

PRECISE4Q



PREDICTIVE MODELLING IN STROKE

DELIVERABLE

Project Acronym: **Precise4Q**

Grant Agreement number: **777107**


Project Title: **Personalised Medicine by Predictive Modelling in Stroke for better Quality of Life**

D1.5 – Empirical study on attitudes towards personalized medicine

Revision: 0.1

Authors and Contributors	Julia Amann (ETH Zurich) Alessandro Blasimme (ETH Zurich) Effy Vayena (ETH Zurich)		
Responsible Author	Julia Amann	Email	julia.amann@hest.ethz.ch
	Beneficiary	ETH Zurich	Phone +41 44 632 46 19

Project co-funded by the European Commission within H2020-SC1-2016-2017/SC1-PM-17-2017	
Dissemination Level	
PU	Public, fully open
CO	Confidential, restricted under conditions set out in Model Grant Agreement
CI	Classified, information as referred to in Commission Decision 2001/844/EC



Revision History, Status, Abstract, Keywords, Statement of Originality

Revision History

Revision	Date	Author	Organisation	Description
0.1	4.11.19	J. Amann	ETH Zurich	Preliminary outline
0.2	15.11.19	J. Amann	ETH Zurich	First draft
0.3	10.12.19	J. Amann	ETH Zurich	Feedback A. Blasimme
0.4	28.12.19	J. Amann	ETH Zurich	Review & Feedback E. Vayena, A. Blasimme
1.0	29.12.19	J. Amann	ETH Zurich	Revision and Final Submission

Date of delivery	Contractual:	31.12.2019	Actual:	31.12.2019
Status	final <input checked="" type="checkbox"/> /draft <input type="checkbox"/>			

Abstract (for dissemination)	<p>This document describes the activities conducted in relation to T1.5. These activities contribute to achieving objective 5-7 laid out in WP1 (Patients' Needs and Ethical Framework).</p> <p>More specifically, this deliverable presents the rationale, design, and findings of a qualitative research study aimed at investigating different stakeholders' expectation and attitudes towards predictive modelling in stroke. The work presented here is instrumental to integrating the perceived ethical priorities of relevant stakeholders in the ethical framework that will be developed within WP1. The document concludes by highlighting some of the challenges and limitations and describes next steps.</p>
Keywords	Ethics, AI, patient autonomy, justice, data protection and privacy, disclosure, responsibility, empathy

Statement of originality

This deliverable contains original unpublished work except where clearly indicated otherwise. Acknowledgement of previously published material and of the work of others has been made through appropriate citation, quotation or both.




Table of Content

Revision History, Status, Abstract, Keywords, Statement of Originality	2
Executive Summary	4
1 Overall objective and scope of the deliverable.....	5
2 Study rationale and background.....	6
3 Methodology.....	7
3.1 Participants	7
3.2 Recruitment	7
3.3 Ethics.....	8
3.4 Data management	8
3.5 Withdrawal of participants	8
3.6 Data collection tools	8
3.7 Data analysis	10
4 Findings.....	11
4.1 Participant characteristics.....	11
4.2 Making sense of AI	11
4.3 Ethically relevant concerns related to the Precise4q scenario.....	12
4.3.1 Patient autonomy and self-determination	12
4.3.2 Disclosure.....	12
4.3.3 Responsibility.....	13
4.3.4 Data protection and privacy	13
4.3.5 Empathy	14
4.3.6 Justice.....	14
5 Challenges and limitations	15
6 Implications and outlook.....	16
7 References.....	17



Executive Summary

PRECISE4Q sets out to minimize the burden of stroke for the individual and for society through multi-dimensional predictive modeling. The qualitative study presented here contributes to this endeavor in that it helps us to uncover ethically relevant issues relating to the use of these novel technologies in stroke research and clinical practice. As such, it supports us in achieving our goal of developing an ethical framework for PRECISE4Q that is agile and responsive to different stakeholders' needs and suitable to guide decision making in stroke research and clinical practice.

The present deliverable outlines activities carried out in relation to Task 1.5 (WP1). More specifically, it describes the rationale, design, and findings of a qualitative study on stroke patients', caregivers', and healthcare professionals' attitudes towards multi-dimensional predictive modelling in stroke. This work is instrumental to uncovering the ethical priorities of relevant stakeholders in the ethical framework that will be developed within WP1.

The overall objective of the study was to generate a theoretically grounded account of how relevant stakeholders' expectations and attitudes towards data-driven predictive modeling in the four phases of stroke (prevention, treatment, rehabilitation, reintegration) are constructed. In our analysis we particularly focused on the identification of ethically relevant issues relating to the use of medical AI in stroke. In light of this objective, we opted for a qualitative research methodology using the sociology of expectations as a theoretical framework.

We performed a review of the pertinent literature to inform study design and data collection. Guided by prior research, we developed the data collection instrument (interview guide and vignette). Both the interview guide and the vignette were further refined in exchange with the Precise4q consortium partners. Based on two pilot-tests and expert feedback, some minor revisions were implemented (i.e. simplification of language). The study was approved by the ETH Ethics Commission (EK 2019-N-88).

Recruitment in Switzerland, Germany, and Austria was supported by our project partners, namely Charité Berlin, as well as patient and professional associations in the three countries. In total, 14 interviews have been conducted, including 6 with patients and 8 with healthcare professionals specialized in stroke. All interviews were audio recorded and transcribed verbatim for analysis. Data is analyzed in the original language (German) using a combination of inductive and deductive thematic analysis.

Findings of this study indicate that despite being a widely used and discussed term, stakeholders' understanding of what AI was somewhat limited. When asked about how they think AI could be used in healthcare and stroke care, specifically, participants' ideas ranged widely in terms of envisioned areas of application and level of automation. The Precise4q scenario for the four phases of stroke (prevention, acute treatment, rehabilitation, reintegration) was judged as realistic and positive by most participants. However, both healthcare professionals and stroke patients acknowledged some legal, technical (infrastructural), and ethical challenges and concerns in relation to its implementation in practice. Ethically relevant concerns evolved around six core themes: 1) patient autonomy and self-determination AI; 2) disclosure; 3) responsibility; 4) data protection and privacy; 5) empathy; and 6) justice.



1 Overall objective and scope of the deliverable

The present deliverable outlines activities carried out in relation to Task 1.5 (WP1). More specifically, it describes the rationale, design, and findings of a qualitative study on stroke patients', caregivers', and clinicians' attitudes towards multi-dimensional predictive modelling in stroke.

The work presented here is instrumental to uncovering the perceived ethical priorities of relevant stakeholders in the ethical framework that will be developed within WP1. The document concludes by highlighting some of the challenges and limitations and describes next steps.




2 Study rationale and background

In recent years, we have observed a growing hype surrounding the application of big data analytics and artificial intelligence (AI) in medicine as tools to foster personalized medicine. This discourse is no longer limited to technical and scientific communities but has reached popular media and in this way also the general public [1]. However, despite the great technical advances of AI-driven applications in medicine [2], particularly in the field of medical imaging [3], their utility and impact on clinical care remains poorly understood due to their rare application in clinical practice [4].

This raises important questions from a bioethical perspective: How will the adoption of AI in clinical practice impact the doctor-patient relationship? What impact will it have on the core ethical principles that should guide clinical care? In order to address these questions, there has been a call for normative guidelines and ethical frameworks that can help to identify and address the ethical challenges of medical AI [5-10]. However, to date, particularly empirical bioethics research on these issues remains scarce [1, 11].

To address this gap in the current bioethics discourse surrounding AI in healthcare, we seek to generate a theoretically grounded account of how different stakeholders' expectations regarding predictive modeling in stroke are constructed. More specially, we focus on identifying relevant stakeholders' ethical concerns and priorities through semi-structured interviews which allow for an in-depth exploration of individuals' personal stories, reflections, and reasoning [12].

The findings generated through this study will contribute to a better understanding of how expectations surrounding AI-models are articulated by the relevant stakeholder groups. The study will also allow us to identify ethically relevant issues relating to the use of medical AI in stroke research and clinical practice. As such, it will support us in our goal of developing an ethical framework for PRECISE4Q that is agile and responsive to different stakeholders' needs and suitable to guide decision making in practice.



3 Methodology

The present study adopts a qualitative research design using semi-structured interviews as method of data collection. It is guided by the sociology of expectations, which recognizes the ‘generative’ role of individuals’ expectations in shaping scientific and technological change [13, 14]. As such, we consider expectations instrumental not only to explore stakeholders’ acceptance of these novel technologies but also to exploring and understanding their ethical concerns and priorities [14].

Interviews are expected to last about 45 minutes and are conducted either face-to-face at the location of choice of study participant (e.g. at the hospital, at the Health Ethics and Policy Lab) or via telephone. All interviews are carried out by a German speaking researcher trained in qualitative research methods. Interviews conducted in Swiss German are translated to High German for analysis. No financial incentives are provided to study participants.

3.1 Participants

We adopted a convenience sampling approach to identify information-rich cases, covering a wide range of stakeholders likely to be affected by the introduction of predictive modeling in clinical practice for stroke. The exact number of participants depends on theoretical data saturation [15] and can only be determined once data collection is underway.

Table 1 Study populations

Study populations
Individuals who suffered a stroke
Caregivers of stroke patients
Healthcare providers specialized in stroke

Inclusion criteria

Individuals are eligible to participate in this study if they meet the following conditions:

- Belong to one of the study populations listed in Table 1
- Aged >18
- German speaking
- Provide informed consent to participate in the study

Exclusion criteria

Patients with aphasia that do not have the ability to take part in the interview.

3.2 Recruitment

For the purpose of this study, we adopt a threefold recruitment strategy: First, study participants are recruited through and with assistance of the Precise4q consortium partners, in particular CUB. Recruitment in Switzerland is further supported through the existing professional networks and ties of the research team at the Health Ethics and Policy Lab at ETH. To complement this recruitment strategy, we contacted medical professional associations and patient associations based in Austria, Germany, and Switzerland to recruit additional study participants. To date, the study has been advertised in the newsletter and website of a large patient association based in Switzerland and a self-help group based in Germany. In addition, we relied on snowballing to identify prospective study participants.



3.3 Ethics

The study was approved by the ETH Ethics Committee (EK 2019-N-88). The committee reviewed the study protocol, data collection instrument, the participant information sheet and the informed consent form.

3.4 Data management

With the explicit permission of the study participants, all interviews are being recorded with an audio recording device and transcribed verbatim for analysis. Analysis is performed in the qualitative data analysis software MAXQDA. All data is securely stored on an encrypted hard drive.

3.5 Withdrawal of participants

Study participants have the right to withdraw from the study at any time and request the deletion of their data, including the audio files and transcription.

3.6 Data collection tools

The interview guide (Table 2) consists of three sections and serves as the framework for the semi-structured interviews. Specifically, we aimed to explore participants' 1) personal experiences with stroke; 2) their understanding of AI and ideas on how it could be applied in healthcare and stroke, specifically; 3) their expectations and attitudes towards the clinical decision support systems (CDSS) envisioned by PRECISE4Q. In this study, we focused on identifying ethically relevant issues relating to the use of these novel technologies in stroke research and clinical practice.

In the first section, we are particularly interested to find out more about how the study participants perceive their respective situation (stroke/caring for someone who suffered a stroke/treating stroke) and the challenges they face in their particular role. In the second section, we aim to determine participants' understanding of what artificial intelligence is and how they think it could be applied in healthcare and stroke, specifically. Before leading over to the third section of the interview guide, the interviewer provides a detailed description of what Precise4q aims to achieve in particular (data-driven predictive models to enable personalized medicine for stroke). Here, we provide the study participants with concrete examples of application in form of a vignette. In doing so, we aim to elicit their attitudes and expectations with regard to a concrete application of AI in stroke, rather than eliciting unsubstantiated expectations based on participants' (potentially) unrealistic ideas of what AI is and what it can do.

The vignette (




Table 3) was informed by D1.2 (Categorized and ranked clinical challenges and needs) and D1.3 (Use cases and their inputs/outputs specifications). It portrays the patient journey of a fictional patient from failed prevention to stroke, acute treatment, rehabilitation, and reintegration. For each of the four phases, it describes how the technologies (i.e. clinical decision support systems) developed within Precise4q would be used by the medical teams and how it would impact patient outcomes.

Both the interview guide and the vignette were further refined in exchange with the Precise4q consortium partners. Prior to initiating data collection, we carried out two pilot interviews to ensure comprehensibility and validity of the data collection instrument. Based on the feedback we received, the interview guide and vignette were slightly modified. In particular, two versions of the vignette were created; one to be used for interviews with healthcare professionals, and once slightly simplified vignette in terms of language to be used when interviewing stroke patients and informal caregivers.

Based on the interviewee's replies to the questions listed below, the interviewer may divert slightly from the interview guide, such that she may, for example, ask for further explanations or clarifications. In particular, in the last section, probes are used to solicit participants' attitudes towards core ethical issues around consent, trust, responsibility, privacy, patient autonomy, and fairness. Upon completion of the interview, each interviewee is asked if they have anything to add or whether they have any additional questions regarding the topic.

With the explicit permission of the participant, interviews were audio recorded and transcribed verbatim for analysis.

Table 2 Interview Guide (translated from the German original)

<p>1. Personal information regarding participants experiences with</p> <ul style="list-style-type: none"> a. stroke (patients) b. caring for a stroke patient (caregiver) c. treating stroke patients (healthcare professionals)
<p>2. Understanding of AI and its application</p> <ul style="list-style-type: none"> a. In general: Do you see applications of artificial intelligence in your daily life? b. In healthcare and specifically stroke: Do you see applications of artificial intelligence in healthcare? Can you think of ways in which artificial intelligence could be used in stroke?
<p>3. Expectations regarding AI in stroke</p> <ul style="list-style-type: none"> a. Do you consider the presented scenario realistic? Why (not)? <ul style="list-style-type: none"> i. Do you have any hopes/concerns in relation to this? b. How do you expect that AI would affect your current situation? <ul style="list-style-type: none"> i. Do you have any hopes/concerns in relation to this? c. How do you think AI will evolve in the future? Why? <ul style="list-style-type: none"> i. Do you have any hopes/concerns in relation to this?




Table 3 Vignette Precise4q (translated from the German original)

Within Precise4q we aim to develop predictive models using artificial intelligence. With these models, it will be possible to personalize treatment and rehabilitation for stroke patients. The predictions will be based on a series of clinical data, such as imaging data, as well as information on medical history and lifestyle data. The goal is that each and every patient should receive the most promising treatment based on their individual profile. It's important to note that these programs are intended to support and not replace the healthcare team. To illustrate this, here a concrete example:

Remo is 52 years old and works in the accounting department of a large company. His job is stressful and Remo works a lot of overtime. He has been a smoker for many years and does hardly any sport. Due to his overweight and elevated cholesterol level, his family doctor advised him to pay attention to his diet and do more sports. When Remo is about to leave the office one evening, he suddenly has trouble moving his right arm and leg, and he also finds it difficult to talk. His work colleagues react quickly and call an ambulance. Remo has suffered a stroke.

Arriving at the hospital, the emergency team employs a **new computer program** that uses simulations to predict the success of various treatment options. The program uses Remo's existing electronic patient file for this purpose. The patient file contains clinical information, such as genetic data, as well as information on Remo's medical history and lifestyle. The rapid response of the emergency team, based on the calculations of the computer program, reduces Remo's risk of suffering permanent function loss.

After his discharge, Remo begins a three-month rehabilitation phase to regain the functions impaired by the stroke through targeted training. In particular, the function of his left leg is still severely impaired and he has difficulty walking without a walking aid. In addition, Remo has difficulties concentrating and is therefore worried that he will not be able to return to his job. Using **computer-aided simulations**, Remo's treatment team can see how different therapy plans would affect Remo's recovery progress. Based on these calculations, they design a rehabilitation program tailored to Remo that adapts daily to Remo's condition to achieve an optimal rehabilitation goal.


To ensure that Remo gets the support he needs at home (e.g. shopping or doing the housework), Remo's treatment team again uses the **computer program**. It helps them to assess whether and how intensive support and care Remo needs in order to follow an independent life. The program also calculates Remo's individual risk of suffering another stroke. Based on this information, Remo's treatment team recommends targeted preventive measures.

3.7 Data analysis

The interviews lasted between 23 and 78 minutes, 39 min on average (9h and 11 min in total). Upon transcription, all text files are imported into the qualitative data analysis software MAXQDA for analysis. Data analysis is carried out in parallel to data collection using a combination of inductive and deductive thematic analysis following Braun and Clarke [12].

In a first step, data is coded deductively to identify stakeholders' concerns related to predictive modeling in stroke based on pertinent ethical principles (see section 3.6) [16]. Following this first coding phase, data is coded inductively based on emerging topics (data-driven coding). In a final step, codes are merged into themes in an iterative process to ensure that themes are coherent and distinctive.

The interviews are analyzed in the original language (German).



4 Findings

4.1 Participant characteristics

In total, we conducted 14 interviews, including 6 with patients and 8 with healthcare professionals specialized in stroke from September to December 2019. Of those interviews, 10 were conducted face-to-face. Table 4 below presents the participant characteristics.


Table 4 Participant characteristics

Healthcare professionals	N=8	Patients	N=6
Professional Role		Age group (mean)	
Physician	7	18-30	
Nurse		31-50	2
Occupational therapist	1	51-60	3
Physio therapist		>60	1
Gender		Gender	
male	5	male	2
female	3	female	4
Years in practice (mean)		Time since stroke	
<1 year	4	<1year	3
1-5 years	3	1-5 years	
6-10 years		>5years	3
>10 years	1		
Setting		Setting	
Prevention		Prevention	
Acute	6	Acute	3
Rehabilitation	2	Rehabilitation	
Reintegration		Reintegration	3

4.2 Making sense of AI

Findings of this study indicate that despite being a widely used and discussed term, stakeholders' understanding of what AI is was somewhat limited. Both patients and healthcare professionals considered that they have little intentional interaction with AI in their daily lives. They did, however, acknowledged that they may be confronted with AI unconsciously. This invisible interference of AI in the background seemed to be of little concern, with none of the participants showing or expressing any form of discomfort in relation to the presumed presence of AI in their life. Some participants also referred to voice assistance (e.g. Siri) as examples of AI in their day-to-day life that they were aware of but chose not to use.

When asked about how they think AI could be used in healthcare and stroke specifically, participants' ideas ranged widely in terms of envisioned areas of application and level of automation. While some participants could see AI systems taking over mundane administrative tasks thus freeing up clinicians' schedules, others imagined sophisticated AI-powered chips that could be implanted in a stroke survivor's brain to restore cognitive functions and mobility. Medical AI was in some instances also equated with healthcare robotics that could one day perform routine tasks in hospitals. Other participants referred to it as a large data management system that could allow physicians to pull all



relevant patient information together in one single program rather than having to open up several different pdf-files or that could help them aggregate and assess relevant research findings.

4.3 Ethically relevant concerns related to the Precise4q scenario

When presented with the Precise4q scenario (i.e. clinical decision support systems for the four phases of stroke), most participants judged it as realistic and positive. They did, however, acknowledge some legal, technical (infrastructural), and ethical challenges and concerns related to the implementation in practice. It is important to note that some of the concerns we identified are not exclusively inherent to medical AI but could be intensified through the introduction of these technologies.

The ethically relevant challenges and concerns we identified evolved around six core themes:

- 1) Patient autonomy and self-determination (see 4.3.1)
- 2) Disclosure (see **Error! Reference source not found.**)
- 3) Responsibility (see 4.3.3)
- 4) Data protection and privacy (see 4.3.4)
- 5) Empathy (see 4.3.5)
- 6) Justice (see 4.3.6)


4.3.1 Patient autonomy and self-determination

Patients as well as healthcare professionals stressed the importance of patient autonomy and self-determination. Indeed, most participants either implicitly or explicitly acknowledged the patient's active role in the decision-making process (patient autonomy and self-determination). As such, there was a consensus that decisions that concern the patient's health and treatment should, whenever the patient's health and cognitive state permit, be taken by the patient and the healthcare professional (shared decision-making). Here, it was, however, noted that if presented with different options, patients may not always necessarily decide for the most promising course of treatment (as seen from the healthcare professional's point of view) due to a misalignment with their individual goals and values.

Most of the participants considered it unlikely that the introduction of AI-driven systems would change anything with respect to the current status quo in terms of patient autonomy and self-determination. AI-driven clinical decision support systems were simply viewed as a new tool that can support clinicians and patients in making an informed decision, recognizing that patients could still decide against the system's recommendations. It was, however, pointed out by some participants that presenting options could, to a certain extent, already impose constraints and limit patient autonomy and that some individuals may not appreciate this form of support.

4.3.2 Disclosure

When asked whether or not a patient should be informed about the use of AI-driven clinical decision support systems, participants' attitudes differed. While most healthcare professionals and patients agreed that the use of AI-driven systems should be disclosed to patients, some healthcare professionals argued that disclosure may not always be necessary or useful. In particular, they deemed disclosure necessary only in cases where a physician disagreed with the system's recommendation or



in cases where a decision was solely based on the system without any human interference or checking (autonomous systems). In support of this viewpoint, they said that also in current practice they would not inform the patient about the respective book or clinical study they used to derive their recommendations. Most patients, on the other hand, expected to be informed about the use of AI-driven prediction tools, highlighting the importance of receiving an explanation and reasoning for treatment recommendations put forward by healthcare professionals. In this context, some patients also recalled and criticized the lack of information provision they experienced during their post-stroke hospitalization. During acute care, disclosure was deemed unnecessary for practical reasons.

Disclosing the use of AI-driven clinical decision support systems was also considered to have a potential impact on trust in the doctor-patient relationship. Participants presumed that factors, including a patient's age and personality, their attitudes towards AI, and their experiences with the healthcare system, more generally, may determine whether medical AI would strengthen or weaken trust in the relationship and subsequently in health relevant recommendations.

4.3.3 Responsibility


Most participants agreed that the treating healthcare professional would be responsible for acting (or not acting) on the algorithm's prediction and resulting recommendations. In this context, they emphasized that healthcare professionals should be able to interfere and make corrections to the recommendations produced by the algorithm, in case they noticed that something was off or that potentially relevant information was missing from the calculation. This underlined the supportive or guiding role (as opposed to replacement) participants attributed to the AI-driven system and the importance of ensuring healthcare professionals' critical appraisal and evaluation of its predictions. Both healthcare professionals and patients voiced concerns about a future where healthcare professionals would blindly trust algorithms, neglecting their own responsibility. In this context, algorithmic transparency or explainability of prediction modelling was named a prerequisite.

When discussing responsibility, only few of the interviewed participants referred to other actors who they considered potentially responsible, including programmers and developers, as well as system administrators. In this case, responsibility was considered to be shared among those actors entering data, those programming the predictive algorithms, and those ultimately executing the resulting recommendations.

4.3.4 Data protection and privacy

When discussing concerns regarding data protection and privacy, many participants explicitly or implicitly acknowledged the patient's right to choose who they want to share information with and to even withhold potentially relevant information from healthcare professionals. Particularly patients from the acute care setting commented on a lack of privacy that existed also independently of AI-driven systems. One of them described privacy as a privilege that was left at the entrance when being admitted to a hospital.

This reflects the distinction we observed between how privacy was understood before the onset of stroke and post-stroke. While participants acknowledged that AI-driven systems could potentially be useful to support stroke prevention, early intervention at onset, and anamnesis (e.g. by helping to



identify the time of stroke), such technologies were viewed as a violation of privacy, particularly by patients.

4.3.5 Empathy


Empathy emerged as a key ethical concern. On the one hand, participants saw the potential of AI-driven systems in helping to free up physicians' time by taking over mundane, routine tasks, including administrative work. Participants also considered it as a possible remedy for staff shortage in healthcare. According to this view, AI-driven systems could help to strengthen the doctor-patient relationship and foster humanistic care and empathy.

On the other hand, particularly patients voiced concerns that the time saved may not necessarily allow for more time spent with patients but may, on the contrary, result in budget and personnel cuts. According to this view, AI-driven systems could pose a threat to maintaining human touch and empathy in healthcare and may prevent healthcare professionals from capturing relevant information (e.g. patients' values and beliefs) that can only be gained from human-to-human interaction.

4.3.6 Justice

With respect to the principle of justice two key dimensions were touched upon by participants. The first one relates to the distribution and equity of access to medical AI. Here, participants raised concerns that predictive modelling may become a luxury of rich western clinics excluding those less fortunate from the benefits of predictive modeling. Particularly financial aspects related to the implementation and maintenance were pointed out as potential threats to improving stroke care for all patients not only the fortunate few.

The second dimension relates to fairness in medical decision-making. Both healthcare professionals and patients considered AI as a neutral, objective and consequently fair system as opposed to healthcare professionals that should not, but may in fact, though subconsciously, treat individuals differently depending on whether or not they like them. Despite the interviewer's attempt to subtly allude to issues related to algorithmic bias and fairness, this was an issue that was not taken up by participants and thus remained undiscussed.



5 Challenges and limitations

Our findings should be considered in light of its limitations.

Given our qualitative study design, we cannot claim that our findings are representative of the respective sub-populations we investigated in this study. We also need to acknowledge some challenges we faced in recruitment. When contacted about the study, two rehabilitation centers in Switzerland voiced concerns that participation in the study would be too demanding and burdensome for their elderly patients. The representative of a German patient association was willing to share our call for participants with the associations large peer-support network via Email but noted that most patients would be more interested to take part in clinical studies (i.e. receive clinical benefits).

Another limitation can be seen in the exclusion of certain groups, including patients with aphasia and those patients suffering from stroke and/or aging related cognitive limitations. These groups were excluded due to our concerns that taking part in the interview would be too demanding or burdensome for them. However, we acknowledge the fact we may have failed to adequately capture their viewpoint. One strategy we adopted to mitigate this concern was to also interview informal caregivers of stroke patients. Recruitment of caregivers, however, proved to be more challenging than anticipated, and has so far been unsuccessful. We are thus currently undertaking efforts to reach this study population through caregiver support groups.

Another group that were excluded are non-German speaking residents. In excluding this group, we may have failed to capture the expectations and attitudes of individuals from different socio-cultural backgrounds and thus potentially diverse concerns and ethical priorities. In light of the major cultural differences that we know exist in doctor–patient communication [17, 18], further work is needed to ensure that an ethical framework for medical AI in stroke is agile and compatible with different socio-cultural values and belief systems. To address this limitation of our work, we intend to consult experts in the field of intercultural health communication.




6 Implications and outlook

PRECISE4Q sets out to minimize the burden of stroke for the individual and for society through multi-dimensional predictive modeling. The qualitative study presented here contributes to this endeavor in that it helps us to uncover ethically relevant issues relating to the use of these novel technologies in stroke research and clinical practice. As such, it supports us in achieving our goal of developing an ethical framework for PRECISE4Q that is agile and responsive to different stakeholders' needs and suitable to guide decision making in stroke research and clinical practice.

Grounded in the findings generated through the work presented here, we aim to conduct further interviews in early 2020. In particular, we aim to recruit informal caregivers of stroke patients, a group that we so far were unfortunately unable to reach but that we consider of great relevance. We also intend to consult experts in the field of intercultural health communication to ensure that the ethical framework for medical AI in stroke that will result from our work is compatible with different socio-cultural values and belief systems.

This work will further inform the development of an interactive deliberative dashboard, which aims to guide researchers and healthcare professionals in navigating novel ethical challenges in the context of big data health research and predictive modeling in stroke (D1.8). The deliberative dashboard will be developed in close collaboration with the PRECISE4Q consortium partners.



7 References

1. McDougall, R.J., *Computer knows best? the need for value-flexibility in medical AI*. Journal of Medical Ethics, 2019. **45**(3): p. 156-160.
2. Jiang, F., et al., *Artificial intelligence in healthcare: past, present and future*. Stroke and vascular neurology, 2017. **2**(4): p. 230-243.
3. Lee, E.-J., et al., *Deep into the brain: artificial intelligence in stroke imaging*. Journal of stroke, 2017. **19**(3): p. 277.
4. He, J., et al., *The practical implementation of artificial intelligence technologies in medicine*. Nature medicine, 2019. **25**(1): p. 30-36.
5. Vayena, E., A. Blasimme, and I.G. Cohen, *Machine learning in medicine: Addressing ethical challenges*. PLoS medicine, 2018. **15**(11): p. e1002689.
6. Vayena, E., et al., *Digital health: meeting the ethical and policy challenges*. Swiss medical weekly, 2018. **148**: p. w14571.
7. Wong, P.-H., *Democratizing Algorithmic Fairness*. Philosophy & Technology, 2019: p. 1-20.
8. Xafis, V., et al., *An ethics framework for big data in health and research*. Asian Bioethics Review, 2019. **11**(3): p. 227-254.
9. Char, D.S., N.H. Shah, and D. Magnus, *Implementing machine learning in health care—addressing ethical challenges*. The New England journal of medicine, 2018. **378**(11): p. 981.
10. Ferretti, A., M. Schneider, and A. Blasimme, *Machine Learning in Medicine: Opening the New Data Protection Black Box*. Eur. Data Prot. L. Rev., 2018. **4**: p. 320.
11. Rigby, M.J., *Ethical dimensions of using artificial intelligence in health care*. AMA Journal of Ethics, 2019. **21**(2): p. 121-124.
12. Braun, V. and V. Clarke, *Using thematic analysis in psychology*. Qualitative research in psychology, 2006. **3**(2): p. 77-101.
13. Brown, N. and M. Michael, *A sociology of expectations: retrospectively prospecting and prospecting retrospects*. Technology analysis & strategic management, 2003. **15**(1): p. 3-18.
14. Borup, M., et al., *The sociology of expectations in science and technology*. Technology analysis & strategic management, 2006. **18**(3-4): p. 285-298.
15. Saunders, B., et al., *Saturation in qualitative research: exploring its conceptualization and operationalization*. Quality & quantity, 2018. **52**(4): p. 1893-1907.
16. Jobin, A., M. Ienca, and E. Vayena, *The global landscape of AI ethics guidelines*. Nature Machine Intelligence, 2019. **1**(9): p. 389-399.
17. Schouten, B.C. and L. Meeuwesen, *Cultural differences in medical communication: a review of the literature*. Patient education and counseling, 2006. **64**(1-3): p. 21-34.
18. Paternotte, E., et al., *Factors influencing intercultural doctor–patient communication: A realist review*. Patient education and counseling, 2015. **98**(4): p. 420-445.