INNOVATION-SANDBOX
for Artificial Intelligence

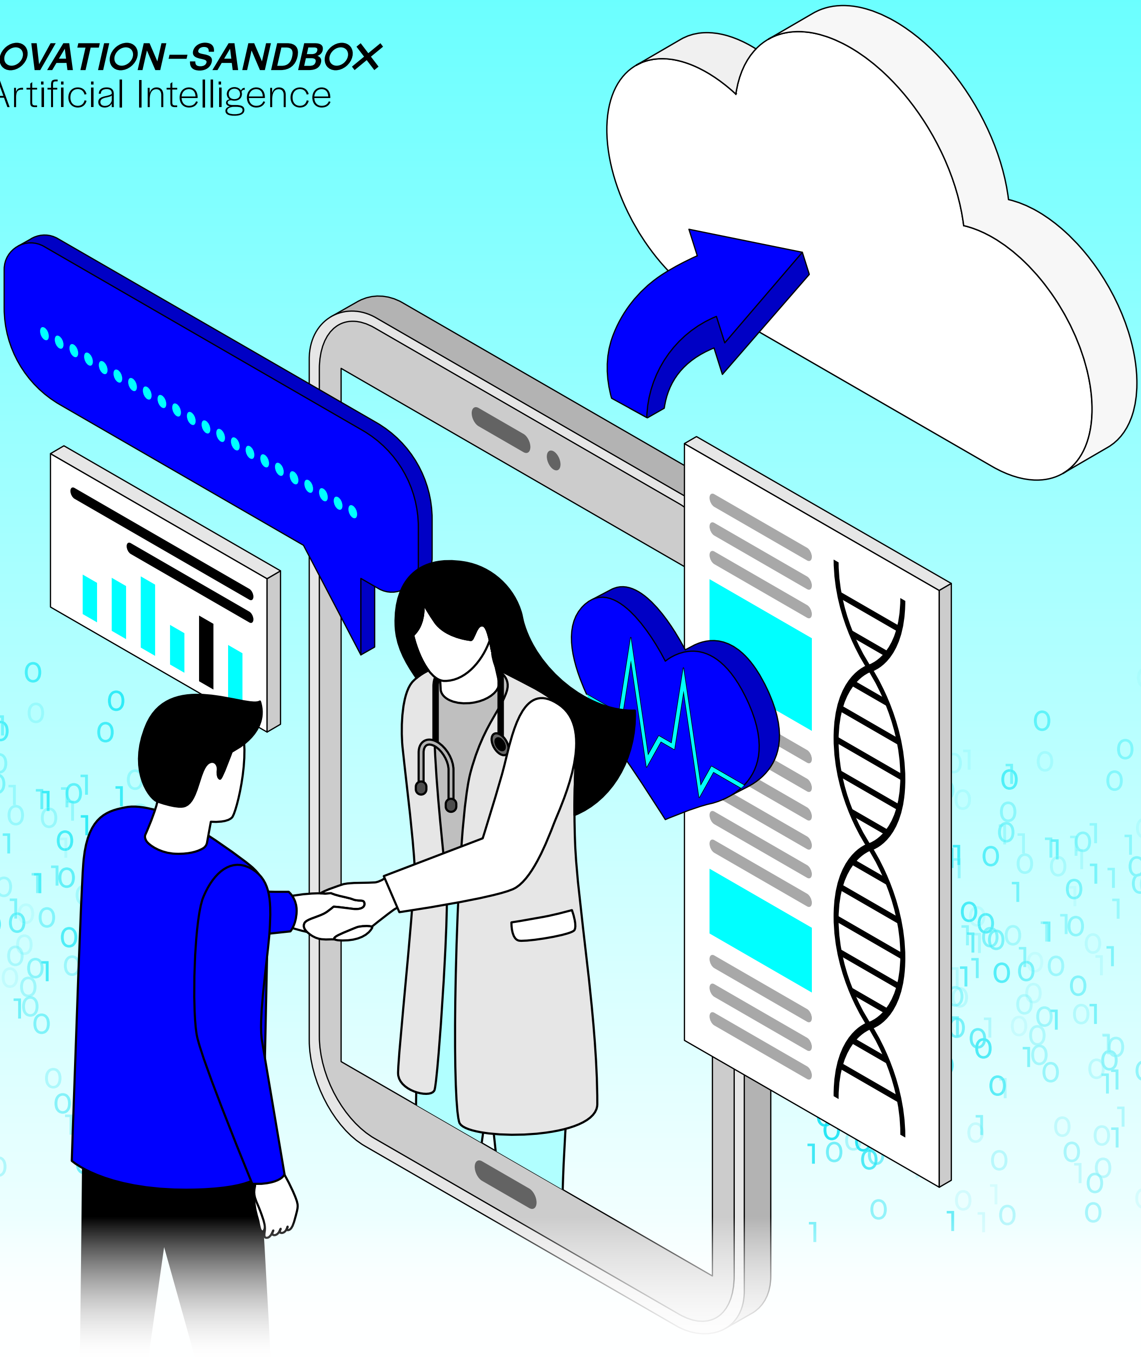

AI in medical documentation
Legal frameworks and recommendations

innovationsandbox.ai

metropolitan konferenz zürich

Kanton Zürich

Innovation Zurich

Increasing administrative requirements are placing a considerable burden on healthcare professionals. The preparation of medical reports takes up a lot of time, with specialist staff often transcribing voice recordings manually and, in some cases, still writing reports by hand. Artificial intelligence (AI) offers great potential here: modern speech recognition systems and large language models (LLMs) can efficiently transcribe, structure and improve the quality of reports. However, the use of AI-based solutions raises key data protection issues – particularly with regard to professional secrecy and the use of cloud services. In addition, it is difficult for many AI providers to estimate when systems of this kind will be classified as medical devices and which regulatory requirements they will need to meet as a result. As part of the Innovation Sandbox for AI, the Office for Economy of the Canton of Zurich and the Center for Information Technology, Society, and Law (ITSL) at the University of Zurich have worked with a number of experts to develop recommendations for the secure use of AI technologies in medical documentation. The results are intended to help to reduce the administrative burden in the healthcare sector while taking into account the highest data protection and security standards. The guide has been developed for providers of AI solutions in particular, but may also contain important information for hospitals, medical practices and other healthcare service providers.

Innovation Sandbox for AI

The project team drew up this document as part of the Innovation Sandbox for AI. The Sandbox is a test environment for implementing AI projects from various sectors. This broad-based initiative from government, business and research promotes responsible innovation by ensuring the project team and participating organisations work together closely on regulatory issues and enable the use of new data sources. The contents of this report are not legally binding and do not represent the official position of any public institutions. Any liability for legal aspects is excluded.

More information

Table of contents

With expert support from

Dr André Baumgart
Head of Quality Management and Patient Safety, VZK (Association of Zurich Hospitals)

Dr Nadine Bienefeld
Associate lecturer, ETH Zurich

Michèle Hess
Lawyer for digitalisation projects, Canton of Zurich Health Department

Dr Rolf Kaufmann
Senior Medical Device Expert

Raffaele Lugli
Head of Process Digitalisation & Innovation Office, Canton of Zurich Health Department

Dr Michael Neugebauer
Senior physician and IT officer, University Children's Hospital Zurich

Corinne Spirig
Chief Operating Officer, digital health center bülach (dhc)

Sebastian Svetel
Chief Information Security Officer, University Children's Hospital Zurich

Dr Atanas Todorov
Chief Medical Officer, Arcondis

Peter Waldner
Head of eHealth, Canton of Zurich Health Department

01.
The potential of AI in medical documentation

The health sector is facing major challenges: rising costs, increasing administrative burdens and an acute shortage of skilled workers are putting pressure on the system. Medical professionals spend a significant proportion of their working time on documentation-related tasks, leaving less time for direct patient care[1]. The consequences are work overload, stress and high levels of staff turnover within the healthcare system. A survey by the Swiss Federal Statistical Office shows that work-related stress is particularly common in the health and social sector, with far-reaching consequences for the quality of care and employee satisfaction.

One particularly time-consuming area is the preparation of ***medical reports***. Medical reports summarise diagnoses, findings, assessments, further procedures and recommendations from specialists. They also serve as a basis for benefit reviews and billing for medical services. Writing these documents requires an in-depth understanding of the context: it is important to describe medical issues precisely, use medical terminology correctly and summarise the subsequent course of treatment in a structured manner. Doctors often dictate their findings and instructions, with trained professionals transcribing, structuring and expanding on them.

In some health facilities, medical professionals still write medical reports by hand. This process is not only labour-intensive but also prone to error. Given the current shortage of skilled workers, such administrative tasks increase the pressure on staff.

Advances in the field of ***artificial intelligence (AI)*** offer great potential for increasing efficiency[2]. ***Speech-to-text models*** (e.g. OpenAI's Whisper or Google's Speech-to-Text AI) and ***large language models (LLMs)*** (e.g. OpenAI's GPT models or Anthropic's Claude) can quickly transcribe voice recordings and format the texts automatically according to medical standards. AI-supported processes enable reports to be created faster, more consistently and with less human effort. LLMs also allow for many other use cases, such as the optimisation of existing reports, suggestions for ***differential diagnoses*** and automated invoicing (see Chapter 4 on distinguishing between different use cases). However, the use of AI also involves risks, such as incorrect transcriptions and factual errors in the automated processing of medical content. Without careful validation, negligent use of AI results can lead to direct patient harm, lack of quality, inefficiency and liability issues. With this in mind, human oversight remains essential.

In view of the technological potential, numerous local start-ups and established companies are

* The terms highlighted in blue are explained in the glossary on page 39

[1] Christino et al. 2013: Paperwork versus patient care: a nationwide survey of residents' perceptions of clinical documentation requirements and patient care (Link).

[2] Perkins et al. 2025: Improving Clinical Documentation with AI (Link).

01. The potential of AI in medical documentation

working on innovative AI solutions to reduce the amount of documentation required. They are already now piloting these solutions in collaboration with a host of healthcare providers. In the process, each organisation carries out its own legal audits, leading to different assessments and conclusions. The use of AI in medical reporting raises complex questions. AI providers and healthcare institutions must comply strictly with data protection, professional secrecy and regulatory requirements for ***medical devices***. What are the requirements for transcribing medical reports using cloud services? Is the ***anonymisation*** of health data in medical documentation through AI a viable option? Furthermore, when does a solution of this nature fall under the regulatory requirements for medical devices, for example if LLMs provide suggestions for differential diagnoses?

To address these questions, the Office for Economy is working with ITSL at the University of Zurich to develop a legal basis for this specific use case as part of the Innovation Sandbox for AI. The project team involved various partners from industry, science and administration in this process. The aim of this report is to create a common basis for all stakeholders in the health sector.

> *«AI-based tools for medical documentation are spreading rapidly in the healthcare sector.»*
>
> *Raphael von Thiessen,*
> *AI Sandbox Programme Manager,*
> *Canton of Zurich*

02. *Data protection, professional secrecy and the cloud*

The previous chapter shows that the use of AI in medical documentation brings a number of benefits, but also raises complex legal issues. In the healthcare sector in particular, personal data, its collection, documentation and further use are subject to strict data protection and confidentiality obligations. The following chapter provides a structured overview of the most important legal requirements, especially with regard to data protection, professional secrecy and the use of cloud technologies. The answers to frequently asked questions are intended to provide clarity and a basis for the legally compliant implementation of AI-based solutions. Chapters 3 and 4 below clarify the classification of different use cases as medical devices.

2.1 Basis

When is data protection law relevant for healthcare institutions?

Data protection law always applies in instances involving personal data. Personal data is all information relating to an identified or identifiable natural person (Art. 5(a) of the Swiss Federal Act on Data Protection (FADP, SR 235.1) or Section 3(3) of the Information and Data Protection Act (IDG) of the Canton of Zurich (IDG, LS 170.4)). This applies regardless of whether the institution is a private practice or a cantonal hospital. The definitions of personal data largely correspond to the various data protection laws (for the applicable legal frameworks, see the following question). A person is identified if their identity can be directly established from the data; in the context of medical reports, e.g. by their name, date of birth or OASI (AHV) number. The person is identifiable if they can be identified through a combination of other data, such as place of residence and occupation.

With respect to medical practices and hospitals, patient data, inter alia, must be classified as personal data. Patient data encompasses all personal information (e.g. contact details, insurance numbers) and health data (e.g. findings, diagnoses) that is collected in connection with the treatment of patients.

If a healthcare provider uses AI systems to create or improve medical reports, personal data is regularly processed, which is why data protection requirements must be observed.

What special features need to be considered when processing health data in medical reports?

Health data is considered to be particularly sensitive personal data according to Art. 5(c)(2) FADP or special personal data according to Section 3(4)(a)(2) IDG, the processing of which is subject to stricter requirements under data protection legislation due to a particular risk of the infringement of privacy rights. As an example, this may include data about a person's state of health, such as medical findings, or treatment data such as therapies, diagnoses, medical history, genetic data, disabilities or information about mental illnesses. Even the appearance of a patient's name in connection with a doctor – when contacting a practice, for example – constitutes health data. This information is particularly sensitive because it is directly related to a person's physical and psychological integrity. This category may include much of the data collected in a health facility.

In addition to data protection considerations, medical reports are subject to professional secrecy (Art. 321 Swiss Criminal Code or cantonal confidentiality obligations for health professions such as Section 15 of the Health Act). Health professionals such as doctors and nurses are obliged to maintain professional or medical secrecy. They must treat all information received as confidential and, in principle, may not disclose any data to third parties. Information and personal data may only be processed by persons entrusted with confidential information and their auxiliaries. Exceptions are made if a legal provision stipulates otherwise, if the data subject has given their consent in individual cases, or if a superior authority releases the party subject to confidentiality in individual cases. However, consent is often not a suitable tool in practice, because this is obtained separately for the respective instance of data processing and an alternative must be offered in the event that consent is refused.

What are the fundamental differences in data processing between a private doctor's practice and a public hospital?

Public-law healthcare providers, such as a cantonal hospital, are subject to public law. This also applies to private hospitals with a public service mandate. The respective cantonal data protection laws therefore apply with regard to data protection, such as the Information and Data Protection Act for the Canton of Zurich (IDG ZH). By contrast, private medical practices are subject to the Swiss Federal Act on Data Protection (FADP) from a data protection perspective. Supervision of compliance with data protection laws is the responsibility of the Federal Data Protection and Information Commissioner (FDPIC) for the FADP and of the cantonal data protection authorities for cantonal legislation.

The funding body of the healthcare provider is relevant to the classification. Hospitals and medical practices operated by **cantonal, communal or public-law bodies** are subject to data protection regulations under public law. Individual or group practices without a public service mandate and private clinics are subject to the provisions of the FADP for private processors. Hospitals in particular can also be structured on a mixed basis (e.g. public-private limited companies). In such cases, it is essential to determine the area in which the data is processed. Whether the institution performs a sovereign (i.e. a state) task is also relevant for the classification. This can lead to a healthcare provider having to comply with different data protection laws depending on the area of treatment.

There are also some special provisions under health law, such as the provisions of the Federal Health Insurance Act (KVG), the Private Insurance Act (VVG) and the cantonal health laws.

When developing AI systems for medical documentation, consideration should be given at an early stage to whether they will be used in private practices (under the FADP) or in public-sector institutions (under cantonal data protection laws). The FADP's requirements for private processors and cantonal data protection laws are based on similar data protection principles, but they differ significantly: in particular, according to the IDG ZH, data processing must have a legal basis and can only be justified with consent in exceptional cases. The different legal requirements should be reflected in organisational measures and in the system architecture, e.g. through adaptable data access rights based on a modular system, logging functionalities or deletion mechanisms. AI providers can therefore ensure that the system is operated in a legally compliant manner depending on the customer group.

What requirements apply to anonymised data?
Some AI providers recommend that their customers anonymise the content of medical reports before they are processed by their system, and in some cases offer their own anonymisation solutions. Data is anonymised if it does not allow any conclusions to be drawn about an individual. Fully anonymised data is no longer subject to the requirements of data protection legislation. It should be noted that data is considered anonymised only if all identifying features have been removed (name, address, date of birth, patient number, and any other combinations that allow conclusions to be drawn). However, the data is solely considered to be anonymised if re-identification is no longer possible. For example, a medical report is not yet considered anonymised if only the patient's name has been removed but information such as the patient's date of birth, place of residence and rare diseases remains unchanged. With medical reports in particular, even seemingly 'harmless' data such as a patient's postcode, gender and clinical picture can be combined to enable identification. In instances involving rare diseases or places with small populations, it is often easy to identify the individual behind the data. Data is likewise not considered to be anonymised if the patient's name has been replaced by a code, but the practice maintains a list of names and codes. In principle, whether data has been anonymised can be verified based on whether it is still possible to determine the identity of the patient with a reasonable amount of effort. If this is possible, this does not constitute anonymisation.

«Legally compliant anonymisation of medical reports is practically unfeasible.»
Stephanie Volz,
Managing Director ITSL, University of Zurich

The anonymisation of medical reports is also challenging because modern technologies and analysis methods (such as AI-supported text analysis or reconstruction attacks) enable conclusions to be drawn about individual persons, even from data sets that have apparently been neutralised . AI-based anonymisation solutions are additionally reaching their limits in the face of increasingly powerful reconstruction methods. If, for example, a free-text medical report is anonymised, combinations of data pertaining to the progression of rare diseases, treatment dates and age may be sufficient to identify the person in question in a small hospital or a specific patient group. In practice, this means the anonymisation of health data requires the removal of extensive parameters, which calls into question the usability of the data.

If special tools are used for anonymisation, it should be noted that data protection regulations must also be observed when processing personal data up to the point that the data is anonymised. Anonymi-

[3] Morris et al. 2024: DIRI: Adversarial Patient Reidentification with Large Language Models for Evaluating Clinical Text Anonymization (Link)

sation tools must likewise be compliant with data protection.

What requirements apply to pseudonymised data?
Pseudonymised data is data that cannot be attributed to a person without the provision of further information. The essential element here is a key that allows the data to be reassigned to a person. Pseudonymised data continues to be considered personal data, meaning the applicable data protection laws must be observed. However, if the key and the data record are kept strictly separate from each other, it may be possible to carry out further data processing than would be the case with plain data (which can be directly attributed to a person; see section 2.3 below).

2.2 Data protection and AI tools

From a data protection perspective, are there special requirements for the use of AI tools for medical documentation by a private healthcare provider?
When using AI systems in medical documentation, private healthcare providers must comply with the requirements of the Swiss Federal Act on Data Protection. In Switzerland, there are currently no specific regulations for AI; the principles of data processing must be observed. AI tools in the medical field are often used via specialist third-party providers which provide the appropriate software and/or infrastructure. From a data protection law perspective, certain provisions regarding the use of contract data processors and the storage of data in a cloud must be taken into account (see below).

What is the situation in a public hospital?
Public hospitals must comply with the provisions of cantonal data protection laws. This also applies to outsourcing, which includes storing data in a cloud. The data protection officers of the Canton of Zurich have created an information sheet for AI-based data processing by cantonal and communal public institutions that sets out how to proceed before using AI applications (including the obligation to submit the project to the data protection officer for prior checking). The information sheet is available at: Information sheet on the procedure for using AI at public institutions.

Do private healthcare providers have to inform patients about the use of AI tools?
It depends on the task performed by the AI tool. Under Swiss law, there is no explicit obligation to provide information about the use of AI tools. However, an obligation of this nature may arise from a **medical** or **data protection** perspective **under certain circumstances**. It is conceivable that, as part of the implementation of the Council of Europe's AI Convention, transparency obligations will be implemented for certain AI providers or organisations that use them.

From a medical law perspective, clarification and information are required if the AI system significantly changes the conventional processing of data or if specific risks to patient safety arise from the use of the AI system, e.g. if the system is used for diagnostics or treatment and qualifies as a medical device.

An obligation to provide information may arise, for example, from the principle of transparency under data protection law, e.g. if patients **interact directly with AI** such as a **chatbot**. A further duty to provide information may arise if organisations **disclose personal data to third parties**, such as a **cloud provider** (see section 2.3 below). No separate clarification or information is required if an AI tool is used exclusively for administrative purposes, if the data processing procedure is comparable to the previous manual processing (e.g. transcription from dictation by AI instead of humans) or if the AI tool has a purely supporting function.

What is the situation for public hospitals?
From a data protection perspective, public hospitals are not obliged to provide information as long as the necessary legal bases for processing exist. Consent (on a case-by-case basis) would only be required if the use of the AI tool entailed a change of purpose. However, a duty of disclosure and information may nevertheless arise from a medical law perspective (see above).

Do patients have to consent if private healthcare institutions use AI tools?
From a medical law perspective, patient consent is required if an AI tool makes clinically relevant decisions or supports doctors' medical decisions, and if particular risks arise during treatment. However, consent is not required if an AI system is used for purely administrative purposes, such as improving a medical report or sending appointment invitations. Consent may also be required if the data collected by the AI tool is used for research purposes.

What is the situation for public hospitals?
Consent is also required from a medical law perspective in public hospitals.

Who is considered to be the controller under data protection law when using an AI tool at a private healthcare facility?
The private healthcare institution is generally considered to be the controller for data processing. The third-party provider that provides software or infrastructure should generally be qualified as a commissioned data processor. As the controller, the healthcare facility is responsible for ensuring compliance with the FADP. Data subjects may only assert their rights of access, rectification, deletion, objection or disclosure of their personal data vis-à-vis the controller. If the commissioned data processor receives such a request, it must forward it to the controller.

What is the situation for public hospitals?
The public hospital is always responsible for data processing; the third-party provider is the commissioned data processor. The public hospital must ensure compliance with the IDG. As far as the rights of data subjects are concerned, the same applies as for private healthcare institutions.

Do private healthcare providers have to carry out a data protection impact assessment?
A data protection impact assessment (DPIA) in accordance with Art. 22 FADP is mandatory for private controllers if data processing involves a high risk to the identity or fundamental rights of the data subjects. In particular, there is a high risk when new technologies are used and when a large amount of particularly sensitive personal data is being processed. Numerous medical AI applications outside the purely administrative sphere are likely to involve a high level of risk and therefore require a data protection impact assessment (DPIA).

Private individuals must consult the Federal Data Protection and Information Commissioner (FDPIC) if, despite the measures provided by the controller, the planned processing still results in a high risk to the identity or fundamental rights of the data subject (Art. 23 FADP).

Do public hospitals also have to carry out a data protection impact assessment?
Yes. In the Canton of Zurich, a DPIA must be carried out before any intended processing of personal data (Section 10(1) IDG). In addition, if AI applications are used, the project must be submitted to the data protection officer for prior checking, as this involves new technologies that pose particular risks to the fundamental rights of the data subjects (Section 10(2) IDG in conjunction with Section 24 of the Ordinance on Data Protection and Information Secu-

[4] Bienefeld et al. 2023: Solving the explainable AI conundrum by bridging clinicians' needs and developers' goals (Link).

02. Data protection, professional secrecy and the cloud

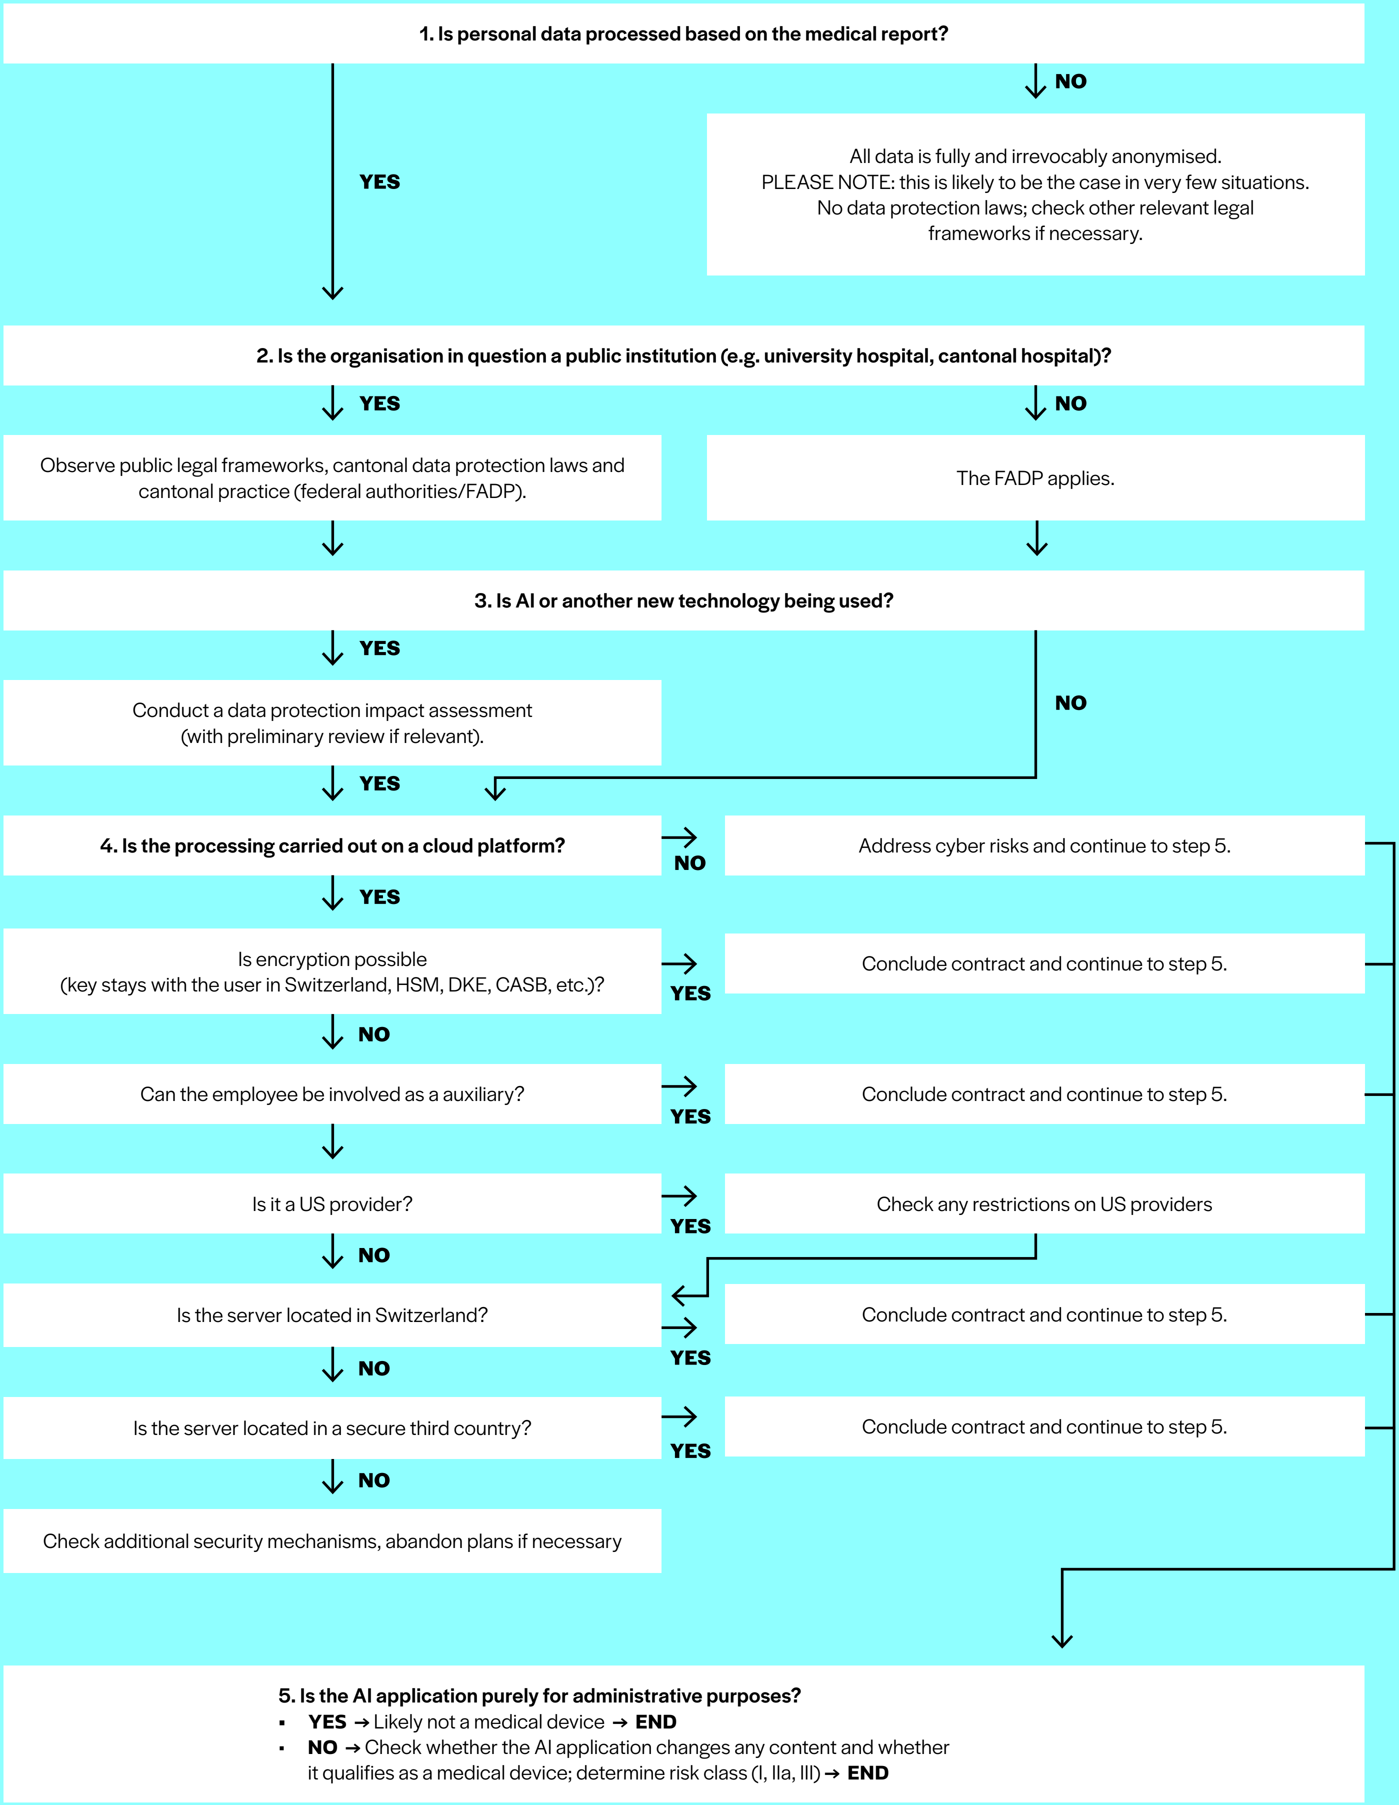

1. Is personal data processed based on the medical report?

NO

All data is fully and irrevocably anonymised.
PLEASE NOTE: this is likely to be the case in very few situations.
No data protection laws; check other relevant legal frameworks if necessary.

YES

2. Is the organisation in question a public institution (e.g. university hospital, cantonal hospital)?

YES

Observe public legal frameworks, cantonal data protection laws and cantonal practice (federal authorities/FADP).

NO

The FADP applies.

3. Is AI or another new technology being used?

YES

Conduct a data protection impact assessment (with preliminary review if relevant).

YES

NO

4. Is the processing carried out on a cloud platform?

NO

Address cyber risks and continue to step 5.

YES

Is encryption possible (key stays with the user in Switzerland, HSM, DKE, CASB, etc.)?

YES

Conclude contract and continue to step 5.

NO

Can the employee be involved as a auxiliary?

YES

Conclude contract and continue to step 5.

Is it a US provider?

YES

Check any restrictions on US providers

NO

Is the server located in Switzerland?

YES

Conclude contract and continue to step 5.

NO

Is the server located in a secure third country?

YES

Conclude contract and continue to step 5.

NO

Check additional security mechanisms, abandon plans if necessary

5. Is the AI application purely for administrative purposes?

- **YES** → Likely not a medical device → **END**
- **NO** → Check whether the AI application changes any content and whether it qualifies as a medical device; determine risk class (I, IIa, III) → **END**

rity (IDV)). The information sheet on the procedure for using AI at public institutions (see above) describes in detail how to proceed and which documents must be submitted.

What specific aspects need to be taken into consideration when using AI tools?
When using AI tools for medical documentation, particular attention must be paid to aspects of transparency, traceability and data-protection-compliant processing of personal data in the context of the cloud. The traceability of data processing, for example through technical logging or visible time stamps, should also be ensured for AI applications that are intended purely to support administrative tasks.

Furthermore, structured technical documentation of the models used is a good way to ensure transparency. ***Model cards*** can be used in this case; these cards contain information on the functionality, training data set and application limits of the LLMs (model card template from Nature). The procedures and metrics used to validate the systems in terms of accuracy, robustness and reliability must likewise be documented. However, it is important to note that such model-based solutions to increase transparency do not directly improve understanding among healthcare professionals . This would require models that make the outputs interpretable and traceable in the respective clinical context.

Another aspect relates to AI systems that continue to learn and change once they have entered into operation (continuous learning). A predefined plan for controlling and overseeing these changes over the system's entire life cycle is helpful to ensure responsible use. It must be guaranteed at all times that the quality, accuracy and safety of the application are not compromised. Continuous learning systems therefore require particularly careful monitoring and clear designation of responsibilities during operation.

Many AI systems for medical documentation are also based on cloud architectures and process large volumes of particularly sensitive personal data, such as health information. This raises questions about data security, data localisation and the possible transfer of data abroad (see 2.3 below).

2.3 Specific aspects relating to the use of cloud services

The cloud plays a key role in the use of AI solutions to prepare and edit medical reports. It allows for the scalable, flexible and secure processing of large amounts of data – provided that the participating organisations comply with data protection and organisational requirements. The following chapter provides an overview of the most important requirements when using cloud services in the context of medical documentation.

What exactly is a cloud provider?
A cloud provider provides IT resources such as storage, computing power or software on its own servers over the internet. They may be large hyperscalers such as Amazon Web Services, Microsoft Azure or Google Cloud, or smaller, specialised companies that also offer cloud services. All providers must comply with specific regulatory requirements.

Are private healthcare providers such as practices or private clinics allowed to store health data in the cloud?
When AI systems are used to prepare or edit medical documentation, e.g. for the transcription or linguistic optimisation of medical reports, data is often processed using cloud services. From a data protection perspective, this is referred to as commissioned data processing (Art. 9 FADP). The responsible healthcare institution – such as a medical

practice – transfers its data processing activities to a cloud provider, which acts as a contracted data processor. The storage or processing of personal data by the cloud provider when using corresponding AI services therefore does not constitute disclosure to a third party. Instead, the cloud provider acts as an extension of the healthcare facility, with the transmission of data falling under a 'disclosure privilege'. The healthcare institution remains responsible for ensuring the contracted data processor complies with the data protection requirements. A contract must therefore be concluded with the contracted data processor to transfer the data protection obligations.

The use of a cloud service may also be precluded by professional secrecy, which prohibits the processing of data subject to professional secrecy by third parties or prohibits the owner of the data subject to professional secrecy from 'disclosing' it to a third party. Technical measures must be taken to prevent this. These include encryption, where the cloud provider does not have a key (current solutions include **confidential computing**, e.g. in conjunction with **HSM, double key encryption or the use of a cloud access security broker (CASB))**. However, these solutions are not always feasible and sometimes entail high costs.

A cloud service could also be used if the cloud provider or the relevant employees are qualified as auxiliaries. For this to be the case, the cloud provider would have to be integrated into the healthcare provider's sphere of responsibility and functional hierarchy through contractual and organisational measures. It is a matter of interpretation as to when exactly this is the case. Any person who collaborates in the activities undertaken by a doctor bound by confidentiality and who requires plain text access for this purpose is to be considered an auxiliary – and is thus functionally understood as being part of the doctor's periphery. Cloud providers (and, through the corresponding transfer of confidentiality obligations, their employees) can become auxiliaries of cloud customers bound by professional secrecy if the employees concerned sign a confidentiality agreement that binds them to the same degree of confidentiality as the medical profession concerned. In addition, the affected employee must be integrated into the functional hierarchy of the healthcare provider. Integration of this kind is often difficult to achieve for major cloud service providers.

Are private healthcare institutions allowed to store data on an internal server (on-premises)?
The storage of data on an internal server is fundamentally possible without any issues and offers benefits in terms of an institution's control over its data. However, when implementing an internal solution, the data security requirements according to Art. 8 FADP must be strictly observed. In particular, this includes measures to protect against unauthorised access, ensure data integrity and guarantee data availability.

Can private healthcare facilities store anonymised data in the cloud?
Anonymised data is no longer considered personal data and can therefore be stored in a cloud without being subject to data protection regulations. When anonymising health data, extensive parameters have to be removed, which restricts the usability of the data (see 2.1 above). The complete removal of such information often means that the data is rendered virtually unusable for AI systems, e.g. for analysis or quality improvement.

Do private healthcare institutions need to obtain consent from the data subject for the use of cloud services?
No, the consent of the data subject is not required as long as the general processing guidelines are complied with (in particular, informing the data subject about the use of cloud services).

02. Data protection, professional secrecy and the cloud

Do special regulations apply when private healthcare providers store data in a cloud abroad?
When transferring data abroad, compliance with data transfer regulations must be ensured in addition to the provisions relating to order processing. Healthcare institutions and AI providers must therefore check whether the countries in which the data is processed offer an adequate level of data protection. This requires transparency regarding the data processing locations and the registered office of the (sub-)processor.

If the data processor or cloud provider is located in a country without a comparable level of data protection to Switzerland, or if the data is processed in such countries, the transfer cannot take place until additional measures are taken to ensure adequate data protection abroad. Countries with an appropriate level of protection are listed in Appendix 1 of the Federal Act on Data Protection (FADP). In addition to the EU member states, certain US providers are classified as secure (see below). However, in the absence of such a decision regarding a provider's adequacy, transfer abroad may still be possible if data protection is guaranteed via alternative measures such as specific or standard data protection clauses. This issue is faced predominantly in conjunction with US providers, though there are other key countries, such as India, that have insufficient data protection laws.

In the context of professional secrecy, there are concerns that transferring data abroad could have an adverse impact on data protection, as it may be easier for foreign authorities to access such data than the Swiss authorities. This is why the argument is sometimes made that foreign service providers cannot be regarded as auxiliaries. However, there is a growing perception that outsourcing is permissible, provided that appropriate safeguards are put in place to ensure confidentiality. The situation varies from country to country and can change rapidly due to political developments (see below on the US).

At the same time, there are legal requirements and enforcement measures in place that prohibit particularly sensitive data (e.g. health data) from being stored abroad. These regulations include, for example, the provisions of the FINMA Circular (2008/7 – Outsourcing – Banks), of the Federal Office of Public Health (FOPH) for national health projects and of the Ordinance on the Electronic Patient Record (Art. 12(5) EPDV).

Does this also apply if private healthcare providers store the data with a US provider?
In the past, authorities have considered collaboration with cloud providers based in the US to be particularly critical due to the CLOUD Act, which allows US authorities to access data under certain conditions. However, since the introduction of the Data Privacy Framework and the amendment of Appendix 1 of the Federal Act on Data Protection (FADP), certified US companies are also considered safe under data protection law. This means data can be transferred to these companies without the need for any additional

«The benefits of AI in documentation are significant – but only if clear rules for data protection, responsibility, and transparency are in place.» *Corinne Spirig, Chief Operating Officer, digital health center bülach (dhc)*

measures under data protection law. In the case of non-certified companies, special measures such as the conclusion of special standardised contractual clauses (EU standard contractual clauses) remain necessary.

However, when data subject to professional secrecy is transferred to the US, special security measures must be taken. It must be ensured that the provider does not have access to the data. This must be achieved through comprehensive technical measures such as encryption, with the provision that the key must remain with the responsible healthcare provider in Switzerland. Possible measures include confidential computing, for example in connection with HSM, double key encryption or the use of a cloud access security broker (CASB). Anonymised data, conversely, can be transferred without the need for any additional data protection measures (see again Section 2.1).

What precautions should be taken?
In summary, private healthcare providers must ensure that contracts with cloud providers include comprehensive measures to protect data security and comply with data protection requirements.

- **Limitation for specific purpose:** the cloud provider may only process the data in the same manner as the controller, i.e. exclusively for the contractually stipulated purposes. They must act solely in accordance with the instructions of the healthcare institution.

- **Data security:** the cloud provider must contractually agree to implement suitable technical and organisational security measures to protect the confidentiality, availability and integrity of personal data. In addition, the data processing location must be contractually agreed.

- **Duty to inform and inspection rights:** to ensure data security, the cloud provider must fulfil their duty to inform medical practices and hospitals and grant them review and inspection rights.

- **Subcontractors:** the use of potential subcontractors must be governed by contract. If the transmission of data is permissible, compliance with data protection requirements must also be governed by a contract.

- **Duty to cooperate:** the cloud provider must support medical practices and hospitals with data protection impact assessments and with requests from data subjects or data protection authorities. However, medical practices and hospitals remain responsible for the implementation of data protection rights such as the right to information, rectification and deletion.

- **Confidentiality:** the cloud provider undertakes to maintain confidentiality and must ensure this confidentiality is maintained within its sphere of influence.

Can a public hospital store its data in the cloud?

Whether and under what conditions a cantonal hospital can store data in the cloud depends on the respective cantonal laws. Cantonal bodies may also commission data processors (Section 6 IDG ZH).

Outsourcing requires a written contract containing certain minimum provisions. Section 25 IDV ZH stipulates a need for clarification regarding the subject matter and scope, the handling of personal data, confidentiality obligations, the handling of requests for access to information, information security measures, controls, sanctions, contract duration and termination of contract. In addition, the General Terms and Conditions of 24 June 2015 on the outsourcing of data processing using IT services (link) apply; these must be declared an integral component of the contract.

Outsourcing must not be in conflict with any confidentiality obligations. In addition to criminal law, professional and official secrecy are often enshrined in cantonal laws.

In the Canton of Zurich, auxiliaries may be involved if the employees concerned are integrated into the client's functional hierarchy. This is provided for in Section 3(1) of the Act on the Outsourcing of IT Services for the Cantonal Administration. It is necessary that the employees concerned are expressly designated for specific data processing, are subject to the cloud provider's right of control and instruction, and are bound by a confidentiality agreement to official and/or professional secrecy.

For public hospitals, the option remains to encrypt the data in such a way that the cloud provider does not have access to the data (see above).

What is the procedure if a public hospital outsources its data to a cloud abroad?

Public hospitals can also outsource data abroad in principle, but this increases the risks for data subjects and additional measures must be observed in line with the provisions on disclosure abroad (analogous to Section 19 IDG ZH in conjunction with Section 22 IDV ZH). Caution should be exercised when outsourcing special personal data; if possible, this should be limited to Europe.

What special circumstances apply if US providers are involved?

Due to the risks, outsourcing to US providers to which the Cloud Act applies should be avoided if special personal data is involved.

What additional precautions do public hospitals need to take?

Public hospitals must also comply with contractual measures to protect data security and comply with data protection requirements. However, in the area of public law, it is recommended that the General Terms and Conditions of the Cantonal Council of 24 June 2015 on the outsourcing of data processing using IT services be directly incorporated. These reflect the conditions required by the IDG ZH. If the General Terms and Conditions are not directly included, the content must nevertheless be reflected in the contract.

03.
Special provisions for medical devices

When generative AI (e.g. LLMs) is used to analyse and prepare medical reports especially, the distinction between purely administrative aids and medical applications is not always clear. In the case of AI software in particular that summarises, prioritises and interprets clinical information, generates relevant diagnostic statements or supports treatment recommendations, the question regularly arises as to whether it should be classified as a medical device. Due to the far-reaching consequences that this classification entails for the development, production, market approval, distribution, sale, maintenance and operation of the software in question, a thorough investigation is strongly recommended. The following section provides guidance on the legal classification and explains which laws and criteria are decisive for classifying a software solution in the field of medical reporting as a medical device.

Which laws are relevant in the field of medical devices?

In Switzerland, medical devices are subject to clearly regulated legislation that is closely aligned with the provisions of the EU. The Therapeutic Products Act is the framework law that establishes the ***legal basis for all therapeutic products***, i.e. for both ***pharmaceuticals*** and ***medical devices***. The ***Medical Devices Ordinance (MedDO)*** and the Ordinance on In Vitro Diagnostics (IvDO) are relevant for medical devices. IvDO is a specific ordinance focusing exclusively on in vitro diagnostic medical devices (IVDs). As a counterpart to MedDO, it regulates everything carried out outside a living organism in a test tube or laboratory device. However, as the name suggests, it is applicable to diagnostics (i.e. tests) and plays only a minor role in the application cases discussed here. The following information therefore focuses on MedDO.

When is something considered to be a medical device?

According to Art. 3 MedDO, medical devices are instruments, apparatus, devices, software, implants, reagents, materials or other objects that are intended for ***use on humans*** and that are oriented towards the individual. Products that are used for the benefit of a population and not an individual are excluded from this definition.

A further requirement for medical devices is that the intended main effect is not achieved through pharmacological, immunological or metabolic agents. Although the mode of action of a medical device can be supported by such agents, if the ***main effec***t is ***pharmacological, immunological or metabolic***, it is ***no longer a medical device*** but a ***medicinal product*** and the relevant regulations apply.

Medical devices (including in vitro medical devices) serve – whether alone or in combination – ***one or more specific medical purposes***, such as the diagnosis, prevention, monitoring, prediction, prognosis, treatment or alleviation of disease, injury or disability, the examination, replacement or modification of anatomy or physiological or pathological processes or conditions, or the gathering of information from the in vitro examination of samples taken from the human body.

Software can also constitute a medical device if it fulfils the above requirements, i.e. if it is intended for human use, its ***main effect*** is not ***pharmacological, immunological or metabolic***, and it has a medical purpose. Software is not a medical device if the processing of medical data is limited to storage, archiving, simple search, communication or lossless compression (MDCG 2019–11). Swissmedic has published the 'Medical Device Software' information sheet, which sets out the distinction between software as a medical device and software as a non-medical device (see information on specific medical devices).

Software that controls, oversees or evaluates medical hardware is also considered to be medical software (e.g. firmware). One particularly challenging question is whether software that displays medical images is limited to storage, archiving, communication, simple search or lossless compression. Generally speaking, software is considered to be a display-only tool rather than a medical device if its only functions are to change the brightness in order to improve the display. However, as soon as a functional change is made – such as subsequently adjusting the contrast of a medical image – this can lead to the system being classified as a medical device. The same applies if it allows length, area or volume to be measured, or if the software automatically searches for suspicious structures in the image.

When does software have a medical purpose?
Software has a medical purpose if it is intended to be used accordingly with respect to humans, for example for the diagnosis, prevention, monitoring, prediction, prognosis, treatment or alleviation of diseases, injuries or disabilities, the examination, replacement or modification of anatomy or of physiological or pathological processes or conditions. The intended purpose is to be determined by the manufacturer. This purpose must be visible on the label or in the instructions for use and must be consistently reflected in the advertising or sales material.

What rules apply if a tool is also to be used in the EU?
If the software is to be used not only in Switzerland but also in the EU, it is also subject to the EU AI Act (EU 2024/1689). Software products that are classified as Class IIa or higher under the Medical Device Regulation (MDR) or Class B or higher under the In Vitro Diagnostic Regulation (IVDR) are automatically deemed to be high-risk products under the ***EU AI Act***. In turn, they are subject to additional requirements regarding technical documentation, and are therefore subject to a conformity assessment procedure.

04. *Assessment of different use cases*

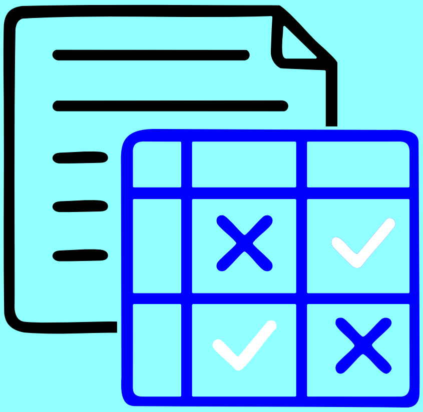

When classifying AI-based software for the preparation of medical reports as medical devices, the main question is whether the software can influence the diagnosis or the perception of doctors. The intended use is important in this context. For example, if the software evaluates medical data for the purpose of hospital statistics or reformats it to aid the reimbursement activities of health insurance funds, it does not perform any medical functions.

Even if software generally falls within the definition of a medical device, it does not qualify as such if it only performs functions such as storage, archiving, lossless compression, simple search or communication of data. These functional exceptions are defined in the applicable law. However, they cannot be unequivocally applied to AI systems such as LLMs, especially with regard to their communication functions or semantic search capabilities. There is currently no uniform interpretation by licensing bodies or supervisory authorities in this respect.

LLMs entail entirely new risks, and current regulatory concepts such as reproducibility, algorithmic transparency and generally accepted metrics for accuracy and robustness are only applicable to a limited extent. Regulatory authorities in various countries are currently developing concepts for the classification and documentation of such systems.[5] However, it will likely take a few more years before a general consensus emerges on these issues and the first guidelines and court judgements are made available.

Different AI use cases in medical documentation

The following scenarios differ in their technical implementation and regulatory classification. AI-supported solutions for medical reports can also combine more than one of the functionalities set out in the scenarios below.

[5] FDA 2024: Total Product Lifecycle Considerations for Generative AI-Enabled Devices (link).

04. Assessment of different use cases

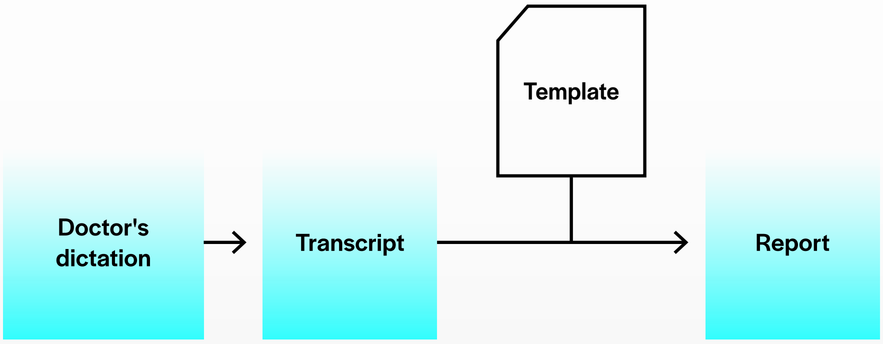

I. Transcription of a doctor's dictation

In this case, a doctor dictates a medical report according to a predefined structure. The AI-supported system converts the dictation into text using speech-to-text technology. In some cases, the transcription is automatically added to a standardised report template, which enhances the consistency of the medical documentation. The completed report can then be reviewed and approved by a medical professional. Automatic transcription significantly reduces manual effort and saves time in daily work routines, thus making this scenario particularly suitable for doctors in private practices and for hospitals that require simple support with documentation.

A key challenge is the accuracy of speech recognition, especially when dealing with specialist medical terminology. Although this information is supposed to be checked by a medical professional afterwards, initial practical experience shows that this step is sometimes omitted due to time constraints – which increases the risk of content-related errors. Review by qualified specialists – ideally medical writing services or appropriately trained medical scribes – is therefore essential. However, the benefits of such systems are clear: they increase efficiency and improve the structural quality of reports.

From a regulatory point of view, the software is usually not considered a medical device if it is used solely for transcription purposes (i.e. it converts speech into text or digitises handwritten notes) without analysing and interpreting medical content or making recommendations. In such cases, the application is viewed as equivalent to an administrative tool. However, as soon as the software amends medically relevant information, for example by highlighting certain terms or generating automated summaries, it can qualify as a medical device. In such cases, a careful examination of the intended use and the actual system functions is required.

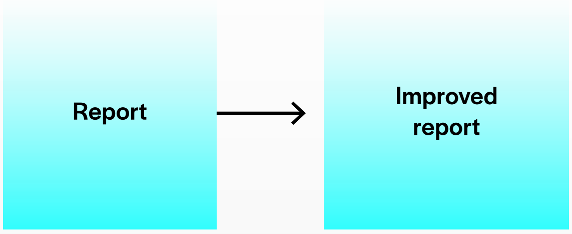

II. Optimisation of an existing report

In this case, LLMs are used to revise the language of existing medical reports or to translate them into other languages. The AI improves the wording, optimises the linguistic style and ensures clear, consistent expression without altering the technical content. In international or multilingual healthcare facilities in particular, this can significantly increase the comprehensibility and quality of medical communication. This scenario is therefore especially suitable for hospitals, international practices and research institutions that seek to achieve efficient linguistic standardisation within medical documents and translation of the same.

The key challenge is to ensure this linguistic optimisation does not result in any changes in meaning. Even if the function of the AI is limited to stylistic adjustments, it cannot be ruled out entirely that rewording may also influence the technical content. Therefore, particular care must be taken when reviewing the revised texts.

From a regulatory point of view, the software is usually not considered a medical device if it only makes linguistic or stylistic improvements or optimises layout and formatting. However, this instance should be considered as more of a borderline case than pure transcription, since linguistic adjustments could indirectly change medical statements. By contrast, if AI independently generates medical content or alters existing diagnoses and treatment suggestions, it must be classified as a medical device. Here, too, the specific purpose and the actual functioning of the software within the context of its use are decisive.

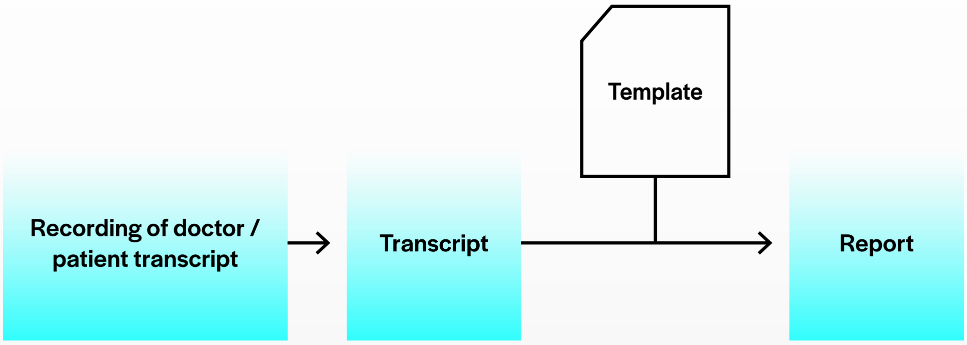

III. Preparation of a report based on patient interaction

In this case, a conversation between a healthcare professional and a patient is recorded and transcribed using a speech-to-text model and converted into a structured report template. ***Ambient clinical intelligence*** is particularly significant for longer, free-form interactions – for example, in psychiatry or GP care. The aim is to automatically capture the content of the consultation and make it usable for medical documentation purposes without the need for manual post-processing.

The use of such systems can significantly reduce the work involved in documentation and make workflows more efficient, especially at facilities with high patient throughput. The AI system must be able to filter medically relevant information from often unstructured, incomplete or contradictory conversations and convert it into a standardised form – without interpreting or changing the subject matter. However, this is precisely where a key challenge lies: in practice, there is an increased risk that LLMs will hallucinate content or assign an incorrect weighting to it. This is particularly critical when AI not only transcribes, but also prepares summaries or structured evaluations. Furthermore, AI systems for transcription capture primarily spoken content, but not clinical observations or preliminary hypotheses made by the medical professional, such as an abnormal appearance or suspected diagnoses. This supplementary information must continue to be documented manually, as this constitutes an essential basis for further medical care.

From a regulatory perspective, if the software only converts spoken language into text, it is not usually considered to fulfil a medical function – similar to the mere transcription of a doctor's dictation. However, as soon as the system begins to summarise, evaluate or even generate diagnostic information, the application becomes much more difficult to delineate. In such cases, there is a risk that medically relevant statements will be falsified, thus requiring a clear demarcation between a medical decision and the AI output. If AI is used to support decision-making or if it changes the technical content of the report, it must be classified as a medical device. The challenge here lies not only in the technical implementation, but also in the legally compliant categorisation of such hybrid functions.

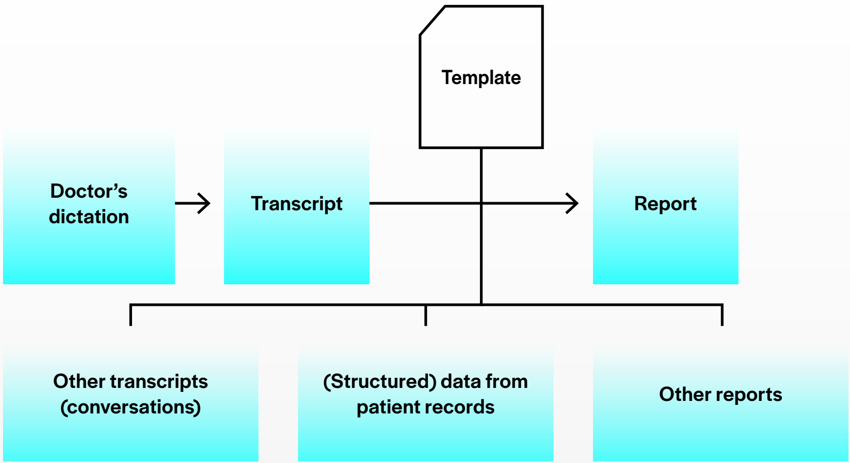
Template
Doctor's dictation
Transcript
Report
Other transcripts (conversations)
(Structured) data from patient records
Other reports

IV. *Reports from multiple data sources*

In this case, AI-supported systems combine different data sources – such as a doctor's dictation notes, electronic patient records and/or phone call recordings – and convert them into consolidated, structured medical documentation. The use of LLMs gives medical professionals access to an integrated presentation of a patient's history, which improves clarity and potentially supports diagnostic and therapeutic decisions.

The opportunities presented by these systems lie in increasing efficiency and providing better access to relevant information across different forms of documentation. At the same time, there are significant challenges with regard to the reliable collation and interpretation of content. In particular, contradictory diagnoses, the possibility of incomplete information or ambiguous wording from different sources can lead to erroneous or misleading information. It is therefore essential for AI systems to correctly assign the origin, context and meaning of the information and present it transparently so medical professionals can verify the content. Sufficient clinical experience is required to be able to make a well-founded assessment of the medical relevance and accuracy of the information generated.

From a regulatory point of view, a distinction must be made between the following: if the software only aggregates existing information, presents it in a structured form and makes its origin traceable – without medically evaluating or altering it – it is likely not a medical device. However, if the system analyses data, for example through trend analysis, or even independently draws medical conclusions, derives diagnoses or formulates treatment recommendations, it is deemed to fulfil a medical purpose. In such cases, it should be classified as a medical device, and the corresponding requirements for safety, traceability and clinical validation must be met.

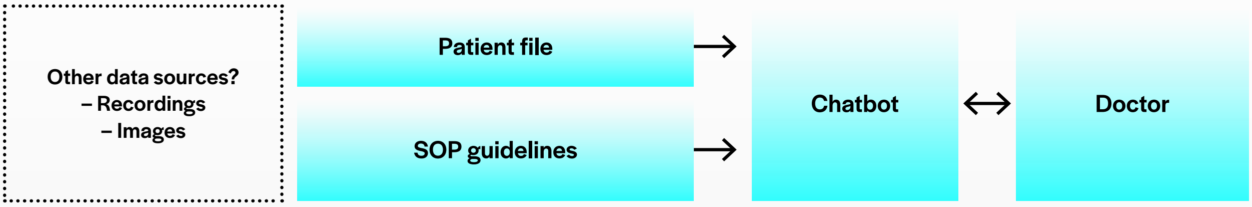

V. Chatbot for patient records

In this use case, an AI-supported chatbot enables medical professionals to retrieve specific information from a patient's file, e.g. to identify relevant findings, recognise correlations or consider differential diagnoses. The query can also be supplemented with medical guidelines and standards that support evidence-based decision-making and can contribute to the standardisation of treatment. Providing relevant information in a timely manner allows clinical decisions to be made more efficiently, which can improve the quality of care, especially in situations where time is critical.

The challenge lies in the meaningful connection and interpretation of different data sources. As with other scenarios involving multiple data sources, this may result in contradictory, incomplete or ambiguous information, thus affecting the chatbot's informative power. The more able the system is to combine complex information and formulate clinically relevant answers, the greater its benefit – but also the greater its risk, especially when it is used in situations that could compromise safety, such as to identify contraindications prior to drug administration.

From a regulatory point of view, the manner in which a system is used and the purpose it serves are the decisive factors. If the system is purely used to perform a context-sensitive search that makes existing content such as document excerpts or guidelines retrievable and reproduces the content without changing or re-evaluating it, it is usually not considered a medical device. However, if the system is used in such a way that it processes information and makes it available in an edited form for the purpose of actively supporting medical decisions (e.g. by formulating possible diagnoses or recommendations for treatment), it is a medical device. In such cases, increased requirements in terms of safety, traceability and clinical validation must also be met, depending on the area of application.

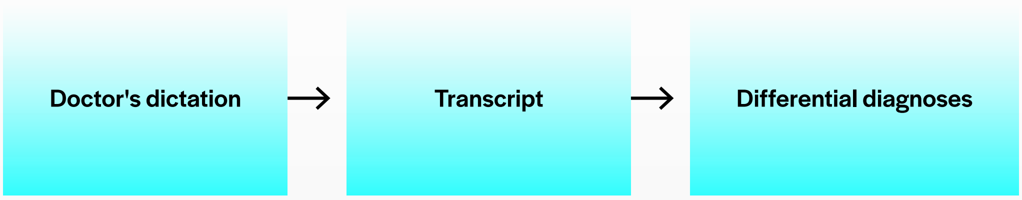

VI. *AI-based suggestions for differential diagnoses*

In this case, a doctor's dictation is first transcribed using a speech-to-text model and transferred to a structured report template. An LLM then analyses the collected information – which may include symptoms, medical history, laboratory findings or other relevant data – and generates a list of possible differential diagnoses. These will be shown to the healthcare professional for further examination, supplementation or rejection. The final decision regarding diagnosis remains expressly with the attending physician.

Such systems can be useful in supporting diagnostic decision-making by identifying relevant but not immediately obvious alternatives, thus increasing the likelihood of diagnostic certainty. They can also help with considering rare or difficult-to-identify diseases at an early stage and generate clinically sound suggestions based on structured information. In diagnostically complex specialist areas in particular – such as internal medicine or A&E – this technology offers the potential to provide informed assessments more quickly. At the same time, however, there is an increased risk of misinformation or hallucinations on the part of the language model. Therefore, it is essential that AI-generated information is clearly marked as a proposal and that oversight by a medical professional is guaranteed at all times.

From a regulatory perspective, this scenario clearly provides support regarding medical decisions. AI intervenes directly in the diagnostic decision-making process by generating and sometimes prioritising diagnoses. This unequivocally indicates that the system has a medical purpose, leading to classification as a medical device according to the definition – presumably in risk class IIa or higher, depending on the clinical field of application and potential risks in the event of malfunction. Such systems are subject to increased requirements in terms of security, transparency, traceability and clinical evaluation. The functioning of the model must be technically documented and transparent, including the data sources used, validation methods and metrics for accuracy, robustness and reliability.

Although existing regulatory frameworks such as the MDR or ISO 82304-1 are open to technology and essentially regulate products rather than technologies, their application to LLM-based systems *presents particular challenges. Due to their dynamism, lack of traceability and potential adaptability, increased requirements regarding documentation and a more complex approval review should be assumed. While the use of such systems is promising, this must be approached with special care. Clinical validation, clear demarcation between medical decisions and AI support, and transparent communication with users regarding the use of AI systems are essential for meeting legal requirements and guaranteeing patient safety. Initial scientific studies highlight the potential of these technologies in supporting differential diagnoses.*[6]

[6] McDuff et al. 2025: Towards accurate differential diagnosis with large language models (link).

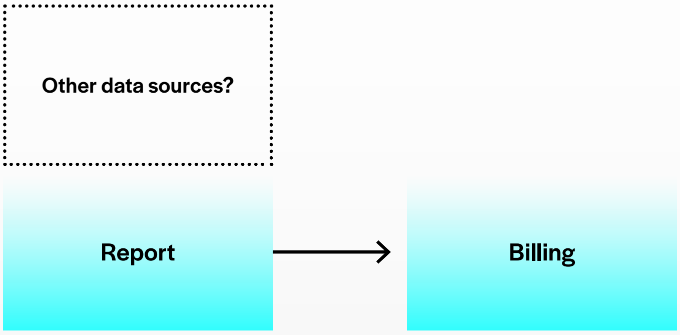

VII. Automatisierte Abrechnung

In this case, AI-supported software automatically generates a benefit statement based on a medical report – and, where relevant, other data sources, such as recordings of consultations. Larger group practices and hospitals in particular can use LLMs to standardise billing processes, thus reducing discrepancies between individual service providers. A typical example is the automated assignment of a suitable ***ICD code*** to a diagnosis in order to support consistent documentation and billing.

The main advantages are increased efficiency and the standardisation of administrative processes. At the same time, the challenge here is that even purely organisational processes can become medically relevant, especially if AI not only processes information but also evaluates and assesses it for its medical relevance.

From a regulatory perspective, software for automated billing is generally considered non-medical, as it primarily serves an organisational or financial purpose. However, if, for example, the software is used to check the medical justification of a treatment on the basis of a diagnosis or to influence treatment decisions, this may constitute a medical purpose. In such cases, careful consideration must be given to whether the use falls within the scope of the Medical Devices Ordinance. Its secondary use is particularly relevant: even if AI is originally used to generate a diagnosis for billing purposes, it can still fulfil a medical purpose, e.g. if the information is included in medical reports without medical validation, thereby influencing treatment decisions.

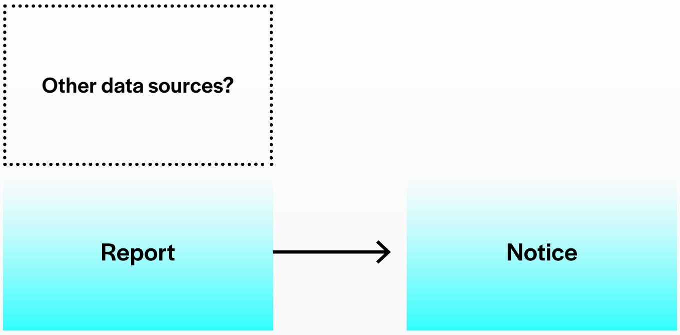

VIII. Automatic requests for patients

An automatic request is an AI-based application that automatically invites patients to attend an appointment based on predefined criteria, medical reports, patient files and schedules. AI supports appointment management and helps to facilitate the efficient use of available resources, particularly in busy practices and hospitals. The aim is to reduce the burden involved in administrative processes and optimise care planning.

The challenge is to distinguish between purely administrative processes and medically motivated decisions. If the request is generated, for example, on the basis of fixed check-ups or known treatment schemes, the functionality remains administrative in nature. The situation becomes more critical if the AI decides independently whether and when to contact a patient on the basis of evaluating medical content, e.g. by identifying an abnormality in a report. In this case, it can be seen as supporting medical decision-making.

From a regulatory point of view, a clear distinction must therefore be made between the following: if automated appointment allocation is based solely on administrative parameters, the system is not a medical device; if it is based directly on a medical analysis – for example, because the AI derives a need for action from a report – it should be assumed to fulfil a medical purpose. This may require the system to be classified as a medical device and meet the corresponding requirements in terms of safety, traceability and clinical evaluation.

04. Assessment of different use cases

The use cases above illustrate how closely the classification of software as a medical device is linked to its specific purpose and functionality. As soon as a piece of software analyses, interprets or potentially influences decision-making processes, it is subject to the requirements of the Medical Devices Ordinance. This makes technically sound implementation, transparent descriptions of functionalities and a clear separation between administrative functions and support for medical decisions all the more important – as the basis for regulatory compliance and patient safety. For applications whose classification is not entirely clear in particular, it is advised to seek support from regulatory consulting firms and to have a regulatory opinion drawn up, for instance.

What are the consequences of classification as a medical device?

As a first step, medical devices must be classified accordingly. According to ***Annex VIII MDR***, they are divided into four different classes based on the risk posed by the product, which is strongly dependent on the patient's condition (I, IIa, IIb, III). The significance of the information provided by the software is also relevant. There are three options in this respect: high (directly deciding on treatment and diagnosis), medium (significantly influencing clinical management) or low (informing clinical management only).

A more detailed list of classification rules with examples can be found in the MDCG Guidelines 2019–11 or in the EU Manual on Borderline and Classification. According to classification rule 11 of Appendix VIII MDR, only products that do not provide information that can be used for diagnostic or therapeutic purposes fall into class I. The following table from Appendix III of the MDCG 2019–11 shows how the legislature envisages the classification of diagnostic or therapeutic software.

		Significance of Information provided by the MDSW to a healthcare situation related to diagnosis/therapy		
State of Healthcare situation or patient condition		**High** Treat or diagnose ~ *IMDRF 5.1.1*	**Medium** Drives clinical management ~ *IMDRF 5.1.2*	**Low** Informs clinical management *(everything else)*
	Critical situation or patient condition ~ *IMDRF 5.2.1*	**Class III** *Category WV.i*	**Class IIb** *Category III.i*	**Class IIa** *Category II.i*
	Serious situation or patient condition ~ *IMDRF 5.2.2*	**Class IIb** Category III.ii	**Class IIa** Category II.ii	**Class IIa** Category Lii
	Non-serious situation or patient condition *(everything else)*	**Class IIa** *Category II ili*	**Class IIa** *Category Liti*	**Class IIa** *Category Li*

Table I: Classification Guidance on Rule 11

04. Assessment of different use cases

What standards must medical devices comply with?
If, according to the applicable regulations, software is classified as a product according to MedDO or IvDO, various regulations and standards must be taken into account. It is important to distinguish between ***system standards*** (requirements of the company or organisation) and ***product standards*** (requirements of the medical device itself).

Essential system standard

SN EN ISO 13485	Medical devices – quality management systems – requirements for regulatory purposes

Key product standards

SN EN 62304	*Medical device software – software life cycle processes*
SN EN 82304-1	*Health software – part 1: general requirements for product safety*
SN EN 62366 – 1	*Medical devices – part 1: application of fitness for use to medical devices*

Certification of the manufacturer according to ISO 13485 is mandatory if the product is classified as class IIa or higher (or IvDO class B or higher). However, there are many other standards, the verification of which depends on the context of the respective software. In the field of machine learning in particular, several new IEC and ISO standards are currently being planned that will also apply in Switzerland.

In view of the increasing number of cyber attacks and the growing importance of cyber security, certification of the product manufacturer according to the ISO 27000 series of standards (e.g. ISO 27001 for information security management) is recommended. This certification does not apply to the medical device itself, but rather to the manufacturer's security management – particularly for cloud-based or networked applications. Alternatively, product-related supplements can be based on the IEC 81001-5-1[7] standard, which is also recognised by European approval bodies.

Which requirements apply to AI software that qualifies as a medical device?
Any enterprise that releases a medical device on the market in Switzerland must first and foremost comply with the ***General Product Safety Regulation (GPSR)***, which can be found in Appendix I MDR. Evidence of compliance with the GPSR comprises a clinical evaluation and a risk analysis demonstrating the safety and performance of the product. In addition, ***technical documentation*** must be prepared containing the information specified in appendixes II and III of the EU-MDR or EU-IVDR.

AI providers must also carry out a ***conformity assessment***. Following the assessment, the products must bear the conformity marking.

Since the EU's mutual recognition of medical devices was revoked in 2021, it is no longer possible to have products of class IIa or higher (or IvDO class B or higher) certified in Switzerland. This requires an

[7] Health software and health IT systems safety, effectiveness and security – part 5–1: security – activities in the product life cycle.

approval body in the EU, which is known as a ***notified body***. Switzerland unilaterally recognises EU approval as a medical device.

Do special circumstances apply if a medical device has been developed by healthcare institutions and is used purely for internal purposes?
Special and facilitated requirements apply to medical devices manufactured and used in healthcare facilities (Art. 9 MedDO and 9 IvDO). In principle, software of this kind is deemed to have been put into operation, and the relevant essential security and performance requirements must be met without restriction. However, the other requirements of MedDO and IvDO do not apply. More information can be found in the MDCG 2023–1 guidelines. Healthcare institutions must give notice of the medical devices manufactured and used before they are put into operation (Art. 18 MedDO and Art. 10 IvDO).

> *«LLMs fit into existing regulatory frameworks – provided their function is clearly defined and documentation is thorough.»*
> *Dr Atanas Todarov,*
> *Chief Medical Officer, Arcondis*

Can generic LLMs be used for medical devices?
From a liability perspective, it is strongly recommended not to use generic LLMs in a medical device, as the manufacturers of such LLMs usually exclude professional health applications from their GTCs (e.g. OpenAI's usage policies). As a provider of an LLM-based healthcare solution, the institution runs the risk of becoming liable for false statements made by the LLM. However, there is now an extensive range of LLMs that have been specifically trained using medical data and are intended for medical applications (e.g. Med-PaLM 2 and BioGPT). One open-source Swiss alternative is Meditron, developed by the Swiss Federal Technology Institute of Lausanne (EPFL). Providers can build on this basic system and use additional technical and organisational measures to improve the accuracy and reliability of the LLM in the respective application context. In addition, the use of ***retrieval-augmented generation*** can help to integrate medical expertise in a controlled manner while reducing hallucinations and context-related responses.

Who is responsible for compliance with the specifications for medical devices?
The manufacturer is primarily responsible for the qualification and classification of its software as a medical device and for its conformity, safety and performance. In addition, other stakeholders throughout the life cycle have responsibilities: importers and distributors must ensure that only compliant products enter the market and comply with appropriate controls and documentation requirements. Operators such as hospitals and doctors' surgeries are obliged to ensure proper installation, use and maintenance and to report incidents. Patients have no direct obligations, but they do have a right to safe and effective products.

AI systems with continuous learning capability are particularly challenging, as their behaviour changes after they have been launched. A full ex-ante assessment of safety and performance is therefore only possible to a limited extent and requires additional control mechanisms when it comes to operation (predetermined change control plans; see Chapter 5.1, point 9).

05. *Recommendations and ideas for the future*

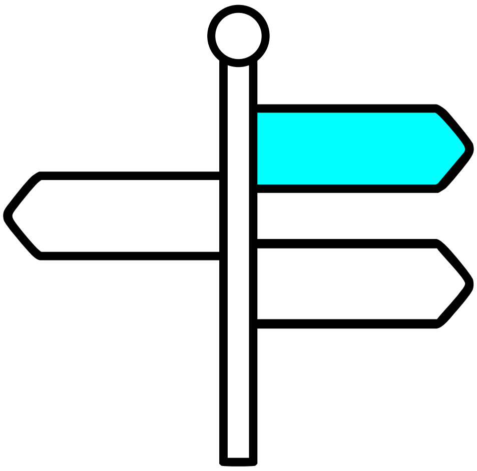

The previous chapters examined the legal prerequisites and regulatory requirements for AI in medical documentation. Building on this, the following recommendations show how such solutions can be implemented in practice in a responsible manner in compliance with data protection and the law.

5.1. Recommendations and best practices

01 ***Clear product qualification and classification***
Every AI solution in medical documentation should be formally checked at an early stage to determine whether it should be classified as a medical device (product qualification) and – if so – assigned to a risk class (product classification, e.g. class III for high-risk products). A clear definition of the intended use minimises subsequent adjustments and regulatory risks.

02 ***Early-stage collaboration with a specialist advisory service***
AI providers should work with specialist advisory services to develop a specific written proposal to classify the AI solution as a (non-)medical device. This creates regulatory clarity at an early stage and reduces the risk of having to make corrections later on or market delays.

03 ***Module-based system architecture for regulatory delimitation***
In solutions with mixed functionality, AI providers should only classify and register those modules that fulfil medically relevant purposes as medical devices. As a result, the rest of the system is subject to less extensive regulatory requirements. However, this does require the regulated modules to be clearly delineated from the non-regulated modules in the software architecture.

04 ***Compliance with the most important standards (e.g. ISO 13485)***
Before a module or system is submitted to an appointed body as a medical device, evidence of a functioning quality management system must be provided. For products that are class IIa or higher (and class B or higher for IVDs), the quality management system must even be certified according to ISO 13485. This step is essential for approval and should be integrated into project planning at an early stage.

05 ***Avoid shadow use through controlled solutions***
The unofficial use of non-certified tools (known as 'shadow use') – such as AI-based transcription or diagnostic applications – carries significant data protection, liability and quality risks.

To address this risk effectively, healthcare facilities should provide actively tested, controlled solutions that are safe and user-friendly. Transparent processes and clear internal guidelines on the use of such tools also help to avoid grey areas and ensure security and compliance.

06 ***Transparency through technical documentation***
The LLMs used must be clearly described in the technical documentation , for example, by using model cards. It must also be clearly documented how the AI was tested and which metrics were used to ensure its reliability, accuracy and robustness. Transparency alone does not mean traceability, i.e. no real opportunity for medical professionals to control or interpret the content of the system's behaviour.

07 ***Traceability and trust through human control***
Even in the case of administrative tools, medical professionals must be able to check, correct and approve content at any time. Ensuring this control function calls for clearly defined human-in-the-loop models that take into account not only usability criteria but also human factors such as excessive trust and loss of competence. For traceability, AI-generated content should be clearly labelled, editable, traceable with protocols and processing sequences, and interpretable within the specific clinical context. In this way, decision-making remains in the hands of human beings, thereby strengthening trust, quality and patient safety.

08 ***Compliance with the terms of use of LLM providers***
Companies that use generic LLMs in AI systems for medical documentation purposes should carefully check whether the respective LLM provider's terms of use explicitly allow the system to be used for medical purposes. Due to possible liability risks, it is generally preferable to use LLMs that have been specifically trained and approved for use in the healthcare sector. The use of retrieval-augmented generation should also be examined to incorporate specialist specifications such as medical terminology, jargon or institute-related guidelines.

09 ***Continuous learning software***
The US Food and Drug Administration enables medical device software that learns continuously and evolves after launch. However, this requires a predetermined change control plan. It is also important to ensure the accuracy and reliability of the software is guaranteed at all times. This kind of approach, which will likely also play a role in the context of the EU AI Act, Art. 43(4), can likewise serve as a model for AI providers in Switzerland.

10 ***End-to-end encryption instead of anonymisation***
Solutions should be designed with end-to-end encryption from the outset, as efforts geared toward pure anonymisation by removing names or birth dates do not adequately protect sensitive patient data. This also reduces technical and organisational complexity. Furthermore, encryption is an effective way to prevent cyber security risks.

11 ***Promising confidential computing***
The use of confidential computing is a forward-looking approach that ensures cloud providers cannot access patient data at any time. This ensures professional secrecy is fully maintained, including for medical reports – regardless of

whether the AI solution is used as an admin tool or as a medical device. Confidential computing offers great potential for both private practices and public health institutions.

12 ***Use of established reference architectures***
The requirements for public hospitals differ considerably from those for private practices. The focus on tried-and-tested reference architectures supports rapid, secure implementation in the context of public institutions in compliance with data protection regulations. This allows for interoperability as well as simple adjustments and updates (institutions such as the Association of Zurich Hospitals are working on various reference architectures).

13 ***Institutionalised round tables between stakeholders***
Platforms for regular dialogue with authorities, technology providers, healthcare providers and patient organisations strengthen understanding, acceptance and transparency. This results in more robust, practical solutions, especially for use cases where a large number of players are working on solutions in parallel.

5.2. Strategic considerations and ideas for the future

In addition to immediately applicable best practices, forward-looking, strategic considerations should be taken into account in order to optimally exploit the potential of AI in the healthcare sector. The following ideas address overarching challenges and offer scope for innovation, ethics and the sustainable, people-centred integration of AI.

01 Open ecosystems for medical reporting data
AI solutions for medical reports should not only be seamlessly integrated into existing hospital information systems and relevant peripheral systems but also be accessible via open, interoperable interfaces with standardised data formats. This is the only way for stakeholders – including attending physicians, nurses, therapists, pharmacies and pharmaceutical firms, researchers, insurance providers and patients themselves – to use medical and clinical information across organisations. In the future, there will be a need for open platforms that prevent vendor lock-in, facilitate competition, promote innovation and guarantee clear rules for data protection, data sovereignty and quality assurance.

02 Joint development of standards across cantons
Based on practical experience with AI in medical documentation, concrete standards can be developed to meet the data protection requirements of different cantonal data protection laws. The aim is to reduce complexity for AI providers, avoid regulatory duplication, and promote more harmonised implementation across Switzerland – ideally with the involvement of expert bodies such as the Swiss Data Protection Conference (privatim) and the cantonal health departments.

03 Ethical use of AI solutions for medical reports
Ethics in AI-generated medical reports means, in particular, transparency regarding the origin and status of the information (human versus AI-generated), the avoidance of discriminatory results through systematic bias analyses, and the assurance that doctors can take responsibility for the content and message of the report at all times. When developing and selecting AI models, attention should be paid to the diversity of training data to avoid unintended bias towards certain patient groups. Human-by-design approaches should be mandatory and take into account factors such as excessive trust and loss of competence in addition to the usual criteria such as traceability. Warnings in the event of uncertainty and targeted training that takes into account not only technical expertise but also the challenges of human-AI collaboration and implementation in complex socio-technical systems are essential. This is the only way to establish AI as a supporting tool rather than as an isolated decision-making system.

04 New value creation models for AI-generated medical reports
To be able to use AI solutions for medical reports across the board, innovative financing and business models are required that reflect the benefits for the healthcare system, e.g. in terms of efficiency gains, quality assurance, financial potential and relief for specialist staff. In the future, hybrid remuneration models could be developed that are tailored to measurable results such as saving time, reducing errors or improving the availability of information. Collaborative approaches are also conceivable, where healthcare institutions, technology providers and cost bearers jointly invest in learning systems, the value of which increases with each use – in pursuit of a sustainable, data-driven healthcare system.

05. Recommendations and ideas for the future

05 Regulatory test environments and classification aids

Regulatory test environments (e.g. sandboxes in the healthcare sector) should provide standardised templates that AI providers can use to submit their solution for legal classification at an early stage – similar to a pre-submission dossier. The template should clearly state the intended use, system boundaries and differentiation from support for medical decision-making. This creates transparency – even before a solution is launched on the market – as to whether it should be classified as a medical device. In turn, this reduces legal uncertainty, speeds up approval procedures and promotes innovation-friendly framework conditions.

06 Institutional reforms for adaptive AI regulation

Continuous learning AI systems require new regulatory approaches that go beyond traditional approval logic. Feasible examples include gradual approvals based on predetermined change control plans (see FDA model) and continuous adjustment based on dynamic risk profiles. On an institutional level, specialised evaluation units – such as multidisciplinary expert committees or sectoral AI contact points – could be set up to guide regulatory decisions on an iterative basis. Complemented by multi-stakeholder platforms, this results in a governance model that is more flexible, practical and capable of learning, and thus better suited to adaptive AI systems.

07 AI transforms medical reporting

When it comes to medical reports, the question arises regarding the extent to which the traditional medical report will retain its current structure in the future. It is conceivable that, with improved data interoperability and integration, the focus will shift from the report itself to the structured presentation and availability of basic medical or clinical information for the respective target group – be it for attending healthcare professionals, patients or health insurance companies. Conversely, a report represents a documented decision by the respective service provider and may therefore continue to play an important role in healthcare, particularly in connection with addressing liability issues.

> *«AI needs new regulatory approaches – flexible, risk-based, and beyond rigid approvals.»*
>
> *Raphael von Thiessen,*
> *AI Sandbox Programme Manager,*
> *Canton of Zurich*

Glossary

Ambient clinical intelligence (ACI)
ACI refers to AI-supported technologies that discreetly record, transcribe and structure medical conversations in the background for the purpose of automating documentation. This significantly reduces the administrative burden on medical staff while improving the quality and traceability of reports.

Anonymisation
Anonymisation means changing personal data in such a way that it can no longer be attributed to a person. This is technically challenging in medical documentation, as many information-related elements allow conclusions to be drawn. Fully anonymised data is no longer subject to data protection law, but this also often leaves AI systems less able to work with it.

Appointed body
An appointed body is an independent testing organisation – appointed and supervised by a public authority – that is responsible for the conformity assessment of certain medical devices and in vitro diagnostic medical devices. It checks whether a product meets the requirements of the relevant EU regulations (e.g. MDR or IVDR) prior to it being released on the market. AI-based medical devices require the involvement of an appointed body if they fall into a risk class requiring external evaluation.

Cloud access security broker (CASB)
A CASB is a security solution that mediates between users and cloud services and provides control mechanisms. It enables visibility, access control, encryption and threat protection when using cloud applications. In healthcare, a CASB can help with meeting compliance requirements and securely control access to medical data in the cloud.

Confidential computing
Confidential computing refers to technologies that protect data, including during processing, in specially secured hardware environments (trusted execution environments [TEE]). This keeps sensitive information inaccessible even to cloud providers or administrators. In the healthcare sector, this technology enables AI applications to be used in conjunction with patient data in compliance with data protection regulations without compromising its confidentiality.

Differential diagnosis
Differential diagnosis is a systematic procedure in medical diagnostics where doctors weigh up and compare the different possible causes of a patient's symptoms. The aim is to identify the most probable diagnosis by means of exclusion procedures and targeted examinations.

Double key encryption (DKE)
DKE is a security approach in which data is encrypted with two independent keys. One key is held by the cloud provider, while the other remains with the data owner (e.g. a hospital). This means the cloud provider cannot decrypt the data without the second key, which provides an additional layer of protection against unauthorised access. This is particularly relevant for sensitive patient data.

EU AI Act
The EU AI Act is a European Union regulation that governs artificial intelligence. It classifies AI systems into risk levels (e.g. low, high or impermissible) and defines specific requirements with regard to development, transparency, security and monitoring. AI applications in the medical sector are generally considered high-risk systems and are therefore subject to particularly stringent requirements.

Glossary

Hardware security module (HSM)

An HSM is a specialised hardware device for the secure management of cryptographic keys. It is often used in security-critical areas such as healthcare to protect key generation, storage and use against unauthorised access. In medical documentation, an HSM can be used to store or transmit sensitive data in encrypted and legally compliant form.

ICD Code

The ICD (International Classification of Diseases) Code is an international standard for the codification of diagnoses and health problems issued by the WHO. It is used for standardised documentation, billing and statistical evaluation in the healthcare sector.

Artificial intelligence (AI)

AI comprises systems that can perform tasks such as language comprehension, pattern recognition or automated decision-making. In medical documentation, AI technologies are used in particular for the transcription of voice recordings (speech-to-text), linguistic smoothing, automated structuring of content and text creation.

Large language model (LLM)

An LLM is an AI language model that has been trained with large amounts of text data to understand and generate human-like language. In medicine, it can be used to analyse, create or translate medical texts and more.

Medical report

A medical report is a structured summary of medical information about patients, usually prepared by specialists as part of diagnosis, treatment or follow-up. It serves as communication between service providers and as documentation for patients, insurance companies and authorities.

Medical device

A medical device is an instrument, a device, software or a material that is intended for medical purposes such as diagnosis, monitoring or treatment in humans and whose principal action is not pharmacological, immunological or metabolic. Regulation and approval in Switzerland are subject to the Medical Devices Ordinance (MedDO).

Medical Devices Ordinance (MedDO)

In Switzerland, MedDO regulates the marketing, monitoring and safety of medical devices. It is based on international standards and sets out requirements for the assessment of conformity, labelling and clinical evaluation. The provisions of MedDO also apply to software that fulfils medical purposes.

Model cards

Model cards are standardised documents that contain information about an AI model – such as its purpose, training data, performance metrics, and known limitations and risks. They promote transparency and help users to better understand and correctly evaluate the use and limitations of a model.

Pseudonymisation

In pseudonymisation, identifying data is replaced by identifiers; personal reference remains possible through separate correlation. Within AI applications, pseudonymisation enables processing to be compliant with data protection while maintaining traceability. Pseudonymised data is still considered to be sensitive.

Glossary

Retrieval-augmented generation (RAG)
RAG refers to a process in which a language model (e.g. an LLM) selectively accesses external sources of knowledge during text generation, such as medical guidelines, coding aids or internal guidelines at an institution. This content is retrieved in real time and integrated into the response, thereby increasing professional precision, reducing the risk of hallucinations and improving traceability. In the healthcare sector, RAG enables targeted adaptation to the specialist jargon and contexts of individual institutions.

Speech-to-text
Speech-to-text refers to the automatic conversion of spoken language into written text using speech recognition technology. In medical applications, speech-to-text models are often used to transcribe dictation, conversations or diagnostic recordings.

Vendor lock-in
Vendor lock-in refers to dependency on a specific provider, for example in the case of cloud services or AI systems. In healthcare, this can lead to limited interoperability, increased switching costs and regulatory hurdles. Therefore, when selecting technologies, attention should be paid to open standards and the possibility of a system change.

Autor

Stephanie Volz
Managing Director ITSL,
University of Zurich

Raphael von Thiessen
AI Sandbox Programme Manager,
Canton of Zurich

Case studies from the Innovation Sandbox for AI
The company MPAssist served as a case study within the Innovation Sandbox for AI. The organisation submitted a project proposal to the Sandbox in summer 2024. MPAssist offers AI solutions for medical reporting. The content of this report was developed on the basis of this specific case study.

Impressum

Publisher
Division of Business and Economic
Development, Canton of Zurich
Metropolitan Area Zurich Association
Innovation Zurich

Project conception and coordination
Raphael von Thiessen
Standortförderung Kanton Zürich
8090 Zürich
raphael.vonthiessen@vd.zh.ch

Concept in collaboration with
Stephanie Volz
Isabell Metzler
Patrick Arnecke
Markus Müller

Author
Raphael von Thiessen
Stephanie Volz

Design
here we are gmbh, here-we-are.ch

Publication
This report is published exclusively in digital format and is available in German and English.

Translation
Supertext AG

Project-Steering

- Office for Economy, Canton of Zurich
- Statistical Office, Canton of Zurich
- Department for Digital Administration, Canton of Zurich
- Office for Economy, Canton of Schwyz
- Zurich Metropolitan Area Association
- ETH AI Center
- Center for Information Technology, Society, and Law (ITSL), University of Zurich
- ZHAW entrepreneurship