is blocked because its vendor is not contracted as a sub-processor, and no process exists to decide on what terms it could be used.
- **A commercial cost to adding a sub-processor.** Customer agreements give customers information and objection rights over new sub-processors, and management expects a substantial share of customers to object to a newly notified model provider.
- **An anonymised dataset that is not.** To feed an LLM-based dashboard, a team shuffles values and shifts dates in a patient-level dataset. The result is very likely still personal data.
- **Contracts silent on AI.** Customer agreements allow data to be used to develop and improve products and services, but do not mention model training.
- **No shared answer to the question: is it AI?** Marketing, engineering and legal use different definitions, and no in-house system has been classified under the AI Act.
- **Gaps in the foundations.** The AI inventory is incomplete, and there is no AI governance committee or accountable AI lead.




Figure 8: Situation 1: why the assistant stays blocked, and the way through

ii. How we address it

In this situation, priorities shift quickly. Contract review matters less than a practical answer to one question: **which data may go into which tool, and under what conditions**. We therefore build the deployer framework around a data classification work instruction and a permitted-use matrix (V.e), supported by:

- An eight-step onboarding procedure with three triage tracks, built into the existing vendor-onboarding process so that vendors receive one questionnaire, not two
- A scored four-dimension assessment in which commitments to customers are a blocking criterion
- An explicit rule that shuffled, perturbed or date-shifted datasets are treated as personal data unless the DPO has approved an anonymisation analysis
- A qualification form that tests each system against Article 3(1) and the Commission's guidelines, then classifies it, including through the Annex I route for AI in regulated diagnostic software
- An assessment of a privately hosted model as an alternative where adding a cloud model provider as sub-processor is commercially unacceptable

Anonymisation is a legal conclusion, not a technical setting.

WHAT THIS SHOWS

For a company that processes data on behalf of customers, the customer contract sets the ceiling. AI governance has to reach procurement and customer contracting, not only IT and legal.

b. Situation 2: An AI Meeting Companion Pilot

AT A GLANCE

TYPICAL PROFILE	Clinical-stage biotech with activities in worldwide	
TRIGGER	Pilot of the AI companion built into the company's videoconferencing platform: transcripts, summaries and action items	
OUR ROLE	As DPO, a data protection analysis, then a wider regulatory analysis with a recommended model for the pilot	

1	Recording law	All-party consent in about a dozen US states
2	Criminal law in Europe	Covert recording is an offence, e.g. section 201 of the German Criminal Code
3	Biometric privacy	Voiceprints protected by BIPA, Texas and Washington
4	GDPR	Lawful basis, DPIA, processors and transfers
5	AI Act	Art. 5(1)(f): no emotion inference from voice at work

Figure 9: Situation 2: five legal layers on one AI-transcribed meeting

i. Pain points observed

- **A patchwork of recording laws.** The company's home state requires all-party consent, as do about a dozen others. A click-through acknowledgment when the feature is enabled records the user's awareness, not the consent of every participant.
- **Criminal law in Europe.** Covert recording of private conversations is a criminal offence in several European countries, for example under section 201 of the German Criminal Code, independently of the GDPR.
- **External participants.** Legitimate interest can support internal meetings. For external participants with no prior relationship with the company, consent is the appropriate basis.
- **A DPIA trigger.** New technology, potential use across all meetings and systematic recording of the same individuals meet three DPIA criteria.
- **Voice as biometric data.** Voiceprints are protected identifiers under several US state laws, and AI note-taking tools are now the subject of biometric privacy litigation.
- **Features that cross a prohibition.** Meeting platforms often offer sentiment analysis. Where emotions are inferred from voice in a workplace setting, Article 5(1)(f) of the AI Act applies.
- **Model providers behind the vendor.** Meeting platforms rely on large model providers as sub-processors. Transfers may be covered by standard contractual clauses and the vendor's Data Privacy Framework certification while its documentation shows no clear opt-out from secondary use of meeting content, which then has to be confirmed with the vendor.

ii. How we address it

The mix of internal, external and cross-border meetings makes meeting-by-meeting compliance impractical. We recommend **one uniform standard for every meeting, set at the strictest applicable rule.**

PRACTICAL SHEET

A uniform consent standard for AI meeting companions

1. **Notice in every meeting invitation** that the meeting will be transcribed by an AI tool, with the purpose and a link to the privacy notice
2. **Opt-in consent in the platform from every participant**, internal and external, before the companion starts, with consent actions logged for audit
3. **Companion activated only if all participants consent**, so that a single refusal keeps it off for that meeting
4. **Sentiment and emotion features disabled** at account level
5. **An internal ground rule** on what may not be discussed while the companion is active, such as patient-level health data
6. **A DPIA** built on the vendor's published DPIA and completed with the company's own context of use
7. **A pilot limited to a defined user group**, without relaxing the consent standard
8. **No use of summaries in performance evaluation** without a separate assessment, since that use would fall under Annex III

WHAT THIS SHOWS

Meeting tools are recording tools first. The decisive obligations come from recording and biometric law, which a data protection review alone would not surface.

c. Situation 3: A Roll-Out Date without a Process

AT A GLANCE

TYPICAL PROFILE	Small, fully remote clinical-development organisation working for biopharma sponsors, with staff in the United States and the United Kingdom and no establishment in the EU
TRIGGER	An enterprise LLM assistant to be enabled for a first group of users within weeks, followed by a review of every application with AI features
OUR ROLE	As external DPO, an AI governance framework, a fast-track assessment of the first tool and AI literacy training

i. Pain points observed

- **A date set before a process exists.** The first users are scheduled to start within weeks. There is no assessment route, and decisions on AI tools sit with two executives.
- **Partial documentation.** A list of AI tools exists and an acceptable-use policy is partly drafted; there is no AI committee or AI lead.
- **Obligations that do not come from the AI Act.** With no EU establishment and no AI placed on the EU market, the main drivers are the UK GDPR, US state law and confidentiality obligations to sponsors.
- **Connectors to sponsor records.** The assistant is to connect to the company's collaboration suite, where sponsor programme records are stored.




Figure 10: Situation 3: a fast track that holds the roll-out date, then a framework the team runs

ii. How we address it

- A documented go decision on the first tool within five working days, with conditions of use
- A framework sized for a small team, which its own team runs for every new tool
- The vendor register and records of processing reused as the backbone of the AI inventory
- Literacy training for first users, a recorded module for later users and an attestation log
- Alignment with ISO/IEC 42001 and the NIST AI Risk Management Framework

WHAT THIS SHOWS

Governance must be proportionate. A small company needs a decision in days and a process its own team can run, not a committee structure designed for a multinational.

d. Situation 4: An AI Assistant on Clinical-Trial Data

AT A GLANCE

TYPICAL PROFILE	Clinical-stage sponsor running a multi-country trial with sites inside and outside the EU and data managed by a CRO
TRIGGER	Clinical Operations wants to use a general-purpose AI assistant to flag possible inconsistencies in pseudonymised EDC listings ahead of an interim analysis
OUR ROLE	As DPO, an opinion on whether the use is permitted and on what must be in place, or documented, before the first upload

1	Data status	Pseudonymised health data, not anonymous: the GDPR applies
2	Contract	Article 28 terms with the AI provider before any upload
3	Transfers	Standard contractual clauses, plus each country's own rules
4	DPIA	An update of the study DPIA, not a new one
5	AI Act	Minimal risk, as long as a person verifies every flag

Figure 11: Situation 4: five checks before the first upload of trial data

i. Pain points observed

- **Pseudonymised is not anonymous.** The listings carry no names and the sponsor does not hold the key, but they contain subject codes, age, sex, ethnicity, visit dates, dosing and clinical results. Identifiability is assessed from each actor's perspective (EDPS v SRB, C-413/23 P), but a sponsor that controls who may access participants' identities must treat such listings as pseudonymised health data.
- **A new processor outside the existing contracts.** The CRO agreements do not cover the AI provider, which processes the data on the sponsor's behalf. Article 28 of the GDPR requires a written contract before the first upload: in practice, a business subscription under the provider's data processing addendum. Consumer and personal accounts fall outside it.
- **A study DPIA that does not mention AI.** The DPIA covers data management and lists suppliers among the recipients, but not an AI tool, its provider or the transfer to the United States. Under Article 35(11), a change in risk calls for a review: an update of the existing DPIA rather than a new one.
- **Participant information that is usually sufficient.** Consent forms usually tell participants that coded data may go to the sponsor's service providers and outside their country. Since Article 13 allows recipients to be identified by category, a new provider used for the same purpose does not in itself require re-consent; where a form lists providers by name, the list is updated.
- **Country rules that differ.** Spanish law, for example, allows pseudonymised health data to be used for research only with a technical and functional separation from those holding the key and a commitment not to re-identify. Outside the EU, several countries require a

regulator's licence or permit, local standard contractual clauses or the participant's specific consent before data leaves for the United States.

- **Traces held by the provider.** Prompts, files and outputs are retained by the provider for a period. Retention terms are not an absolute guarantee: in 2025, a US court ordered an AI provider to preserve user conversations, including deleted ones, in litigation to which its customers were not parties.
- **A classification that depends on the use.** Flagging possible data-quality issues is neither prohibited nor high-risk under the AI Act: it falls outside Annex III and is not a medical device, since it makes no diagnostic, treatment or eligibility decision. The sponsor is a deployer whose main obligation is AI literacy, until the use extends to eligibility, dosing or safety assessment.

ii. How we address it

Our answer is usually yes, with conditions. We record it in a DPO opinion that sets out the conditions and the countries cleared, and **the conditions are met before the first upload, not after it.**

PRACTICAL SHEET

Eight conditions before trial data goes into an AI assistant

1. **A business account in the sponsor's name**, under the provider's data processing addendum with standard contractual clauses and no training on customer data; personal accounts prohibited
2. **An upload template limited to the fields the review needs**, without the identification key, free text or unnecessary dates
3. **A workspace configured before use**: restricted access, shortest retention, feedback sharing off
4. **An update of the study DPIA** and of the record of processing, with the DPO's advice
5. **Country-by-country clearance**: data from a country is uploaded only once its transfer and local requirements are met
6. **Participant information checked**, and provider lists updated where a country requires it
7. **Every flag verified against the EDC** by Clinical Operations before any query, with a record
8. **Files and conversations deleted** at the end of each review cycle

Beyond this use, we recommend a light framework of four elements: tool assessment, data classification, a register of AI use cases and user training. It is the deployer track of Chapter V in its simplest form.

WHAT THIS SHOWS

Using AI on clinical-trial data is rarely prohibited. It adds a new processor, a new transfer and a new risk to an existing study, and it must be cleared country by country before the first upload, not after it.

e. Twelve Recurring Pain Points

Across these four situations, twelve pain points recur. Each has an identifiable root cause and a control that resolves it.

Pain point	Root cause	Control that resolves it
Tools blocked indefinitely	No decision route or owner	Intake, triage and decision rights (V.b, V.c)
Shadow AI on personal accounts	Governance slower than demand	Fast track for low-risk tools (V.c)
Customer objections to new sub-processors	Contracts not checked before connection	Blocking contractual criterion and notification check (V.d)
Sensitive data pasted into prompts	No rule users can apply	Permitted-use matrix (V.e)
Anonymised data that is not	Technical transformation mistaken for a legal test	DPO approval of any anonymisation claim (V.e)
Over-broad connector access	Delegated permissions granted by default	Least privilege, read-only by default (V.d)
Recording without valid consent	Account-level acknowledgment treated as consent	Uniform all-party opt-in standard (IV.b)
Prohibited features left on	Vendor defaults not reviewed	Feature review at onboarding; off by default (V.d)
Disputes over what counts as AI	No agreed definition	Qualification against Art. 3(1) (V.h)
No evidence for auditors or partners	Decisions taken informally	Decision register and training records (V.f, V.g)
Proprietary data reused to train models	Consumer terms or default settings accepted	Business terms that exclude training, checked at assessment (V.d)
Trial data uploaded before every country is cleared	One global approval for a multi-country study	Country-by-country clearance and DPIA update (IV.d)

Table 3: Recurring pain points and the controls that resolve them


CHAPTER



V. The iliomad Operating Model

Two tracks on one foundation, built into processes a life-sciences company already runs.

READ THIS CHAPTER IF you need to design or upgrade the process that decides on AI tools.



a. Two Tracks, One Foundation

The model separates two questions that are often confused: how the organisation governs the AI it **uses**, and how it governs the AI it **builds**. Most companies need the first immediately; companies that develop software, devices or analytics also need the second. Both tracks share one policy, one inventory, one committee and one scoring scale, and both are built into existing processes rather than added beside them.




Figure 12: The iliomad operating model: two tracks on a shared foundation

b. Roles and Decision Rights

Every step has a single accountable owner. When one person holds several roles, as is common in smaller companies, accountability for high-risk decisions remains separate.

Role	Responsibilities
Executive management	Approves the AI policy and risk appetite; appoints the AI governance lead; accepts residual risk on escalated decisions
AI governance committee	Approves enhanced-track tools and any new AI sub-processor before customers are notified; reviews the inventory and incidents at least quarterly; settles disputes on what counts as AI
AI governance lead	Owens the policy, procedures and inventory; runs a monthly review; coordinates literacy. The AI Act does not require this role; we recommend it, and it can sit alongside the DPO
Business requester	Submits the intake; defines purpose, users and data
Assessment coordinator	Runs triage and the assessment; maintains the decision register
Legal, privacy and DPO	Contract, transfer, lawful basis and DPIA review; accountable for standard-track decisions
IT security	Security review of the tool and each integration; access configuration
Procurement	Turns contractual conditions into mandatory contracting terms; links the vendor record to the AI decision

Table 4: Roles in the deployer governance model

The procedure plugs into vendor onboarding at five points. Intake is joint, and existing vendor questionnaires are reused. Contractual conditions pass to procurement as mandatory terms. Vendor onboarding cannot close until the AI decision is recorded, and offboarding a vendor triggers a review of its AI tools.

c. Intake and Triage

Every request follows the same eight-step procedure. It starts with a short intake form completed by the requester, whose answers determine how deep the review goes.

STEP 1 Intake form	STEP 2 Triage	STEP 3 Vendor evidence	STEP 4 Assessment
STEP 5 Decision	STEP 6 Contract terms	STEP 7 Configuration	STEP 8 Register and review

Figure 13: Eight-step onboarding procedure for third-party AI tools


Eight yes or no questions route each request to one of three tracks. The fast track exists so that low-risk tools are decided in days; the enhanced track reserves committee time for the requests that need it.

EIGHT TRIAGE QUESTIONS 1. Personal data in inputs or outputs? 2. Health, genetic or other special category data? 3. Customer or sponsor data? 4. Connection to enterprise systems? 5. Audio or video recording? 6. A new sub-processor or model provider? 7. Outputs used in decisions about individuals? 8. Use in regulated (GxP) records or submissions?	FAST TRACK No to all questions: decided in days
	STANDARD TRACK Personal data or connectors: decided by legal, privacy and the DPO
	ENHANCED TRACK Customer data, special category data, recording, new sub-processor or decisions about people: decided by the committee

Figure 14: Triage: eight questions, three tracks

d. Four-Dimension Assessment

Each tool is assessed against four dimensions before it connects to enterprise systems. Criteria are scored on the three-point scale iliomad uses in its DPIAs: **0** acceptable, **1** could be improved, **2** not acceptable. Some criteria are **blocking**: a score of 2 prevents approval whatever the overall result.



Blocking criteria: **commitments to customers and sponsors, prohibited practices and use of company content for model training** that the contract does not exclude.

Figure 15: The four assessment dimensions and their blocking criteria

Scores convert into a decision in a fixed order, so that two assessors reach the same outcome on the same facts.




Figure 16: From scores to decision

PRACTICAL SHEET

Six questions to ask the vendor

They settle most assessments early. Send them with the intake, not after it.

1. Where are inputs, outputs and logs hosted?
2. How long are they retained, and can zero data retention be contracted?
3. Is customer content used to train or improve models, and is that excluded in the contract?
4. Which systems does the tool connect to, and with what permissions?
5. How can participants opt in or out during live use?
6. Can specific data categories be excluded from processing?

e. Permitted-Use Matrix

Policies fail when users cannot apply them at the moment of use. The matrix gives users a single rule: **identify the tier of the data, identify the status of the tool and read the answer.** When a dataset mixes tiers, the highest tier applies.

PRACTICAL SHEET

Eight rules that apply in every tool, whatever its status

1. Enter only the data the task requires.
2. Never enter credentials, keys or tokens.
3. Approve connectors per data source and scope, read-only by default.
4. Take no decision about an individual without the defined human oversight.
5. Verify outputs: AI output is not evidence in a regulatory submission, audit or customer deliverable without documented human verification.
6. Label AI-assisted content where required.
7. Do not store prompts or outputs containing T2 to T4 data in uncontrolled locations.
8. Treat anonymisation as a legal conclusion, not a technical setting.

Within those rules, the matrix sets what each tool status allows for each data tier.

Data tier / tool status	Not approved	Approved: general use	Approved: specific purpose	Approved: customer data	Private hosting / development
T0 Public	Yes	Yes	Yes	Yes	Yes
T1 Internal	No	Yes	Yes	Yes	Yes
T2 Confidential	No	Yes (1)	Yes	Yes	Yes
T3 Personal data	No	Yes (2)	Yes (3)	Yes (3)	Yes (5)
T4 Health and customer data	No	No	No	Yes (4)	Yes (5)

(1) With manager approval of the use case; no source code of regulated algorithms; no customer contract or tender text without Legal approval.

(2) De-identified data only, with DPO approval.

(3) Within the approved purpose and its record of processing; HR uses also require a DPIA, information to staff and, where relevant, a fundamental rights impact assessment.

(4) Only if all four conditions are met: customer contracts permit the sub-processor; notification is complete and any objection period has expired; a DPIA is complete; and the committee has approved the tool on the enhanced track.

(5) Within the approved purpose, with a DPIA, role-based access and, where customer data is used for development, a contract check.

Figure 17: Permitted use by data tier and tool status (all uses limited to the approved purpose)

Entering T3 or T4 data into a tool where it is not permitted, or discovering an unapproved tool in use, is handled as an **incident** and assessed under the breach procedure (*Articles 33 and 34 GDPR*).

f. Decision Register and Life Cycle

Every decision is recorded in a register with one of four outcomes: **approved**, **approved with conditions**, **rejected** or **escalated**. The register also holds the scores per dimension, each condition with its owner and deadline, the decision-maker, the next review date and a link to the vendor record.

Approval is tied to a purpose, a user population and specific connections; any change requires a new assessment. The reasons for a rejection are recorded so that the same request is not resubmitted unchanged. Where a customer contract conflicts with an internal approval, the contract prevails.

Five events trigger a re-assessment, and approved tools are reviewed periodically even without a trigger.




Figure 18: Life cycle of an AI tool decision

g. AI Literacy by Role

Literacy is what makes the matrix work in practice, and we build it by role. Attendance records and attestations provide the evidence that Article 4 now calls for: proof that measures were taken.

Audience	Content
All staff	General awareness: the approved tools, the permitted-use matrix and the eight usage rules
Tool owners and human overseers	Targeted module: conditions of approval, oversight duties, incident reporting
Developers	Targeted module: qualification, classification and the provider track
Management	Targeted module: risk appetite, escalated decisions and accountability
Clinical research and medical devices teams	Domain module, because their risks are specific

Table 5: AI literacy by role

h. When You Also Build AI

Companies that develop AI, including analytics, software as a medical device or AI components in IVDs, need a provider track. It reuses the deployer model's policy, inventory and scoring scale and adds three steps. For research-use models that carry no regulatory

obligation, applying a lighter version of the same method is a commercial decision: biopharma customers increasingly expect it.

1. **Qualification first.** A form tests each system against the seven elements of the AI system definition in Article 3(1), using the Commission's February 2025 guidelines. Inference is the decisive element; rule-based logic, basic statistics and simple heuristics fall outside. The conclusion is written so that a non-specialist can follow it, and marketing claims are aligned with it.
2. **Classification in sequence.** Territorial scope; prohibited practices; the Annex I route, under which AI in a medical device or IVD is high-risk only where the product requires notified-body assessment; the Annex III route; transparency duties; and general-purpose AI.
3. **Obligations mapped onto the existing quality system.** Each AI Act requirement is attached to an existing process: risk management to ISO 14971, logging to IEC 62304 and human oversight to IEC 62366-1, for example. The company then keeps one technical documentation set and one declaration of conformity.


CHAPTER

VI

VI. A 90-Day Implementation Plan

What to do first, second and third.

READ THIS CHAPTER IF you want to put the model in place this quarter.



Most organisations can deploy the model in about three months. **Sequence matters more than pace:** decisions on pending tools must not wait for the full framework.

a. The Roadmap

DAYS 1 TO 30 Decide and contain	■ Name the AI governance lead and the decision-maker for high-risk tools ■ Build the AI inventory from the vendor register and records of processing ■ Decide the tools already in the queue through a documented fast track ■ Publish interim usage rules and switch off emotion and sentiment features ■ Check customer and sponsor contracts for sub-processor and AI training terms
DAYS 31 TO 60 Build the process	■ Adopt the AI policy and data classification ■ Set up intake, triage and the four-dimension assessment inside vendor onboarding ■ Publish the permitted-use matrix ■ Open the decision register ■ Deliver general awareness training with attestations
DAYS 61 TO 90 Run and prove	■ Hold the first committee meeting and the first monthly review ■ Deliver role-based modules ■ Start the provider track where you build AI: qualification first ■ Report the first indicators to the executive team

Figure 19: A 90-day implementation plan

b. Three Pitfalls to Avoid

PITFALL 1 Waiting for the full framework Every week without a decision on pending tools feeds shadow AI. Decide them through the fast track first.	PITFALL 2 Copying a multinational's structure A small company needs a decision in days and a process its own team can run (Situation 3).	PITFALL 3 A policy without a matrix A policy users cannot apply at the moment of use will not be followed. The matrix is the policy in practice.
--	--	--

c. Measuring Whether It Works

A governance model is followed when it is fast, and trusted when it produces evidence. Five indicators, reviewed monthly by the AI governance lead and quarterly by the committee, show both.

Indicator	What it shows
Time to decision, by track	Whether the fast track is fast enough to beat shadow AI
Share of inventoried tools with a registered decision	How much of actual AI use is governed
Open conditions past their deadline	Whether approvals with conditions are honoured
Incidents and unapproved tools discovered	Where the matrix or the training is not working
Training completion and attestations	The evidence Article 4 calls for

Table 6: Five indicators to steer the model

d. Five Decisions for the Executive Team

PRACTICAL SHEET

Five decisions for the executive team

1. **Name one person accountable for AI governance**, and decide who signs off high-risk tools.
2. **Decide which data tiers may never leave the company's environment.**
3. **Check customer and sponsor contracts** for sub-processor and AI training terms before connecting any tool to their data.
4. **Switch off emotion and sentiment features** in meeting and contact tools.
5. **Set a target time for fast-track decisions**, and keep to it.

CHAPTER

VII

VII. Conclusion

Approve AI tools quickly and be able to defend those approvals later.

The AI Act timetable has moved; the use of AI in life-sciences companies has not waited for it. The obligations that bind deployers today come from the prohibitions already in force, data protection law, recording and biometric statutes and, most often, contracts with customers and sponsors. None of them is met by a policy alone.

What works is an **operating model**: decisions with an owner and a deadline, an assessment proportionate to risk, a rule users can apply in seconds and a register that shows what was decided and why. Built into existing processes, it lets a company approve AI tools quickly and defend those approvals later.

And it starts with the company's own data: knowing what is shared with AI tools, on what terms and for what reuse. Compliance for high-risk systems then builds on that foundation.

Annex: References

Primary sources cited in this paper, as available on 6 October 2026.

1. *Regulation (EU) 2024/1689* laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L, 12 July 2024. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
2. *Regulation (EU) 2026/1744* amending Regulation (EU) 2024/1689 (Digital Omnibus on AI), adopted 8 July 2026, published in the Official Journal on 24 July 2026, in force on 27 July 2026.
3. European Commission, *Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689*, C(2025) 884 final, 4 February 2025.
4. European Commission, *Guidelines on the definition of an artificial intelligence system established by Regulation (EU) 2024/1689*, C(2025) 924 final, 6 February 2025.
5. European Commission, *Guidelines on the scope of obligations for providers of general-purpose AI models under the AI Act*, July 2025.
6. *Regulation (EU) 2016/679* (General Data Protection Regulation). <https://eur-lex.europa.eu/eli/reg/2016/679/oj>
7. European Data Protection Board, *Opinion 28/2024 on certain data protection aspects related to the processing of personal data in the context of AI models*, 17 December 2024.
8. CNIL and Conseil de l'IA et du numérique, exploratory note on agentic AI and personal data, July 2026. <https://www.cnil.fr/sites/default/files/2026-07/ia-cianum-cnil.pdf>
9. European Medicines Agency and US Food and Drug Administration, *Guiding principles of good AI practice in drug development*, January 2026. https://www.ema.europa.eu/en/documents/other/guiding-principles-good-ai-practice-drug-development_en.pdf
10. German Criminal Code (*Strafgesetzbuch*), section 201, violation of the privacy of the spoken word.
11. Illinois Biometric Information Privacy Act, 740 ILCS 14.
12. *Cruz v. Fireflies.AI Corp.*, putative class action complaint, Illinois, December 2025. <https://commlawgroup.com/wp-content/uploads/2025/12/Fireflies.ai-Complaint-1.pdf>
13. Colorado Senate Bill 26-189 on automated decision-making technology, signed 14 May 2026, effective 1 January 2027.
14. California Privacy Protection Agency, CCPA regulations on automated decision-making technology, risk assessments and cybersecurity audits, effective 1 January 2026.
15. ISO/IEC 42001:2023, Information technology, Artificial intelligence, Management system.
16. National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, January 2023.
17. Court of Justice of the European Union, Case C-413/23 P, European Data Protection Supervisor v Single Resolution Board, judgment of 4 September 2025.
18. *The New York Times Company v. Microsoft Corporation and OpenAI*, US District Court for the Southern District of New York, preservation order of May 2025.
19. Spain, Organic Law 3/2018 of 5 December 2018 on the Protection of Personal Data and the Guarantee of Digital Rights, Additional Provision 17.


About iliomad Health Data

iliomad Health Data is a regulatory and compliance consultancy dedicated to the life-sciences sector. Its team of certified practitioners supports more than 30 pharmaceutical, biotech and healthtech organisations in more than 75 countries as external Data Protection Officer, as EU, UK and Swiss Data Protection Representative and as AI governance partner.

iliomad’s AI work follows the dual approach described in Chapter I.

PROTECT YOUR DATA AI Governance Data governance applied to AI: governance framework, AI inventory and risk tiering, control of the data shared with AI tools, third-party AI tool onboarding and AI literacy.	QUALIFY HIGH-RISK AI AI Compliance AI Act applicability and classification, provider and deployer gap assessments, fundamental rights impact assessments articulated with DPIAs and technical documentation for AI in medical devices and IVDs.		
30+ pharmaceutical, biotech and healthtech clients	75+ countries covered	10 certified compliance practitioners	ISO 9001 quality and 27001 information security management systems

We are practitioners rather than theorists: we run compliance programmes day to day, own the documentation and are measured on the outcome. Our work is delivered under an ISO 9001 quality management system and an ISO 27001 information security management system.



PLANNING AN AI TOOL ROLL-OUT?

Let us assess the tool with you before it connects to your systems.

Or qualify an AI system you are developing before it reaches the market.

contact@iliomad.fr | www.iliomadhealthdata.com

iliomad Health Data SAS, 64 avenue Pierre Grenier, 92100 Boulogne-Billancourt, France
iliomad Health Data UK Ltd, Salisbury House, Station Road, Cambridge CB1 2LA, United Kingdom

