

# Helse Bergen Exit Report

The use of artificial intelligence (AI) in the follow-up of vulnerable patients

# Innhold

---

- SUMMARY ..... 3**
  - Findings in brief..... 3
  - Going forward ..... 3
- ABOUT THE PROJECT ..... 5**
- ALGORITHMS AND DATA FLOW ..... 7**
  - Treatment-oriented health registry (Care Pathway Database) ..... 7
- OBJECTIVE OF THE SANDBOX PROCESS ..... 10**
- IS THE DEVELOPMENT AND USE OF AI IN THE PROVISION OF MEDICAL CARE LEGAL? ..... 11**
  - General requirement for legal basis ..... 11
  - Supplementary legal basis and principle of legality ..... 11
  - Is there a legal basis for the use of personal data for the development and use of AI in clinical practice? ..... 12
  - Prohibition against automated decision-making vs. decision-support systems..... 17
- HOW SHOULD INFORMATION BE TAILORED TO THE PATIENTS?..... 19**
  - Transparency in the development phase..... 19
  - What information must be provided in the development phase? ..... 20
  - Information relating to the right to object to personal data being used for development ... 20
  - Transparency in the application phase..... 21
  - Transparency in the continuous machine learning phase ..... 22
  - Proposal concerning medical record entries..... 22
- HOW CAN WE ENSURE THAT THE ALGORITHM PROVIDES A FAIR RESULT? ..... 25**
  - Data Protection Impact Assessment (DPIA) ..... 25
  - Built-in data protection ..... 26
  - Data quality and the requirement for data minimisation ..... 26
  - The algorithm’s accuracy ..... 27
  - Correct use of the algorithm in practice ..... 27
- GOING FORWARD ..... 28**

## Summary

---

Helse Bergen wishes to use artificial intelligence (AI) to establish an automated warning system for patients with a high probability of readmission. The warning system is intended to help clinicians identify patients who require additional follow-up to avoid being readmitted to hospital after a short period of time. The objective of the sandbox project was to clarify the legal position regarding the use of AI and to explore how patients' rights may be protected.

### Findings in brief

- **Legality:** Helse Bergen has a legal basis for the development and use of AI as a decision-support tool as part of its provision of healthcare services to patients under the EU's General Data Protection Directive (GDPR) and a supplementary legal basis pursuant to Norway's healthcare legislation.
- **Transparency:** Patients should receive general information that patient data are being used to develop AI tools. In the application phase, an entry should be made in each patient's medical records explaining the result in a manner that is intelligible for both clinicians and patients.
- **Fairness:** Distortions in any algorithm's underlying data can lead to discrimination and prevent patients from being treated equitably. To avoid this, routine quality controls of the algorithm should be performed. The warning system is a decision-support solution. To ensure human control of the decision-making process, uniform guidelines must be drawn up for how healthcare personnel should interpret and apply the results of AI tools used in patient care.

### Going forward

This exit report helps to illustrate relevant data privacy considerations and risks associated with the use of AI in clinical care of patients. The project has restricted itself to AI used for clinical decision-support. Fully autonomous AI systems, without human intervention, will require further study.

The project has concluded that, with regard to documenting its efficacy and utility, the same quality standards must be demanded of AI as of other equipment, programmes or procedures that provide decision-support in a clinical setting. Meeting the GDPR's requirements relating to information, consent and access to one's own health data lays the foundation for the development and application of new, AI-based clinical tools.



## What is the sandbox?

In the sandbox, participants and the Norwegian Data Protection Authority jointly explore issues relating to the protection of personal data in order to help ensure the service or product in question complies with the regulations and effectively safeguards individuals' data privacy.

The Norwegian Data Protection Authority offers guidance in dialogue with the participants. The conclusions drawn from the projects do not constitute binding decisions or prior approval. Participants are at liberty to decide whether to follow the advice they are given.

The sandbox is a useful method for exploring issues where there are few legal precedents, and we hope the conclusions and assessments in this report can be of assistance for others addressing similar issues.

## About the project

---

A small minority of patients account for the bulk of the resources expended in the health service. According to Helse Bergen, 10 per cent of its patients account for over half (53 per cent) of all bed days and have a high rate of readmission.

Similar conditions have also been documented elsewhere in Norway and worldwide. Despite this existing knowledge, little is being done to prevent readmissions.

[Read more about major users of somatic specialist health services in this report from Helse Bergen \(pdf\)](#) (in Norwegian).

### Readmission

Readmission is defined as being admitted to hospital for emergency treatment within 30 days of discharge from hospital.

A study has shown that, based on routine data, it may be possible to predict and identify which group of patients has the highest risk of frequent readmission. It turns out, however, that such a model works well at the group level but is less accurate at the individual level.

[The study, published in the Lancet in 2018, may be found at \[www.pubmed.gov\]\(http://www.pubmed.gov\) \(external link\)](#)

In this project, Helse Bergen wishes to explore whether it might nevertheless be possible to predict, at the individual level, which specific patients will, in all probability, be readmitted to hospital. This could provide a solid foundation for initiating preventive measures and, hopefully, result in fewer stressful days in hospital and better use of the health service's resources.

Helse Bergen wishes to test a two-step model to enable it to predict which patients will most likely be readmitted.

In stage 1, Helse Bergen wishes to establish an automated warning system for patients with a high probability of readmission. The prediction will be made automatically with the help of AI, which will base its result on previous admission records, which are included as standard data variables in Helse Bergen's electronic patient records system. The result is available the moment the patient arrives at the hospital's A&E department.

In stage 2, the objective is for hospital staff and, potentially, municipal healthcare providers, to use the automated prediction and their own informed judgement to

score the patient on the basis of their level of functionality/frailty and initiate preventive measures for those patients who need them.

## Frailty score

The frailty score is a numerical assessment of the patient's functional level based on the patient's admission history and the doctor's own judgement.

[See also "Measuring frailty - when, why and how?" on the Journal of the Norwegian Medical Association. \(external link, in Norwegian\)](#)

Helse Bergen wishes to develop a simple and user-friendly solution to identify at-risk patients based on the prediction in stage 1 and the frailty score in stage 2. The prediction will be integrated into and collated with the existing patient records/medical follow-up system.

It will be a relatively simple task to share the two-stage model to be tested and reused in other hospitals. This is because all hospitals in Norway have well-structured table data available for each patient, via data that is derived from the hospitals' electronic patient records and reported onwards to the Norwegian Patients Registry (NPR). This data contains information about previous hospital stays, with data variables linked to gender, age, date of admission, primary and secondary diagnoses, etc.

The clinical scoring means that patients can be identified as early as possible and offered a more suitable treatment pathway. This may involve interventions and follow-up by a Patient-Centred Health Team (PCHT). Such schemes, which involve the patient receiving a coordinated and holistic healthcare package delivered by a cross-functional team from the municipality and the hospital, have been successfully trialled by other hospital trusts. This has, for example, reduced the mortality rate for elderly patients and those with multiple comorbidities, while also reducing the number of emergency admissions and bed days in hospital.

[Read the article "Flere pasienter lever bedre med nytt tilbud" at \[www.forskning.no\]\(http://www.forskning.no\)](#) (in Norwegian).

The project's utility is linked to two main factors:

1. Patients may receive better care and avoid readmissions, thereby reducing the total number of days they spend in hospital.
2. Hospital trusts will be able to reduce the number of bed days for major users and, to a greater extent, divert resources to better personalised follow-up of this patient group in partnership with municipal health services. This will also facilitate follow-up care in the municipality and thereby contribute to more efficient use of society's resources.

## Algorithms and data flow

---

The algorithm is built up around the treatment-oriented health registry, the Care Pathway Database, which is described in this chapter.

(You will find a figure outlining the Care Pathway Database on this page.)

Personal health data for all patients treated by Helse Bergen are retrieved from the electronic patient records system (DIPS) and stored as a structured dataset in the same format as data variables included in the Norwegian Patient Registry (NPR). These variables are well documented and described, and are understood by all Norway's health trusts to have the same meaning. This enables other health trusts to easily reuse the same algorithm, since all health trusts have a duty to supply data to the NPR in this format by means of an NPR report.

The Care Pathway Database will give the algorithm access to a large number of episodes, with each episode constituting an admission to a department or the entire hospital stay if the patient remains in only one department. Each episode has associated data variables, such as the date of admission/discharge, gender, age, primary diagnosis, indicators of urgency (elective or emergency treatment) and care level (in-patient/out-patient). The algorithm counts the number of previous primary diagnoses, length of hospital stay, as well as previous readmissions that have taken place before each individual episode. Using this information, the algorithm learns which patients are at high risk of readmission.

When a patient is admitted to the hospital, Helse Bergen wishes to make a prediction about this patient's risk of readmission. If the risk is deemed high, the doctor will be notified by means of an entry in the patient's notes recorded in DIPS. As time goes by, the patterns underpinning the risk of readmission may change. It is therefore important to put the algorithm through subsequent rounds of training (continuous machine learning) at regular intervals, in order to update it on the basis of the new patient data.

The fully developed algorithm model scores the patient's likelihood of being readmitted to hospital in the future. A software robot cuts information generated by the algorithm and pastes it into DIPS as an entry in the patient's medical records. Doctors rely only on the information presented in DIPS, from where they will retrieve any prediction concerning a particular patient.

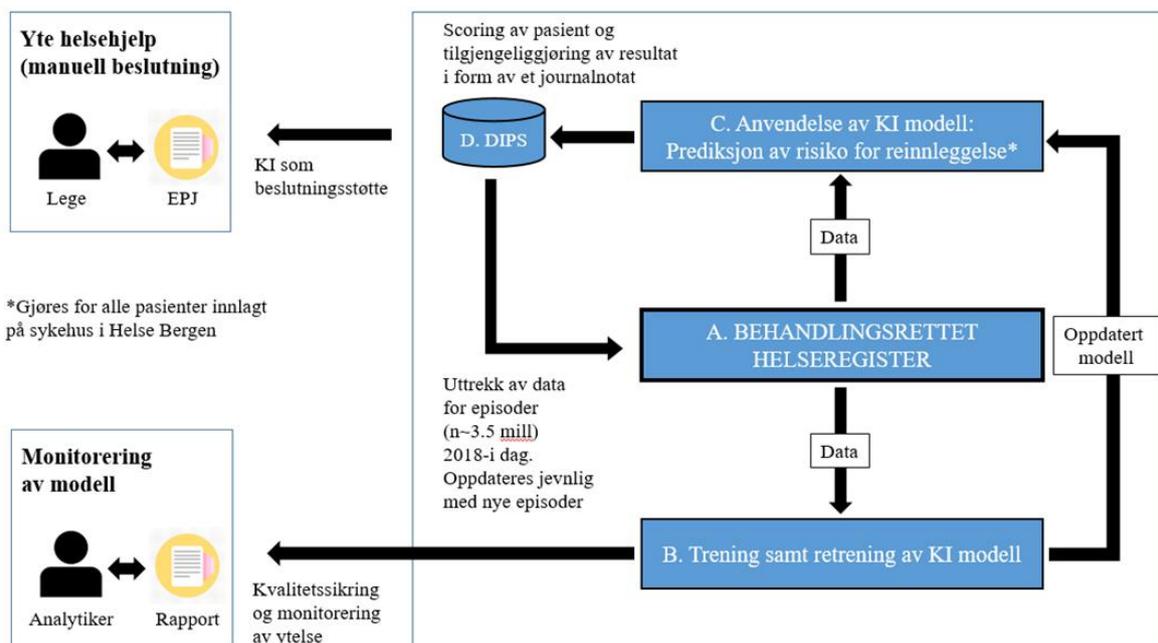
### Treatment-oriented health registry (Care Pathway Database)

The establishment of an internal treatment-oriented health registry (subsequently referred to here as the Care Pathway Database) is crucial if Helse Bergen is to be able to make use of health data derived from the medical records of a large number of patients. Such registries must be warranted in law, cf. [Section 6\(1\) of the Norwegian Medical Records Act](#), and must meet the requirements for design and organisation set out in [Section 7 of the Norwegian Medical Records Act](#).



## Section 2 of the Medical Records Act

Section 2(d) of the Medical Records Act defines a “treatment-oriented health registry” as a “patients’ medical record and information system or other registry, catalogue, etc., in which health data are stored systematically, such that data relating to the individual may be retrieved, and which are intended to provide the basis for medical assistance or the administration of medical assistance to individual people”. Section 8 of the Act imposes a duty on undertakings which provide medical assistance to ensure that they have treatment-oriented health registries.



The figure above shows the flow of data in the project:

- A. All data are stored in a treatment-oriented health registry (Care Pathway Database).
- B. The artificial intelligence (AI) algorithm is trained using data from the database.
- C. An updated AI model makes predictions about the risk of readmission for patients admitted to hospital.
- D. The readmission risk prediction is documented in DIPS, is available to the treating physician in the patient’s electronic medical records (EMR) and may be used to support decisions relating to patient care. Patients presumed to have a high risk of readmission may be offered a personalised care package. The quality of the algorithm is regularly assessed and monitored to verify its performance and utility.

Helse Bergen has created several prototypes for the categorisation of patient groups with respect to readmission. Based on NPR data from 2018 to 2020, the project has uncovered several variables specific to individuals, which correlate statistically to the risk of readmission. Provisional results show that by applying machine learning to these variables, it is possible to predict, to an accuracy of 0.75, whether a patient arriving at the hospital will experience a readmission within 30 days. This figure takes into account both sensitivity and specificity, and shows that, in around  $\frac{3}{4}$  of all cases, it will be possible to identify individual patients who will experience or will not experience a readmission within 30 days of their discharge date.

During the autumn of 2022, Helse Bergen will discuss the model's results with doctors who work on a daily basis at the Admissions Clinic at Haukeland University Hospital. In addition, Helse Bergen will work with both clinicians and patients to determine how the results should be communicated in a user-friendly fashion via patients' notes in DIPS. Efforts will also be made to explore the possibility of trialling the use of PCHTs in a pilot project in a few municipalities.

Helse Bergen considered it necessary to clarify issues relating to privacy and data protection before embarking on any trial of the model in clinical practice.

## Objective of the sandbox process

---

Helse Bergen entered the sandbox with an AI tool that was practically ready to go into operation. At that time, a thorough assessment of the legal requirements for the development and use of the tool had already been performed.

The objective of the project's participation in the sandbox was to:

- Clarify Helse Bergen's opportunities for utilising AI, where doing so is lawful and responsible.
- Shorten the path from idea to the implementation of AI in other areas of Helse Bergen in particular and the health sector in general.

In other words, the project may be beneficial and have transferable value for other projects undertaken by Helse Bergen and other health trusts, as well as other public sector entities.

The aim of the regulatory sandbox is to find good, privacy-friendly solutions for the development and use of AI. In this project, we have discussed the legality of using confidential health data for the development of algorithms, and the extent to which it is permitted to use such tools in the provision of clinical care. Furthermore, the rights of patients have taken up a large part of the agenda. In this respect, discussions have especially focused on how the right of access and the right to be forgotten may be safeguarded in practice, as well as what information about the algorithm the patient must be given in order to aid in the provision of medical assistance.

These discussions are grouped and summarised in the exit report under three main topics, which are described below. The three main topics are:

1. Is the development and use of AI in the provision of medical care legal?
2. How should information be tailored to the patients?
3. How can we ensure that the algorithm provides a fair result?

In Part 2, we will discuss issues linked to transparency and how the algorithm's results must be explained.

In the final part, we present an assessment of Helse Bergen's approach to the requirement that such a model must be fair, and how the requirements for the system to have built-in data protection may help achieve this.

## Is the development and use of AI in the provision of medical care legal?

---

The use of AI and new technologies in the health sector raises new questions related to the collection and use of vast quantities of personal data. In Helse Bergen's sandbox project, we have explored what leeway exists for the use of personal data in a clinical decision-support tool based on AI.

### General requirement for legal basis

The EU's General Data Protection Regulation (GDPR) states that the processing of personal data shall be lawful only if and to the extent that at least one of the legal bases set out in [Article 6\(1\)\(a to f\)](#) applies.



### What constitutes the processing of personal data?

In short, personal data can be defined as any information that can be linked to a natural person, directly or indirectly. It does not matter what format the information is in. Text, images, video and audio are all included. The term “processing” covers everything done with respect to the data. This includes collection, structuring, amendment and analysis.

[Article 9\(1\)](#) establishes a general prohibition against the processing of particular categories of personal data. This includes health data, because this type of personal data is considered to be of a particularly sensitive nature. Article 9(2) lists a number of exceptions to this prohibition. The exceptions listed in this provision are exhaustive.

### Supplementary legal basis and principle of legality

In some cases, Article 6 (3) and Article 9(2) of the GDPR require a supplementary legal basis in member state law. This means that the data controller must be able to demonstrate that the processing of the personal data concerned has a legal basis in both the GDPR and in national law.

## Data controller

The term “data controller” is used in health-related legislation to denominate the person deemed to be the controller pursuant to the GDPR. According to Article 4(7) of the GDPR: “‘controller’ means the natural or legal person, public authority, agency or other body which, alone or jointly with others, determines the purposes and means of the processing of personal data”.

One question is how clear the supplementary legal basis has to be. Article 6(3) may provide some guidance when the legislator is framing such a legal basis. The provision establishes that a specific statutory provision is not required for each and every processing situation as long as the purpose of the processing is laid down by national law, or the purpose is necessary for the exercise of official authority. As we shall see more clearly in the following, both Article 6(1)(c) and (e) will be relevant legal bases for the development and use of Helse Bergen's algorithm.

According to the Norwegian Data Protection Act's preparatory works, processing on the basis of Article 6(1)(c) of the GDPR must have a supplementary legal basis in which its purpose is specified. However, it is sufficient that the supplementary legal basis imposes a legal obligation on the data controller, the fulfilment of which requires the processing of personal data.

[See Prop. 56 LS \(2017-2018\) pt. 6.3.2 at \[www.regjeringen.no\]\(http://www.regjeringen.no\) \(pdf\)](#) (in Norwegian)

For Article 6(1)(e) of the GDPR, it is sufficient that the data controller needs to process the personal data for the exercise of that authority which follows from the supplementary legal basis. Thus, the supplementary legal basis does not need to expressly regulate the specific processing of personal data.

### **Is there a legal basis for the use of personal data for the development and use of AI in clinical practice?**

For the purposes of the following presentation, it is natural to address the question of legal basis in relation to the algorithm model's two main phases: (1) the development and continuous machine learning phase, and (2) the application phase, in which the algorithm model is used in clinical practice.

#### **(1) Legal basis for the development and continuous machine learning phase**

The algorithm has been developed on the basis of historic health data from almost 200,000 patients treated by Helse Bergen in the period from 2018 to 2021. This means that the health data collected about one patient is used for the purpose of

providing medical assistance to others in addition to the specific patient concerned. During the continuous machine learning phase, new health data are continuously fed into the algorithm, so that it is always updated and develops in line with any changes that arise. The processing of health data in the development and continuous machine learning phase requires a legal basis in the GDPR and national law.

Formål	Beskrivelse	Pers- data	Registrerte	GDPR Art. 6	GDPR Art. 9	Suppl. hjemmels- grunnlag
Utviklings- og etterlæringsfasen						
Trening av modell (utvikling)	Analyse, trening av modell, testing og validering		Pasienter	6(1)(e)	9(2)(h)	Hpl. § 29;
Utvikling av anonyme testdata	Utvikling av anonyme/syntetiske data til bruk for systemutvikling og testing ved behov		Pasienter	6(1)(e)	9(2)(i); 9(2)(h)	Hpl. § 26
Utarbeidelse av anonym statistikk	N/A		Pasienter	6(1)(e)	9(2)(i); 9(2)(j)	Hpl. § 26
Trening av modell (produksjon)	Kvalitetssikre modell i produksjon/feilsøk ved avvik		Pasienter	6(1)(e)	9(2)(h); 9(2)(i)	Hpl. § 29; Sphl. § 2-2; Sphl. § 3-2

**Article 6(1)(e)** of the GDPR establishes a requirement that the processing is “necessary” for the performance of a “task carried out in the public interest” or in the “exercise of official authority vested in the controller”. The purpose of developing the algorithm is to reduce the number of readmissions, improve patient follow-up and use health sector resources more efficiently.

**Article 9(2)(h)** may apply when the processing of personal data is “necessary” for the purpose of “preventive or occupational medicine” or for the “provision of health or social care or treatment or the management of health or social care systems”.

**Article 9(2)(i)** may also be used for the processing of health data when it is “necessary for reasons of public interest in the area of public health”. The term “public interest” is not defined in any more detail. It is, however, exemplified in terms of “ensuring high standards of quality and safety of health care and of medicinal products or medical devices”. It may be natural to use Article 9(2)(i) in connection with internal quality assurance on the model during its application (the continuous machine learning phase).

Article 6(3) and Article 9(2)(h) and (i) establish that the processing of health data must also have a supplementary legal basis in national law.

#### *Supplementary legal basis in healthcare legislation*

The starting point for all processing of health data by the health service is that healthcare personnel are subject to a duty of professional secrecy, cf. Section 21 of the Health Personnel Act and Section 15 of the Medical Records Act. The duty of professional secrecy is intended to protect the patient’s fundamental right to privacy and maintain public confidence in the health service. All exemptions from the duty of professional secrecy require the patient’s consent or that the exemption is warranted in law.

A typical and practical exemption from the duty of professional secrecy is the sharing of information with other healthcare personnel when it is “necessary” to provide “adequate medical assistance”, cf. Section 25 and Section 45 of the Health Personnel Act. A corresponding legal basis is required to use one patient’s health data to develop an algorithm that is intended to be used to provide medical assistance to other patients. Since it is not particularly practical to ask every single patient for their consent to waive the duty of professional secrecy, any such secondary use of health data must be covered by a statutory exemption from that duty.

Section 29 of the Health Personnel Act, which entered into force in the summer of 2021, permits the duty of professional secrecy to be waived in order for information from patients’ medical records and other treatment-oriented health registries to be made accessible. Under particular conditions, and upon application, the provision may be used as a legal basis for the development and continuous machine learning of AI-based decision-support tools, cf. **Section 29(1)(a)** of the Health Personnel Act.



## Section 29 of the Health Personnel Act

“The Ministry may, upon application, determine that the duty of confidentiality pursuant to Section 21 shall not prevent information from patients’ medical records and other treatment-oriented health registries being made available when

a.) the information shall be used for an expressly specified purpose related to statistics, health analyses, research, **the development and use of clinical decision-support tools**, quality improvements, planning, management or emergency preparedness **to promote health, prevent disease and injury or provide better medical and social care services.**”

(our emphasis)

In the Act’s preparatory works, the Ministry of Health and Social Care stated that:

“A specific assessment must be made in which the benefits to society must be weighed against the infringement of the individual’s privacy. Considerations relating to the duty of professional secrecy and the patient’s right to protection from the spread of [personal] information shall weigh heavily.”

[See Prop. 63 L \(2019-2020\), Chapter 16.1, page 129 \(pdf\)](#) (in Norwegian)

To ensure that the underlying data are always representative, the model must be constantly supplied with new health data. Such continuous machine learning may be

authorised pursuant to Section 29 of the Health Personnel Act, cf. the word “development”.

When processing health data for internal control and quality assurance, it is natural to use Section 26 of the Health Personnel Act as the legal basis and exemption from the duty of professional secrecy. In this context, quality improvement is understood to mean verification that the algorithm’s use in clinical practice is fair and reasonable.

The introduction of Section 29 of the Health Personnel Act in the summer of 2021 may be seen as a signal that the authorities wish to facilitate increased use of AI in clinical practice. The Norwegian Directorate of Health is the relevant authority for adjudicating applications for dispensation.

## (2) Legal basis in the AI-application phase

In this phase, the decision-support tool is used as part of the health service’s clinical practice. The term “clinical” implies that [the purpose of the tool being developed is to provide practical medical assistance](#). The algorithm will analyse the individual patient’s data and score the patient’s likelihood of being readmitted in the future. A software robot will then cut the score generated by the algorithm and paste it into the patient’s notes in the medical records system DIPS, where it will be visible to the treating physician.

Formål	Beskrivelse	Pers- data	Registrerte	GDPR Art. 6	GDPR Art. 9	Suppl. hjemmels- grunnlag
Automatiske beslutninger	N/A	N/A	N/A	N/A	N/A	N/A
Yte helsehjelp til pasient	Scoring av pasient mot modell, og tilgjengeliggjøre resultat for helsepersonell		Pasienter	6(1)(c)	9(2)(h)	Sphl. § 2-2; Pbrl. § 2-1b
Administrere helsehjelp til pasient	N/A	N/A	N/A	N/A	N/A	N/A
Dokumentere helsehjelp til pasient	Ivaretagelse av dokumentasjonsplikt		Pasienter	6(1)(c)	9(2)(h)	
Styring og ledelse	N/A	N/A	N/A	N/A	N/A	N/A
Intern kvalitetssikring	Kvalitetssikring av modell i produksjon		Pasienter	6(1)(e)	9(2)(i)	

By using the individual patient’s health data, the algorithm will predict the risk of the patient being readmitted to hospital in the future. Health trusts have a duty to provide adequate health and social care services, and must therefore process the data deemed necessary and relevant to provide such services. For this processing of personal data, it is relevant to examine **Article 6(1)(c) of the GDPR** in further detail. This provision requires that

“[...] processing is necessary for compliance with a legal obligation to which the controller is subject”.

Because the decision-support system processes patient health data, one of the exemptions laid down in **Article 9(2)** must be applicable for the processing to be lawful. **Article 9(2)(h)** establishes a right to process health data if the processing is “necessary” in connection with “preventive or occupational medicine” or “the provision of health or social care or treatment”.

A supplementary legal basis may be found in **Section 2-2 of the Norwegian Specialist Health Services Act**, cf. Section 2-1, which requires the specialist health service to provide adequate health services to citizens:

“The health services offered or provided under this Act must be adequate. The specialist health service shall organise its services such that the personnel who perform the services are able to comply with their statutory duties, and such that the individual patient or user is provided with a holistic and coordinated suite of services.”

The duty to provide an adequate service is a legal standard in healthcare legislation. A legal standard is dynamic and will develop in line with developments in society and new standards for medical assistance. This creates a higher degree of unpredictability for the patient. However, access to and the processing of personal data are a prerequisite for and a natural part of, the duty to provide adequate medical assistance. According to the Health Personnel Act’s preparatory works, when the requirement for providing adequate services is discussed, the substance of the duty to provide adequate services must be assessed on the basis of legitimate expectations, the health personnel's qualifications, the nature of the work and the situation in general.

[See Section 4.2.5.3 of Proposition no 13 to the Odelsting \(1998-99\) at www.regjeringen.no](http://www.regjeringen.no) (in Norwegian)

The duty of health personnel to provide an adequate service applies irrespective of the patient's voluntary will or capacity to exercise their own autonomy. The requirement to provide adequate care may be understood as a duty to develop and offer health services based on new knowledge and technologies, including the use of AI where this is expected to be beneficial.

**Section 39 of the Health Personnel Act** establishes that the person providing medical assistance has a duty to document the care provided. This duty implies an individual duty to record information about the patient and the medical assistance provided, which is *relevant* and *necessary*, see also Section 40 of the Medical Records Act and the Medical Records Regulations. The duty of documentation is grounded in the duty to provide adequate care.

A specific legal basis must exist if data are to be processed over and above that which is deemed necessary and relevant for the provision of medical and social care services and the health personnel’s duty to document the same in the specific case concerned.

Furthermore, **Section 19 of the Medical Records Act** establishes that healthcare providers have a duty to:

“... ensure that relevant and necessary health data are available for healthcare personnel and other cooperating personnel when this is necessary to provide, administer or verify the quality of the medical assistance provided to the individual.”

This provision may be used as a legal basis as long as the processing of the data is **necessary** and **relevant** for the provision of the service.

The right and duty to share patient data with cooperating personnel is also laid down in Section 25 and Section 45 of the Health Personnel Act.

In summary, there are several provisions in the legislation regulating the provision of health services that permit personal data to be used for patient care (particularly Section 2-2 of the Specialist Health Service Act and Section 19 of the Medical Records Act). A fundamental precondition for health personnel being able to provide adequate medical assistance (cf. Section 2-2 of the Specialist Health Service Act) is access to relevant and necessary patient data.

## **Prohibition against automated decision-making vs. decision-support systems**

Article 22 of the GDPR prohibits decisions based solely on automated data processing. For such a prohibition to be applicable, the decision must be made without human intervention. In addition, the decision must produce legal effects concerning the data subject or have a similarly significant impact on them. The requirement for human intervention implies that the decision-making process must include an element of real and actual assessment performed by a human being.

The algorithm model in Helse Bergen's project is limited to being decision-support tool, which will only be used as a guide to health personnel in their assessment of patient follow up. [The Health Personnel Act's preparatory works](#) make it clear that the term “decision-support tool” shall be broadly understood and that it encompasses all types of knowledge-based aids and support systems, which may provide advice and support, and may guide healthcare personnel in the provision of medical assistance. This includes the development and use of systems built on artificial numerical analysis and systems built on machine learning. The algorithm's recommendation will thus form only one of many factors that determine the measures to be implemented.

It must nevertheless be emphasised that use of the algorithm in clinical practice is based on a profiling of the patient, cf. Article 4(4) of the GDPR. Profiling is defined as “any form of automated processing of personal data consisting of the use of personal data to evaluate certain personal aspects relating to a natural person, in particular to analyse or predict aspects concerning that natural person's performance at work, economic situation, health, personal preferences, interests, reliability, behaviour, location or movements”. Even though profiling alone does not trigger the prohibition

laid down in Article 22, profiling may indicate that there is a high risk associated with such processing. This is particularly true of profiling that incorporates health data.

Although the algorithm is used as a decision-support system, there is a risk that healthcare personnel may rely wholly on the algorithm's result, without performing any independent assessment of their own. In that case, use of the AI tool will, in practice, be fully automated. Measures to reduce the risk of this occurring include ensuring that healthcare personnel receive sufficient information and training before use of the tool is adopted, presenting the results in terms of probable outcomes rather than categorical outcomes, being transparent with regard to the algorithm's accuracy and introducing procedures and control mechanisms to uncover errors in the model.

## How should information be tailored to the patients?

---

Transparency is a fundamental principle in the GDPR. Transparency about the processing of personal data may help to uncover errors and unfair discrimination. It contributes to trust and enables the individual to uphold their rights and protect their own interests.

Providing the patient with information is a fundamental precondition for effective medical assistance. In many cases, the requirements set out in data protection and health-related legislation will conform and complement each other. Nevertheless, there are some differences that we will illuminate here. When using a clinical decision-support system, the clinician will play a key role in providing relevant information to the patient.

### Transparency in the development phase

Article 13 and Article 14 of the GDPR require the data controller to inform the data subject when their personal data are collected and used in connection with the development of algorithms. Article 13 applies to data **collected from the data subjects themselves**, while Article 14 regulates cases where the data are **collected from other sources**, such as third parties or publicly available data.

In this project, the personal data will be collected from **sources other than** the data subject. The data are obtained from the patient records system and have been generated by the hospital through diagnoses, registration of admissions, etc. The information that the patient must receive is therefore regulated by Article 14.

The data used concern person-specific admission records from almost 400,000 patients. From these, the following data are extracted for training purposes:

- gender
- age
- no. of bed days
- previous readmissions
- primary and secondary diagnoses
- length of hospital stay

Obtaining this information individually can be a resource-intensive process.

Article 14(5)(a to c) of the GDPR establishes certain exemptions from the general rule. An exemption may be made if the data subject already has the information, the provision of such information proves impossible or would involve a disproportionate effort. The same applies if obtaining or disclosing the data is expressly laid down by EU or member state law to which the controller is subject and which provides appropriate measures to protect the data subject's legitimate interests.

It is worth mentioning here that Section 29 of the Health Personnel Act establishes that a duty to submit an application to the Norwegian Directorate of Health,

requirements relating to data minimisation and a duty of confidentiality are appropriate measures.

In this project, Article 14(5)(b and c) may be applicable during the development phase.

However, it must always be clear to the individual data subject that their data are being used for the development of AI, cf. Article 12(1) and the general principle of transparency laid down in Article 5(1)(a) of the GDPR. This may be achieved through the publication of general information concerning data processing at publicly available areas, such as the entity/data controller's website. This is especially important in order to build trust when new technologies, such as AI, are used. In keeping with the transparency principle, Helse Bergen will, when the algorithm is used, present the underlying data, the result and any uncertainties relating to the individual patient in an entry in that patient's medical records, which the patient can view at [www.helsenorge.no](http://www.helsenorge.no).

## **What information must be provided in the development phase?**

Patients should receive general information that patient data are being used to develop AI tools. It is therefore important that this information is available to the data subject before any further processing takes place.

When personal data are processed in connection with the development of AI systems, it is particularly important that information on the following is provided:

- Which types of personal data are processed.
- The purpose for which the algorithm is being developed.
- What happens to the data once the development phase has finished.
- Where the data are obtained from.
- The extent to which the AI model processes personal data and whether measures have been initiated to anonymise the data.

## **Information relating to the right to object to personal data being used for development**

AI is a new technology and there may be many reasons for why patients do not want their health data to be used for the development of AI tools. Patients should be informed in general terms that they have the right to object to their health data being used for the development of AI tools. This information may, for example, also be disclosed on the entity/data controller's website.

The right to object to one's own health data being used in an AI tool is enshrined in Section 7 of the Medical Records Act, cf. Section 17, and Section 5-3 of the Patients' Rights Act.

The right to object under Article 21 of the GDPR does not apply to treatment-oriented health registries authorised pursuant to Article 6(1)(c).

## Transparency in the application phase

In **the application phase**, the information that must be supplied will depend on whether the AI model is being used for **decision support** or for **fully automated decision-making**.

For automated decisions which have a legal effect or significantly affect a person, specific information requirements apply. If processing can be categorised as automated decision-making or profiling pursuant to Article 22, there are additional requirements for transparency, cf. Article 13(2)(f) and Article 14(2)(g). This includes:

- Information that the data subject is the subject of an automated decision.
- Information that they have the right not to be subject to an automated decision under Article 22.
- **Meaningful information** about the AI system's underlying logic.
- The significance and expected consequences of being subject to an automated decision.

Although it is not an explicit requirement that the data subject be provided with supplementary information when the AI system is used as a decision-support tool, the Norwegian Data Protection Authority **recommends** that supplementary information be provided in such cases, particularly when health data are used for the purpose of profiling. This is supported by Article 14(2)(g), together with Article 12(1) and Recital 60. It is also in line with the [guidelines published by the European Data Protection Board](#) (EDPB), see page 25, which also provides guidance on what a meaningful explanation of the logic may contain.

A meaningful explanation will depend not only on technical and legal requirements, but also linguistic and design-related considerations. The intended target group for the explanation must be assessed. In this case, the target groups will be healthcare personnel and patients.

In many cases, a complex mathematical explanation of how the algorithm and machine learning works will not be the right approach to take. Rather, the data controller should focus on the information being clear and intelligible for the data subject. For example, the controller may disclose:

- The categories of data that have been or will be used in the profiling and decision-making process.
- Why these categories are considered relevant.
- How a profile used in the automated decision-making process is constructed, including any statistics used in the analysis.
- How it is used in a decision that affects the data subject.

Such information will generally be more relevant for the data subject and contribute to greater processing transparency. It may also be useful to consider visualisation and interactive techniques to augment transparency about how the algorithm works.

This information will also be useful for healthcare personnel, helping to build their confidence in the system and ensuring a genuine evaluation of its output.

Other requirements concerning the disclosure of information about the reasons for automated decisions may also apply to public sector entities. Such requirements may, for example, be found in the Public Administration Act or sector-specific legislation.

## Transparency in the continuous machine learning phase

In connection with the use of personal data collected during the application phase for **continuous machine learning**, the requirement to provide information will largely be the same as the requirements in the initial development and application phases.

In this solution, the proposal is for periodic rounds of machine learning, in which case the requirements will be the same as for the development phase.

### What information requirements are laid down in the healthcare legislation?

Healthcare personnel have a statutory duty to provide patients with information, and patients have a right to receive information that can fulfil the transparency requirement.

Section 4-1 of the Patients' Rights Act establishes that consent to medical assistance is valid only if the patient has received the necessary information about their own medical condition and what the medical assistance entails. The same may be deduced from Section 4 of the Health Personnel Act concerning the requirement to provide adequate and caring medical assistance.

Furthermore, it follows from Section 3-2 of the Patients' Rights Act that the patient must have the information necessary to obtain an understanding of their own medical condition and what the medical assistance being provided entails. The patient must also be informed about potential risks and side-effects.

These duties mean that information must be provided to the extent necessary for the patient to understand what the medical assistance provided entails.

Our assessment is that the decision-support tool which Helse Bergen has developed does not represent a high risk to the data subject's rights and freedoms, and must be treated in the same way as other technical and medical aids (blood pressure monitor, X-ray machine, blood tests, etc.) with regard to the information that must be provided. Information will be provided and assessed by healthcare personnel, who will also make the final decision.

## Proposal concerning medical record entries

A meaningful explanation will depend not only on technical and legal requirements, but also linguistic and design-related considerations. An evaluation must also be performed of which target groups the explanation is aimed at, which may mean different explanations being given to healthcare personnel and patients. If the

explanations appear to be standardised, it could undermine their significance for the decision-making process. Social factors such as trust in the entity, the significance of the score and trust in AI systems in general may also influence whether an explanation is perceived as being meaningful.

**Information given to patients by healthcare personnel must therefore be personalised. This may also be relevant for medical record entries in order to communicate the information in an intelligible fashion. Here are some examples of how medical record entries may be designed:**

Journalnotat for vurdering av risiko for reinnleggelse

**Hovedresultat: Pasient Kari Normann er blitt vurdert til å ha høy risiko for reinnleggelse.**

Kunstig intelligens har vært brukt for å vurdere risikoen for reinnleggelse av pasient Kari Normann. Følgende sykdomshistorikk har vært benyttet i vurderingen:

inndato	utdato	hoveddiagnose	antall liggedøgn akkumulert	reinnlagt
2023-05-07	2023-05-15	IXXX	9	Nei
2023-05-07	2023-05-07	IXXX	10	Ja
2023-07-01	2023-07-07	IXXX	17	Nei
2023-07-25	2023-07-26	IXXX	19	Ja
2023-07-26	2023-07-31	JXX	25	Nei
2023-09-12	2023-09-16	IXXX	30	Nei
2023-09-21	2023-09-27	NXXX	37	Ja
2023-10-01	2023-10-07	NXXX	44	Nei
2023-11-03	2023-11-03	TXXX	45	Ja
2023-11-03	2023-11-10	NXXX	53	Nei

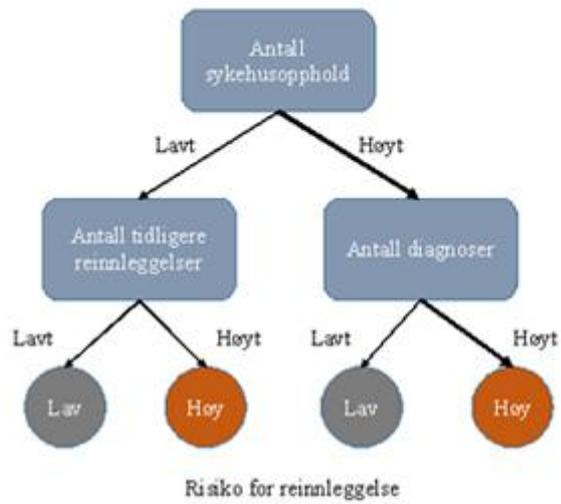
*Tabell 1 Oversikt over tidligere avdelingsopphold for pasienten.*

En reinnleggelse er definert som en øyeblikkelig hjelp innleggelse innen 30 dager etter en tidligere innleggelse. Følgende opplysninger fra pasientens journal i Helse Bergen brukes for å vurdere om pasienten har høy risiko for reinnleggelse:

- ✓ Antall liggedøgn på sykehus
- ✓ Alder per dags dato
- ✓ Antall sykehusopphold
- ✓ Antall hoveddiagnoser
- ✓ Antall bidiagnoser
- ✓ Antall tidligere reinnleggelser

Systemet rapporterer et hovedresultat basert på diagrammet som er angitt i figuren under. Variabelen i hver boks blir scoret som lav eller høy, og angir veien til et hovedresultat.

Resultatet for Kari Nordmann er fremkommet basert på skåringene som er brukt og angitt av de tykke pilene gjennom diagrammet.



Figur 1 Beslutningsdiagram for å vurdere risikoen for reinnleggelse.

Mer informasjon om algoritmen og resultater av klinisk utprøving finnes her ([Her settes inn en link til HBE sine nettsider med beskrivelse av treffsikkerhet til algoritmen + vitenskapelig studie som vi tenker å gjennomføre](#))

Algoritme: *Reinnleggelse\_risiko*

Dato: 01.03.2025

Sted: Haukeland Universitetssykehus

## How can we ensure that the algorithm provides a fair result?

---

When discussing fairness in this sandbox project, our starting point has been three main principles for responsible AI: it must be lawful, ethical and robust.

These main principles are based on the [“Ethical Guidelines for Trustworthy AI”](#), prepared by an expert group appointed by the European Commission. The same principles are also reflected in the Norwegian government’s [National Strategy for Artificial Intelligence](#) from 2020.

A good starting point for upholding these principles is the performance of a Data Protection Impact Assessment (DPIA). A DPIA is a process intended to describe the processing of personal data and assess whether or not it is necessary and proportional. It will also help to manage the risks that such processing poses to the individual’s rights and freedoms, by assessing these and determining risk-reducing measures.

### Data Protection Impact Assessment (DPIA)

If it is probable that a type of processing of personal data will entail a high risk to people’s rights and freedoms, the controller must assess the planned processing activity’s impact on privacy. The Norwegian Data Protection Authority has drawn up a list of processing activities that always trigger the need to perform a DPIA. The list is available on the Data Protection Authority’s website. Among other things, it states that the processing of personal data by means of innovative technologies, such as AI, as well as the processing of special categories of personal data (such as health data) for the purpose of algorithmic training, require a DPIA to be performed.

In its DPIA, Helse Bergen identified several different risk factors. Two of these in particular indicated a risk that the solution may not adequately fulfil the requirement for fairness:

1. False negative and false positive results

The risk that the algorithm predicts so-called false negative or false positive results. This means that someone who should have been given additional follow-up does not receive it, or that a person who does not need additional follow-up is offered it. In the first instance, the patient will be offered a follow-up that accords with current practice, which cannot be said to be associated with a high risk. The second instance represents no risk to the patient, but is undesirable from the perspective of the hospital’s use of resources.

2. Demographic distortion

The risk that the algorithm discriminates against certain groups in society. If the algorithm prioritises certain groups, other groups will feel they are being marginalised or subjected to unfair and inequitable treatment.

Below, examples from the DPIA discussion will be used to illustrate how fairness and data protection have been built into the algorithm.

## **Built-in data protection**

In its guidelines, the European Data Protection Board (EDPB) states that built-in data protection is one of several factors included in the fairness principle, alongside respect for the data subject's rights and freedoms, such as freedom from discrimination, the data subject's expectations and any broader ethical implications of the data processing. Built-in data protection has been a recurring theme in this sandbox project, both in discussions relating to discrimination and distortion in the algorithm and in discussions relating to how patients' rights and freedoms under the data protection regulations may be upheld.

[Read more about built-in data protection in the EDPB's guidelines.](#)

Article 25 of the GDPR establishes a duty to ensure effective data protection in the development of solutions or technical systems through the implementation of technical and organisational initiatives – in other words, a duty to ensure built-in data protection. On its website, the Norwegian Data Protection Authority underlines that the requirement for built-in data protection must be met before personal data are processed, and that mechanisms to ensure built-in data protection must be maintained for as long as the processing takes place.

Helse Bergen has itself developed the algorithm used to predict the risk of readmission. Helse Bergen has therefore had plenty of opportunities to plan necessary measures, such as data minimisation, pseudonymisation and data security initiatives, to meet the requirements for built-in data protection from the outset. Below, we present some examples of relevant measures in this project.

## **Data quality and the requirement for data minimisation**

The Norwegian National Strategy for AI highlights distortions in the underlying data as a particular obstacle to inclusion and equitable treatment. This is explained as follows: [“datasets used to train AI systems may contain historic distortions, be incomplete or incorrect](#). Poor data quality and errors will embed themselves in the algorithm and may lead to incorrect and discriminatory results.

One way of avoiding distortion in the selection (bias) is to ensure that the underlying data are adequate and relevant for the algorithm's predefined purpose. The data minimisation principle states that personal data may be lawfully processed only to the extent necessary to fulfil the intended purpose. The data minimisation principle restricts the use of large quantities of personal data in the development of an algorithm if the purpose may be achieved using a smaller dataset.

The purpose of Helse Bergen's algorithm is to quantify the risk that a patient may be readmitted to hospital in the future. As mentioned earlier, Helse Bergen planned to use only historic data that has proved to have a statistical correlation with the risk of readmission. The underlying data comprises the previous admission records of a large number of patients.

The project quickly found that a few key parameters made the algorithm's predictions as accurate as the use of a large number of parameters with little relevance for readmission. The data variables currently being used include previous admissions, number of bed days, gender, age, indicators of urgency and the number of primary and secondary diagnoses. Information about the patient's diagnoses is listed only as a number and is not specified by type. In addition, Helse Bergen decided that the algorithm should be used only in connection with those patient groups where there was a risk of frequent readmission, with the focus on the patients who had recently been admitted to hospital.

The algorithm will be trained continuously through the input of new data from the patient records system DIPS. To ensure that the Care Pathway Database is updated at the time the algorithm makes its predictions, the algorithm must be run frequently. In this way, any changes in DIPS will be included in the basis for the decision.

## **The algorithm's accuracy**

There will always be a risk of the algorithm predicting so-called false negative or false positive results. This is an issue to be found in the vast majority of algorithms which use AI. In this case, it means that someone who should have been given additional follow-up does not receive it, or that a person who does not need additional follow-up is offered it. In the first instance, the patient will be offered a follow-up that accords with current practice, which cannot be said to be associated with a high risk. The second instance represents no risk to the patient, but is undesirable from the perspective of the hospital's optimal use of resources.

Because the underlying data change over time, there is a risk that the algorithm's accuracy will also change. To uncover any reduction in accuracy or the emergence of distortions in the data on which the predictions are based, Helse Bergen plans to perform routine quality assurance. In the longer term, Helse Bergen will be able to extract historic data about which patients were readmitted and document the extent to which the algorithm succeeded in identifying them in advance.

## **Correct use of the algorithm in practice**

The algorithm developed by Helse Bergen is intended to be used as a decision-support tool and will be supplemented by the healthcare professionals' own expert assessments. However, there will always be a risk that the algorithm's output may be used uncritically and function in practice as a fully automated system. A decision without real human intervention will, in practice, come under the prohibition against fully automated decisions established in Article 22 of the GDPR.

To ensure human intervention in the decision, Helse Bergen plans to draw up a set of uniform guidelines after involving user committees. Questions to be clarified include how the result should be interpreted and the weight that should be attached to the recommendation. Increased awareness of the algorithm's accuracy on the part of healthcare personnel will make it easier to uncover any errors or distortions that may arise, and then make the necessary adjustments.

## Going forward

---

Although artificial intelligence has long been used in the field of biomedical and clinical research, it has not been used to any great extent in the clinical care of patients. In this project, we have attempted to examine relevant privacy considerations and risks arising from the use of AI in the provision of patient care.

We have focused on AI used for clinical decision support and take the view that fully autonomous AI systems, without human intervention, would require further study.

One of our main conclusions is that AI used for decision support in a clinical setting does not, in this case, differ significantly from the current use of existing procedures or medical-technical equipment. However, the same quality requirements must be demanded of AI systems as of other procedures or medical-technical equipment. As long as the requirements for data protection, information, consent, access to one's own health data and right of co-determination in one's own treatment are met, new clinical tools based on AI and that adequately protect the patient's data, may be introduced. Such tools may be algorithms implemented via third-party solutions or by means of the hospital's own resources.

The use of AI in the health service has considerable potential, and this project has illuminated only a single application in which AI may be of benefit. Use of AI can be beneficial for clinical decision-support, resource planning, reporting or preventing unwanted incidents. AI-based systems may be more or less invisible algorithms that monitor the hospital's operation, giving healthcare personnel or various units access to information that is of interest as and when they need it, and that can be used actively for the provision of better patient care. AI's ability to provide decision-support and structure complex data variables may help to further improve health services, and potentially reduce the pressure on staff.



**Besøksadresse:**  
Trelastgata 3, Oslo

**Postadresse:**  
Postboks 458 Sentrum,  
0105 Oslo

postkasse@datatilsynet.no  
Telefon: +47 22 39 69 00

[datatilsynet.no](https://www.datatilsynet.no)  
[personvernbloggen.no](https://www.personvernbloggen.no)