



Information that matters: Exploring information needs of people affected by algorithmic decisions

Timothée Schmude^{a,b,c,*}, Laura Koesten^a, Torsten Möller^{a,b}, Sebastian Tschitschek^{a,b}

^a Faculty of Computer Science, University of Vienna, Vienna, 1090, Austria

^b Research Network Data Science, University of Vienna, Vienna, 1090, Austria

^c Doctoral School Computer Science, University of Vienna, Vienna, 1090, Austria

ARTICLE INFO

Keywords:

Explainable AI
Understanding
Information needs
Affected stakeholders
Question-driven explanations
Qualitative methods

ABSTRACT

Every AI system that makes decisions about people has a group of stakeholders that are personally affected by these decisions. However, explanations of AI systems rarely address the information needs of this stakeholder group, who often are AI novices. This creates a gap between conveyed information and information that matters to those who are impacted by the system's decisions, such as domain experts and decision subjects. To address this, we present the "XAI Novice Question Bank", an extension of the XAI Question Bank (Liao et al., 2020) containing a catalog of information needs from AI novices in two use cases: employment prediction and health monitoring. The catalog covers the categories of data, system context, system usage, and system specifications. We gathered information needs through task based interviews where participants asked questions about two AI systems to decide on their adoption and received verbal explanations in response. Our analysis showed that participants' confidence increased after receiving explanations but that their understanding faced challenges. These included difficulties in locating information and in assessing their own understanding, as well as attempts to outsource understanding. Additionally, participants' prior perceptions of the systems' risks and benefits influenced their information needs. Participants who perceived high risks sought explanations about the intentions behind a system's deployment, while those who perceived low risks rather asked about the system's operation. Our work aims to support the inclusion of AI novices in explainability efforts by highlighting their information needs, aims, and challenges. We summarize our findings as five key implications that can inform the design of future explanations for lay stakeholder audiences.

1. Motivation

Society is in the process of negotiating which AI systems can be used responsibly in public institutions and organizations (Züger and Asghari, 2023; Lucaj et al., 2023; Lima et al., 2023). These systems should be transparent, accountable, and support human oversight, according to the legal frameworks that are being developed by regulatory bodies (European Commission, 2021). Research in the field of explainable AI (XAI) addresses the realization of these values by examining how AI systems can be made understandable to stakeholders and how the given information enables specific actions, such as the contestation of decisions (Alfrink et al., 2023) and the assignment of accountability (Langer et al., 2021). As explanations thereby assume a key role in the deployment of responsible AI systems for the public (Züger and Asghari, 2023), the question of what constitutes a *good* explanation becomes central.

Good explanations of AI systems are aligned with the information needs of the stakeholders (Miller, 2019; Mueller et al., 2019; Byrne,

2023; Ehsan et al., 2022) and are a means to help them achieve their aims, which can depend on the stakeholder's role (Langer et al., 2021; Barredo Arrieta et al., 2020), AI literacy (Long and Magerko, 2020), and domain knowledge (Wang and Yin, 2021). However, information needs of people who are impacted directly and personally by algorithmic decision-making (ADM) systems (European Parliament. Directorate General for Parliamentary Research Services, 2019; Golpayegani et al., 2023), such as domain expert users and decision subjects, have not yet been explored in-depth in XAI research. As a result, current explanations rarely meet the needs of these stakeholder groups. This is critical, as providing explanations for these stakeholder groups can be essential to reduce the risks of power asymmetry (Ananny and Crawford, 2018), inequality (Lopez, 2019), and informational unfairness (Schoeffler et al., 2022) in deployments of public AI (Kuziemiński and Misuraca, 2020; Henman, 2020; Brown et al., 2019).

* Corresponding author.

E-mail address: timothee.schmude@univie.ac.at (T. Schmude).

<https://doi.org/10.1016/j.ijhcs.2024.103380>

Received 16 January 2024; Received in revised form 15 July 2024; Accepted 22 September 2024

Available online 26 September 2024

1071-5819/© 2024 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

However, building explanations that aid domain experts and decision subjects in understanding ADM systems can be challenging. First, these stakeholder groups are diverse and can include “AI novices”,¹ i.e., people with little technical literacy who have heterogeneous information needs and prior knowledge, which necessitates accessible and adaptive explanations (Shulner-Tal et al., 2022; Conati et al., 2021; Kim et al., 2024). Second, the effects of explanations are influenced by perceptions of the deploying institution (Brown et al., 2019; Ehsan et al., 2021), meaning that if the institution is not trusted, the ADM system and its explanations are not trusted either (Woodruff et al., 2018; Scott et al., 2022). Third, understanding is an interconnection of knowledge and reasoning (Grimm, 2019; Keil, 2006) involving numerous cognitive processes (Keil, 2006; Lombrozo and Wilkenfeld, 2019) that can barely be captured by standardized large-scale tests (Sato et al., 2019). Lastly, AI lifecycles encompass large amounts of information (Dhanorkar et al., 2021) and, therefore, require a preselection of what is deemed relevant.

These interactions between cognitive, perceptual, and organizational dimensions can be difficult to address in explanation design. Recent XAI research, therefore, advocates for providing tailored explanations, based on the stakeholder’s aims (Dhanorkar et al., 2021; Freiesleben and König, 2023) and information needs (Liao et al., 2020, 2021), instead of explaining what is assumed to be relevant (Miller, 2023). This study thus follows research that aims to build “human-centered” explanations of AI by making them more accessible and better suited to human cognitive processes (Shin, 2023; Ehsan et al., 2022).

This paper contributes to including domain experts and decision subjects in explainability by presenting an empirical study of their information needs, understanding, and perceptions. We use two ADM use cases as examples: an employment prediction algorithm and a health wristband for geriatric care. To collect participants’ information needs, the explanation setting is flipped: Participants ask questions about the algorithmic systems to evaluate their adoption and are given verbal explanations in response. Participants thus independently choose which information they inquire about, reducing the amount of pre-selected information in favor of demand-based information. We analyze the impact of question-driven explanations on participants’ self-reported understanding and decision confidence, as well as the role of “explanatory stances” (Keil, 2006) in their information acquisition. Lastly, we analyze participants’ perceptions of risks and benefits before and after receiving explanations and describe perceived changes. Our work is guided by the following research questions:

- [RQ1] *Information Needs*: What information do AI novices who could be affected by algorithmic decisions need in order to decide about adopting an ADM system?
- [RQ2] *Understanding*: How do question-driven explanations interact with participants’ understanding?
- [RQ3] *Perception*: How do question-driven explanations impact perceptions of the systems’ risks and benefits?

The contributions of our work include (i) a catalog of affected stakeholders’ information needs summarized in the “XAI Novice Question Bank” (Fig. 3)²; (ii) recommendations on how to address understanding challenges such as blind spots and outsourcing and on how to use explanatory stances in question-driven explanations; and (iii) reflections on the relevance of information about intention in explanations for affected stakeholders and considerations of how to incorporate insights from “civic education” (Lupia, 2016) into explainability.

¹ We use the term in reference to “visualization novices” (Burns et al., 2023) and synonymously with “lay people” (Shen et al., 2020; Szymanski et al., 2021; Schmude et al., 2023).

² In reference to Liao et al. (2020)’s XAI Question Bank.

2. Background and related work

We structured our literature review in three sections. The first section describes related work in Human-Centered AI on high-risk systems and stakeholder-focused explainability (Section 2.1). The second section covers the interplay between explanations and understanding as well as definitions of understanding and approaches to analyze it (Section 2.2). The third section describes the importance of people’s perceptions of ADM systems’ risks and benefits and the relationship between explanations and the assessment of normative values (Section 2.3).

2.1. Human-centered AI (HCAI)

The notion of human-centered AI aims to provide principles and frameworks that enable “ethical, interactive, and contestable use” of AI systems, covering both interaction design approaches as well as the incorporation of ethical values (Capel and Brereton, 2023). Shneiderman (2022), in line with other scholars (Araujo et al., 2020; Xu, 2019; Shin, 2023; Shin and Shin, 2023), emphasizes that human-centered AI systems should (1) build on user observation and stakeholder engagement and should (2) empower rather than replace people, i.e., be “accessible and controllable” (Shin and Shin, 2023). Prioritizing people, in this sense, means giving attention to users and other stakeholders throughout the development, deployment, and evaluation of AI systems (Shin, 2023) while observing human values such as fairness, trust, and accountability (Shneiderman, 2022). Similar notions have been brought forward in multiple responsible and trustworthy AI guidelines, such as the Montreal Declaration for Responsible Development of AI, the Vienna Manifesto on Digital Humanism, and, more recently, the EU AI Act (European Commission, 2021).

Human-centered design becomes especially important when considering “high-risk” AI applications — systems that pose risks to the health, safety, or fundamental rights of persons (European Commission, 2021). Rendering high-risk systems transparent, understandable, and explainable to users and preserving their control is seen as a mitigation of these risks (Adadi and Berrada, 2018), but the practical realization proves challenging (Ananny and Crawford, 2018). As people who are affected by ADM systems tend to be AI novices, explanation approaches need to be tailored to their needs, but as Shin (2023) notes, there are few “well-established ways of incorporating lay users into the process of designing algorithms”. Further, no universal explanation method exists that can cover the needs of all stakeholder groups equally (Phillips et al., 2021). Therefore, the question of how AI novices’ information needs can be analyzed and incorporated into explanation approaches is an open challenge in XAI research (Shulner-Tal et al., 2022).

To this end, several studies have examined how explanations can help lay users to understand algorithmic systems: Szymanski et al. (2021) examined explanations for model predictions and found that lay users preferred a visual format over a textual one but had difficulties correctly interpreting the visual explanation, which led to the development of a new explanation combining both formats. Cheng et al. (2019) found that interactive explanations and white-box models improved lay participants’ comprehension; however, a later study by Bove et al. (2022) could not replicate the results regarding interactive explanations and found the adverse effect. Bove et al. (2022) suggest that the diverging results could have stemmed from the difference in task domains (student admission vs. car insurance) and emphasize that context influences the explanation’s effect. Lastly, Bertrand et al. (2023) used interviews to elicit the information needs of domain experts and end-users about a financial robo-advisor, which they then used to develop and test explanations for the system’s recommendations in a large-scale survey. Their analysis showed that feature-based SHAP explanations, one of the most well-known explanation methods, did not increase users’ understanding nor enable a suitable trust calibration. In contrast, our study takes a more general perspective on the qualitative

exploration of participants' global information needs and perceptions when they vote on the adoption of a high-risk ADM system.

These multi-faceted insights indicate how difficult it can be to tailor explanations of algorithmic systems to lay people. For this reason, this study aims to take a closer look at the information needs of AI novices and how they can be categorized. To elicit these needs, participants in our study assume the role of a decision-maker who can ask questions and receive explanations in return, as described in Section 3. We thus “flip” the explanation setting in order to answer RQ1-Information Needs.

2.2. Explanations and understanding

One purpose of explanations is to increase the recipient's *understanding* of a system (Langer et al., 2021). This increased understanding can help people to pursue their own goals, such as to assess the system's values (Kim et al., 2023), consider its recommendations (Lee et al., 2019), or contest its decisions (Alfrink et al., 2023). Shin et al. (2022a) further state that the understandability of explanations impacts how users evaluate both the quality of explanations and the system's fairness, accountability, and transparency. However, despite this vital role of understanding, there is no general conception of what is most important about an AI system, nor how this understanding can be evaluated (Schmude et al., 2023).

Numerous definitions of understanding have been proposed in the literature (Grimm, 2019; Baumberger et al., 2017; Zagzebski, 2019; Keil, 2006). For this work, we use a definition that aims to preserve its epistemic aspects while allowing for operationalization in an empirical setting. We thus define understanding as (i) connecting and applying information, whereas “knowledge” only stores information (Grimm, 2019; Baumberger et al., 2017), (ii) involving different varieties depending on whether people (intention, norms) or objects (causality, natural laws) are understood (Grimm, 2019), and (iii) being the attempt to grasp the underlying structure of a phenomenon by way of simplification (Zagzebski, 2019).

2.2.1. Gaps in understanding and explanatory stances

To integrate the epistemic aspects of understanding practically, we use models provided by the cognitive sciences. Keil (2006) and Lombrozo and Wilkenfeld (2019) state that explanations increase understanding by calling attention to *gaps in understanding*. These gaps can be dealt with by (i) filling them with new information, (ii) “outsourcing” them to another mind, or (iii) deciding that they are irrelevant. In this sense, developing understanding means dealing with understanding gaps until one feels to have reached a “working understanding” (Keil, 2003, 2006). This process can be perturbed by the “illusion of explanatory depth” (Rozenblit and Keil, 2002), an overestimation of one's understanding, which can lead to incomplete understanding (Keil, 2006; Lombrozo and Wilkenfeld, 2019; Chromik et al., 2021).

For our analysis, we further use the concept of three *explanatory stances* (Keil, 2021, 2006; Lombrozo and Gwynne, 2014; Lombrozo and Wilkenfeld, 2019), introduced by Dennett (1998): a theory aiming to capture how people make sense of information and generalize from it in order to develop predictive strategies. The three explanatory stances are described below in the order *mechanical*, *design*, and *intentional* stance.

- The *mechanical stance* focuses on low-level technical detail, such as “parts, processes, and proximate causal mechanisms” (Lombrozo and Wilkenfeld, 2019). For example, explaining why an alarm clock buzzes in a mechanical stance would state that the circuit connecting the power source and buzzer was completed (Lombrozo and Wilkenfeld, 2019). In contrast to functional and intentional explanations, mechanistic³ explanations do not take into account aims or beliefs (Páez, 2019).

- The *design stance* focuses on functions and goals beyond mechanical interactions in a system (Keil, 2006) and describes why things happen by looking at their functional purposes (Lombrozo and Gwynne, 2014). Explaining why an alarm clock buzzes from a design stance would state that it functions according to its design: waking its owner at the set time (Lombrozo and Wilkenfeld, 2019; Dennett, 2006).
- The *intentional stance* focuses on a system's purpose as defined by the human beliefs or aims ingrained in it (Keil, 2006; Dennett, 1998; Yurrita et al., 2022). Explaining why an alarm clock buzzes in an intentional stance might refer to its purpose of waking its owner, whose desire is to have enough time before leaving for work.

Previous XAI literature used explanatory stances to examine how explanations are perceived and evaluated (Páez, 2019; Zerilli, 2022; Miller, 2019; Byrne, 2023). For example, Byrne (2023) describes that explanatory stances shed light on how people explain the behavior of algorithmic systems to themselves, which might differ depending on whether they interpret the decisions as results of an algorithmic process or as depictions of decisions by a human (e.g., a job counselor). The stances, therefore, aid in identifying whether the users' information needs focus on the systems' technical or intentional aspects and thus make it easier to match their information needs to a specific explanation (Byrne, 2023). In our approach, we use the stances to differentiate if participants asked for information about technical details, functional relations, or the sociotechnical context.

2.2.2. Question-driven explanations

Questions assume a central role in XAI, as explanations of AI systems can be defined as responses to questions that are asked by the user. Miller (2019) defines three classes of these questions (what, how, and why), which are meant to cover different levels of causal reasoning through their responses. A comparable classification approach has been developed earlier by Lim and Dey (2009), who defined five intelligibility questions (what, why, why not, what if, and how to) to capture questions that “end-users of novel systems may ask” to convey to users the practical functionality of these systems. A practical implementation of question-driven explanation was pursued even earlier by Ram (1989), who built AQUA — a system that used questions about a story's content to generate explanations and advance its knowledge.

From a didactic perspective, Wiggins and McTighe (2005) describe questions to be “doorways to understanding”, through which learners explore concepts, theories, and problems that reside within the content. When we thus define explanations as answers to questions, it is noteworthy that these questions often remain tacit in XAI. In response, Liao et al. (2021) propose a question-driven design process for explanations, which is grounded in the needs and questions of users. The concept further builds upon their previous work on the XAI Question Bank (Liao et al., 2020), a collection of expert users' information needs represented as prototypical questions (e.g., “What kind of output does the system give?”). In our setting, we apply the question-driven design process by providing participants with explanations as responses to their questions. We aim to examine how participants' understanding is changed by these explanations and whether they allow for a better matching between information and explanatory stances. To this end, we aim to answer RQ2-Understanding.

2.3. Perceptions of risks and benefits of algorithmic systems

The acceptance and responsible use of algorithmic systems in society majorly depends on how these systems are perceived (Jakesch et al., 2022; Shin, 2022; Züger and Asghari, 2023). If the “technical affordances” of a system diverge from the “social needs” (Ehsan et al., 2023), i.e., if the system does not or cannot do what is expected from it, it can result in users' distrust in the deploying institution

³ Both the terms ‘mechanical’ and ‘mechanistic’ are used in the literature (Keil, 2006; Lombrozo and Wilkenfeld, 2019; Keil, 2021). Where ‘mechanical’ refers to the stance, ‘mechanistic’ refers to the explanation.

and their rejection of the system (Jakesch et al., 2022). Explainability is understood to influence these perceptions, taking a “facilitating” role in how users evaluate an algorithm’s features and their attitudes towards it (Shin, 2021; Langer et al., 2021). Shin (2023) describes this as a “dual-step flow”, in which explainability allows users to determine their perceptions of the system’s fairness, transparency, and accountability (non-functional attributes), which then, in turn, lets them calibrate their trust in the system and their evaluation of accuracy and personalization (functional attributes). Trust, in this sense, takes a mediating role in that it connects explainability to the system’s evaluation by the user (Shin, 2022).

Explanations thus play an essential role in helping people adjust their perceptions of whether an algorithmic system adheres to normative values. But to fulfill this facilitating function, the information provided in explanations needs to match “what users really wish to understand” (Shin, 2023). Ensuring this fit can further improve the perceived “informational fairness” of algorithmic systems, meaning that information is consistent, reveals relationships between input and output, and includes actionable insights (Schoeffer et al., 2022). However, several studies have shown the complex interrelations that exist between normative values, trust, acceptance, and explanations (Shulner-Tal et al., 2022; Shin et al., 2022b; Schoeffer et al., 2022), and the meaning of normative values such as fairness, trust, and transparency is not strictly fixated in colloquial speech, making it difficult to compare self-reports by participants (Schmude et al., 2023). Motivated by Corvite et al. (2023)’s work, we therefore use an evaluation based on perceived risks and benefits of algorithmic systems. By using a two-dimensional scale that depicts both risk and benefit simultaneously (described in Section 3), we further aim to let participants consider the positive and negative consequences of deploying a specific algorithmic system. This focus on the explanations’ effect on these perceptions is reflected in RQ3-Perceptions.

3. Method

To elicit the information needs of AI novices, we conducted an interview study in German with 24 participants recruited from a local job agency and a local apartment complex inhabited mainly by retirees (Section 3.4). We used a qualitative approach for this study to allow for comprehensive conversations with participants and the collection of data that could be analyzed thematically (Braun and Clarke, 2006). The interviews were thus designed to develop a theoretical account of information needs that is grounded in empirical observation (Urquhart et al., 2009) (Section 3.3). The focus of this study is therefore on the conceptualization and exploration of AI novices’ information needs, which aims to lay the groundwork for further qualitative and quantitative examination in future work.

We presented participants with one of two algorithmic decision-making (ADM) use cases (Section 3.1): an employment prediction algorithm or a health wristband for geriatric care. Participants were tasked to decide whether the presented use cases should be adopted. Before deciding, participants had the opportunity to gather information about the systems by asking questions in two 15-minute inquiry phases. In response to their questions, they received verbal explanations from the main author of this paper. The study closed with an interview at around 50 min. The data was analyzed using inductive and deductive analysis (as described in Section 3.3).

The procedure was designed to fulfill multiple study goals. Firstly, the decision task was meant to give participants a sense of importance for their decision and to motivate them to acquire relevant information. Secondly, by letting participants inquire freely and flexibly about the use cases in the first phase and providing the XAI Question Bank (Liao et al., 2020) only in the second phase, we could examine both participants’ intuitive information needs and the bank’s impact on their inquiry. Third, by eliciting participants’ understanding of the ADM systems and their perceptions of the risks and benefits, we can examine the interactions between these aspects and participants’ information needs. The whole study setup is described in Section 3.2, and the study procedure is depicted in Fig. 1.

3.1. Use cases

The two use cases were selected as representative forms of so-called “high-risk” ADM systems, i.e., systems that pose risks to the health, safety, or fundamental rights of people (European Commission, 2021). In particular, they are examples of the domains *employment* and *biometrics*, which are defined as areas with high-risk potential in the upcoming EU AI Act (European Commission, 2021).

Both use cases are thus examples of more general application areas of high-risk algorithmic systems. While models, data, and context are specific to each use case, the overarching themes of information inquiry and the utility of question-driven explanations are broadly applicable. They can be transferred to the explanation design for other algorithmic systems. In the following, each use case is briefly described, and its selection is motivated.

3.1.1. Use Case A: AMS (employment prediction) algorithm

Description. The AMS algorithm⁴ is a scoring system meant to predict the employability of job-seekers in Austria. It was developed by a private company for the Austrian Public Employment Agency between 2015 and 2021 but was never used as a live system and was put on hold in 2021 due to privacy objections (Allhutter et al., 2020b). To predict employability chances, the system used a logistic regression model trained on historical data to calculate how job-seekers’ demographic features influenced their chances (such as age, education, nationality, etc.⁵). This model would produce a short-term and long-term employment score for each job-seeker. The scores would be given to the job-seeker’s counselor at the employment agency to aid in deciding about suitable support measures. Counselors could overwrite the system’s group assignment of job-seekers but would need to give a reason for doing so (Allhutter et al., 2020b; Holl et al., 2018).

Choice of use case. Algorithmic tools for the profiling of job-seekers have been deployed in various countries, including Germany (Bundesagentur für Arbeit, 2021), Austria (Allhutter et al., 2020b), Poland (Niklas et al., 2015), and the Netherlands (Desiere and Struyven, 2021). The AMS algorithm is thus an example of a more general trend: the introduction of automation and algorithmic procedures into public employment agencies to increase productivity and efficiency (Allhutter et al., 2020a). Profiling tools like the AMS algorithm assist in assessing job-seekers and resource allocation and often split the responsibility for tasks between the algorithmic system and caseworkers (Scott et al., 2022). However, the introduction of these applications also repeatedly highlights which parts of human interaction cannot be easily replaced by automation, such as “the desire to be seen as a whole human being” (Scott et al., 2022). As the AMS algorithm, therefore, conjoins typical procedures, interactions, and challenges of ADM systems in public employment, it presented a suitable use case for our study.

3.1.2. Use Case B: Health wristband for geriatric care

Description. The health wristband is an assistance tool for the long-term care of older adults, which tracks and evaluates movement and acceleration data and raises an alarm when the wearer trips or falls. The system does not require manual activation by the wearer but relies on trained machine learning models to detect events (erratic movement, heart rate anomalies, etc.) in the recorded data, making it viable for use in private homes. Further, the wristband collects data on the wearer’s sleeping, movement, and eating behavior, allowing for remote monitoring of patients. The wristband thus fulfills a two-part purpose for the care organization: First, harmful events can be detected immediately, and thus, little time is lost when dispatching emergency services. Second, patients’ data can be used for the organization of care services by making overall coordination more efficient. Systems with

⁴ AMS stands for the Public Employment Agency (Arbeitsmarktservice).

⁵ The full list of features is given in the supplementary material.

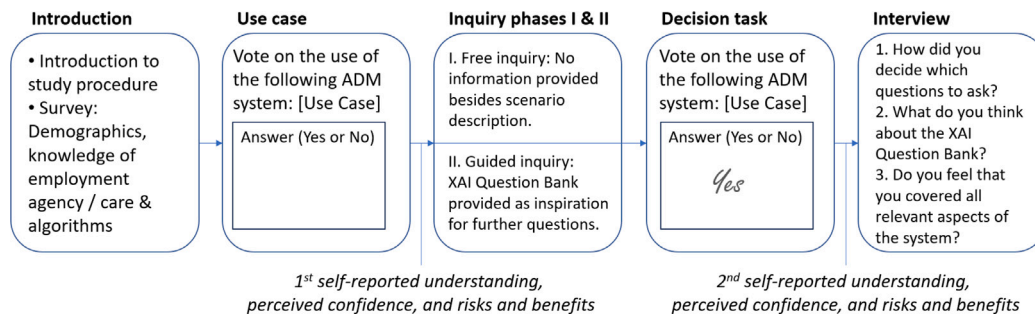


Fig. 1. Depiction of the study procedure. Participants received a short introduction to the study, filled out a questionnaire on demographic information and prior knowledge, and then were presented with one of two use cases. Participants could then ask questions about the given use case in two 15-minute inquiry phases. In the second of these phases, they received the XAI Question Bank (Liao et al., 2020) for reference. They then voted on adopting the discussed system and answered three interview questions. Participants were asked two times for their self-reported understanding, perceived decision confidence, and perceived risks and benefits of the system: before the inquiry phases and after making the decision.

comparable functionality have reportedly been employed in Sweden since 2016 (Kaun and Taranu, 2020; Sahlen et al., 2018).

Choice of use case. Wearable devices for identifying and monitoring patients are already in use in many care organizations (Kutsarova, 2020; Kaun and Taranu, 2020; Sahlen et al., 2018), such as nursing homes (Chang et al., 2023) and psychiatric wards (Nahavandi et al., 2022). Wristbands, in particular, can capture signals such as heart rate and body temperature while tracking movement and sleep (Seneviratne et al., 2017). This data can then be used to detect anomalies in breathing, physical movement, or other vital signs and trigger the appropriate reaction. This facilitates health status monitoring without keeping patients in the hospital or nursing homes, making them useful for various scenarios (Kutsarova, 2020). Besides being an example of wearable devices already well-established in care contexts, the wristband also incorporates a direct trade-off between perceived privacy and safety, as it can help detect risky situations in exchange for continuous monitoring (Kutsarova, 2020). Consequently, some users might reject these health devices due to feeling under surveillance (Nahavandi et al., 2022). A further known issue is prediction inaccuracy for certain population groups due to biases in the training data (Gerke et al., 2020). As the wristband is thus representative of known benefits and risks of wearable devices in care contexts, we selected it as the second use case for our study.

3.2. Study procedure and setup

The study procedure consisted of several steps that are depicted in Fig. 1. This section describes the use cases and decision task, the inquiry phases I and II, the question-driven explanations during the study and the reporting mechanisms for understanding, confidence, and risks and benefits.

3.2.1. Use case and decision task

At the beginning of the study, participants were informed that they would have to vote over the adoption of the ADM system after the inquiry phases, which introduced an incentive to gain information. After participants were given a consent form and a questionnaire on demographic information and previous knowledge, they were presented with the decision use case, including a brief description of the use case and its planned deployment. The description rendered the participants as part of a small citizen referendum asked to provide their vote on the system. A small box was printed below the use case description for participants to fill in “yes” or “no” after the inquiry phases, acting like a voting ballot. The design was inspired by ballots used in popular votes in Switzerland (Swiss Federal Council, 2023).⁶

⁶ Popular votes in Switzerland require the voter to write “Yes” or “No” on the ballot, which can be argued to be more deliberate compared to a tick box.

3.2.2. Inquiry phases I & II

After being briefed about the decision use case, participants had 30 min to ask the study examiner questions and receive verbal explanations about the ADM use case before making their decision. The time available to participants for inquiries about the use cases was split into two phases of 15 min. In the first phase, participants asked questions about the use case from their own intuition. In the second phase, participants were provided with a reduced version of the XAI Question Bank (Liao et al., 2020),⁷ a collection of questions on AI systems sourced from AI developers and designers, which they could use as a reference for further questions. As the original XAI Question Bank includes a comprehensive collection of questions, we reduced it by summarizing questions that aim towards similar topics. Participants were given a pen and paper to take notes during their inquiry.

3.2.3. Question-driven verbal explanations

Participants received verbal explanations from the study examiner as responses to their questions. The explanations were based on information prepared in advance and available to the study examiner during the study (as depicted in Fig. 2). For each use case, two pilot studies were conducted to collect a stock of questions that would likely be asked in the subsequent studies and which served as guidance for preparing information. When questions could not be answered due to a lack of information, the examiner asked how and why the requested information would be relevant to participants (for examples, please refer to the supplementary material).

3.2.4. Self-reported understanding, confidence, and risks & benefits

Participants reported their understanding, decision confidence, and perceptions of the system’s risks and benefits (and an explanation for their answer) at two points: first, before the inquiry phases, and second, after making the decision. Self-reported understanding and decision confidence were elicited through 5-point Likert scales. To encourage participants to speak about the trade-offs between the systems’ risks and benefits, these were split into societal and personal dimensions. For each elicitation of risks and benefits, two visual analog scales (10 cm) (Rugg and Petre, 2007) were combined to form a two-dimensional Cartesian coordinate system with the origin located at 5 cm on both axes. By denoting perceived risk on the x-axis and perceived benefit on the y-axis, the coordinate system formed four quadrants of varying risk and benefits combinations (as depicted in Fig. 4). Participants drew marks on the axes, measured, and then rounded to the first decimal place.

⁷ A detailed tabular explanation for the reduction is given in the supplementary material.

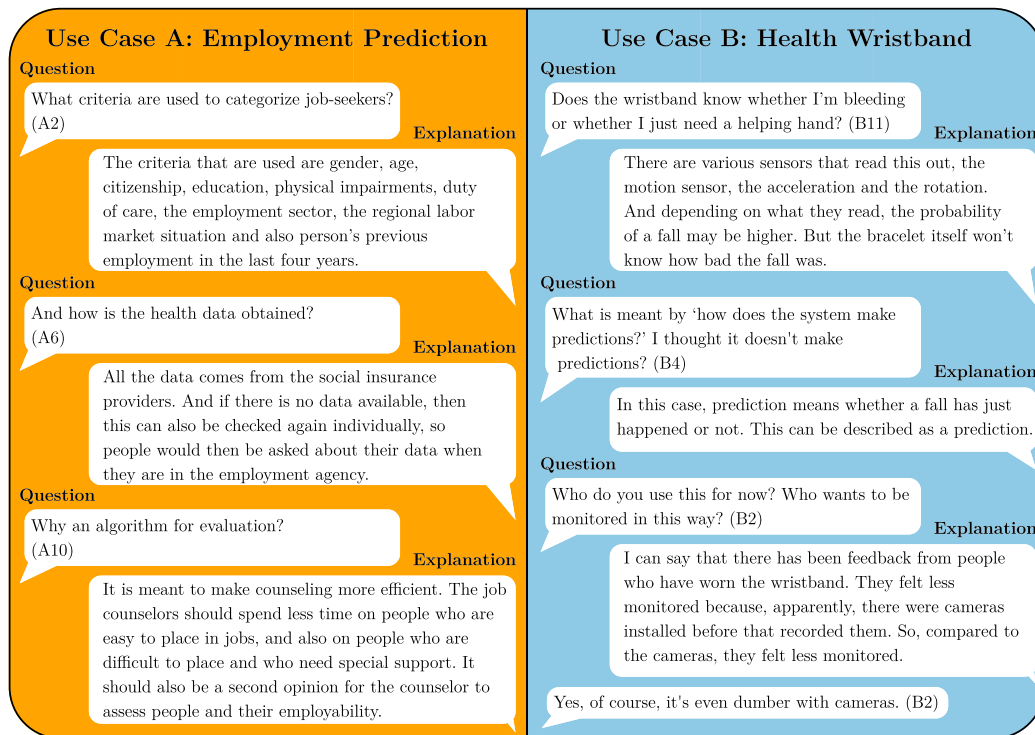


Fig. 2. Interview transcripts showing examples of question-driven explanations. The study examiner responded to participants' questions with verbal explanations during the inquiry phases I and II. All explanations were based on publicly available information about the systems and were phrased so as not to convey any personal opinion or judgment.

3.3. Interview method and analysis

Interview method. The study employed two interview techniques throughout the study procedure (Fig. 1). In inquiry phases I and II, participants were free to ask any question about the ADM use case that they found relevant. This led to a form of *unstructured interview*, in which not the examiner but the participant was driving information acquisition. At this point, the examiner only asked questions to clarify participants' inquiries and to determine why they asked for certain information. In the last part of the study, the examiner conducted a *semi-structured interview* to gather participants' responses on three key aspects: their overall reasoning and strategy for asking questions, their perception of the XAI Question Bank's (Liao et al., 2020) usefulness, and their perception of whether they acquired a complete picture of the system. At this point, the study examiner followed up on thematic aspects that arose during the interview and further connected participant responses to the quantitative items elicited before to facilitate qualitative exploration (Weiss, 1995).

Analysis. Regarding RQ1-Information Needs, we use thematic analysis (Braun and Clarke, 2006) to examine participants' information needs inductively and deductively using the XAI Question Bank (Liao et al., 2020). All questions raised by participants in the two inquiry phases were coded according to the category of inquiry. Categories already present in the XAI Question Bank (Liao et al., 2020) were assigned deductively, and topics not yet covered were defined inductively from the data. The XAI Question Bank (Liao et al., 2020) thus served as a starting point for a comparative analysis between participants' information needs and those of AI practitioners, highlighting similarities and discrepancies between the groups. A team of three coders from the authors' research group checked an intermediary version of the coding for inter-coder agreement by independently analyzing the transcripts and comparing code categories.

For RQ2-Understanding, we evaluated participants' responses on the self-reported understanding and decision confidence Likert scales and

combined this with a deductive analysis using the frameworks of "explanatory stances" (Keil, 2006) and understanding challenges (Wiggins and McTighe, 2005; Keil, 2006) (described in Section 2).

For RQ3-Perception, we analyze how explanations provided during the inquiry phases impact participants' perceptions of the systems' risks and benefits by comparing their reports on the risks and benefits coordinate system (Section 3.2.4) before and after inquiry and analyzing their corresponding articulations. Further, we analyze how participants' perceptions influence which questions are asked by comparing their reports on risks and benefits and information needs between use cases.

3.4. Participants

Recruitment. Table 1 provides an overview of the study participants. We recruited participants from two main locations: In the employment prediction use case, we collaborated with the local employment agency Job-TransFair, which established contact with job-seekers and personnel counselors. In the wristband use case, we used street sampling in an apartment complex known for the quality of living it offers to retirees.⁸ The recruitment criteria were (i) no or little previous experience with either algorithms or AI systems (reported in the questionnaire) and (ii) identification with either the domain expert or decision subject stakeholder group as described below. All interviews were conducted in person, either in office spaces or cafés, except for one conducted online. Participants were compensated with € 12 in cash, which was disclosed only after completion of the study so as not to incentivize participation solely for compensation.

Choice of participants. In both use cases, participants were selected such that they were representative for one of two roles: **domain**

⁸ While we contacted many institutions active in geriatric care work, none wanted to collaborate. The reasons for this are unknown. However, the difficulty of recruiting retirees for study participation over official channels is noteworthy.

Table 1

Details on the study participants. The ID indicates which use case participants were presented with: A = *employment prediction*, B = *health wristband*. The decision column shows how participants voted regarding the adoption of the corresponding ADM system.

ID	Age	Education	Occupation	Decision	ID	Age	Education	Occupation	Decision
A1	54	University	Personnel counselor	Yes	B1	62	A level	Retired	Yes
A2	61	Vocational university	Personnel counselor	No	B2	68	A level	Retired	Yes
A3	54	Secondary school	Job-seeking	Yes	B3	57	Secondary school	Retired	Yes
A4	43	A level	Personnel counselor	Yes	B4	46	Apprenticeship	Geriatric care	No
A5	27	Apprenticeship	Job-seeking	No	B5	80	Secondary school	Retired	Yes
A6	53	University	Personnel counselor	No	B6	68	University	Retired	Yes
A7	63	University	Personnel counselor	n/a	B7	58	University	Freelancer	Yes
A8	56	Apprenticeship	Job-seeking	No	B8	70	Vocational university	Retired	Yes
A9	50	A level	Job-seeking	No	B9	36	University	Geriatric care	Yes
A10	63	Secondary school	Cleaning service	No	B10	66	Secondary school	Retired	Yes
A11	40	A level	Personnel counselor	Yes	B11	89	Apprenticeship	Retired	Yes
A12	25	Secondary school	Job-seeking	Yes	B12	60	University	Doctor	Yes

experts or decision subjects. Following the XAI stakeholder framework (Barredo Arrieta et al., 2020; Langer et al., 2021), we define domain experts as people who are competent in the field that the ADM system is used in (e.g., employment and healthcare) and who would be its likely users. As domain experts, we recruited personnel counselors for the employment prediction use case, and care workers and a doctor for the health wristband use case. All domain experts had multiple years of experience in their field and could readily imagine how the corresponding ADM system would integrate with their work.

We describe decision subjects as people who are impacted by an automated decision and might seek ways to contest or change it but who are not necessarily users of the system (Alfrink et al., 2023; Yurrita et al., 2022). We recruited job-seekers and people who had previously been job-seeking as decision subjects for the employment prediction use case, as well as retired people and their relatives for the health wristband use case. Most of these participants already had experience with systems comparable to our use cases, such as system providing statistical competence analyses of job-seekers and wearable alarm systems for retired people in geriatric care, which, however, did not include machine learning.

4. Results

In this section, we present our results, structured according to our research questions: information needs of AI novices (RQ1, Section 4.1), participant understanding and decision confidence (RQ2, Section 4.2), and perceived risks and benefits (RQ3, Section 4.3).

4.1. RQ1-information needs: What information do AI novices who could be affected by algorithmic decisions need to decide about adopting an ADM system?

To address RQ1, we coded participants' questions raised in the inquiry phases to identify reoccurring question categories. The found categories are *system context*, *system usage*, *data*, and *system specifications*, and are each composed of several prototypical questions. An overview is provided in Fig. 3, a tabular representation in Table 2. In the following, each of the four question categories is described and split into subcategories. We denote the total number of questions asked in each subcategory in brackets.

4.1.1. System context : Intention, deployment process, consequences, development, target group, responsibility, and ethical considerations

This category captures questions that elicit background information on the system's development and deployment (69 questions total).

Intention: What is the intention of deploying it? Questions about the intention of the system's deployment were asked for both use cases (18 total), with a majority in employment prediction. These questions aimed at the underlying reasoning for introducing the system. Here we differentiate (i) practical intention, such as data collection, "How did they come up with the idea of GPS? What's the intention behind it?" (B4);

and (ii) structural intention, such as societal benefit, "What is the benefit for society? What did they expect from this algorithm?" (A1).

Deployment process: How does the deployment process work? Questions about the general deployment process were primarily asked in the employment prediction use case (13 of 19 total). This includes inquiries about how the system would be tested and how comparable systems were deployed in other countries.

Consequences: What are the consequences after deployment? Questions on what would follow from the system's deployment were asked in low counts in both use cases (9 total). Topics of inquiry included the effect on the available workforce and care ratio and tended to focus on risks and challenges:

Won't I lose labor potential if I only decide based on educational background and so on? If this algorithm is introduced, will it not lead to the loss of workforce? (A1)

Development: How is the system developed? Questions on the system's technical development emerged mainly in the employment prediction use case (8 of 9 total). They covered, e.g., the contracted company, duration of development, and the collaboration between developers and domain experts. Personnel counselors primarily asked these questions, who related the information to their experience with contracted software developers and compared it to known procedures.

Target group: Who is the system's intended target group? Questions on the system's target groups emerged mainly in the wristband use case (8 of 10 total) and covered, e.g., the target group's physical mobility and financial situation. In contrast, the two inquiries in the employment prediction use case considered target groups more on a societal level: "Would the system be applied for every population group? Can we say that every citizen would then be covered?" (A11).

Responsibility: Who is responsible for the system's deployment? These questions (4 total) addressed responsibility for the system's deployment, including the political decision, financing, project management, and technical implementation.

Ethical considerations: Is the system's deployment right? Questions in this subcategory addressed perceived social wrongs or difficult trade-offs on personal and societal levels, which might be aggravated by the ADM system's introduction (9 total). Personal topics included perceptions of unfairly distributed resources, perceived inaccessibility to own data, and over-reliance on care personnel for patients' safety ("If I know that the person falls all the time, what do I do? If I let them walk around anyway, am I then responsible?" (B1)). Societal topics included the (in)sensibleness of sanctioning job-seekers, reflections on whether an ADM system could be more profitable if it did not aim for people's well-being and the (in)validity of quantifying people's attributes ("Do you believe that you can categorize every human being in such a scientific way? By some gradations and variables?" (A4)).

XAI Novice Question Bank

System context

- ① ② What is the intention of deploying it?
- Ⓐ Ⓑ How does the deployment process work?
- Ⓓ Ⓔ What are the consequences after deployment?
- ② Ⓕ How is the system developed?
- ⑤ Who is the system's intended target group?
- ① ② Who is responsible for the system's deployment?
- ⑤ ⑥ What do other people think about the system?
- ① ⑥ Is the system's deployment right?

System usage

- ④ Ⓒ How is the system operated?
- Ⓒ Ⓓ How is the system used by and on people?
- Ⓐ Ⓑ How is the system integrated into ex. structures?
- Ⓓ Ⓔ How can the system be misused?
- ③ ⑤ How would the system handle this case?
- ③ What are the costs?

Data

- ③ What kind of data was the system trained on?*
- ③ What is the source of the training data?*
- ③ Are the data correct / representative / safe?

System specifications

- ③ What features does the system consider and why?*
- ③ What is the scope of the system's capability?*
- ③ What kind of output does the system give?*
- ③ How does the system learn?*
- ③ What is the system's overall logic?*
- ③ What kind of algorithm was used?*
- ③ What kind of mistakes is the system likely to make?*
- ③ What are the limitations of the system?*
- ③ How reliable are the predictions?*
- ③ What does [a machine learning concept] mean?

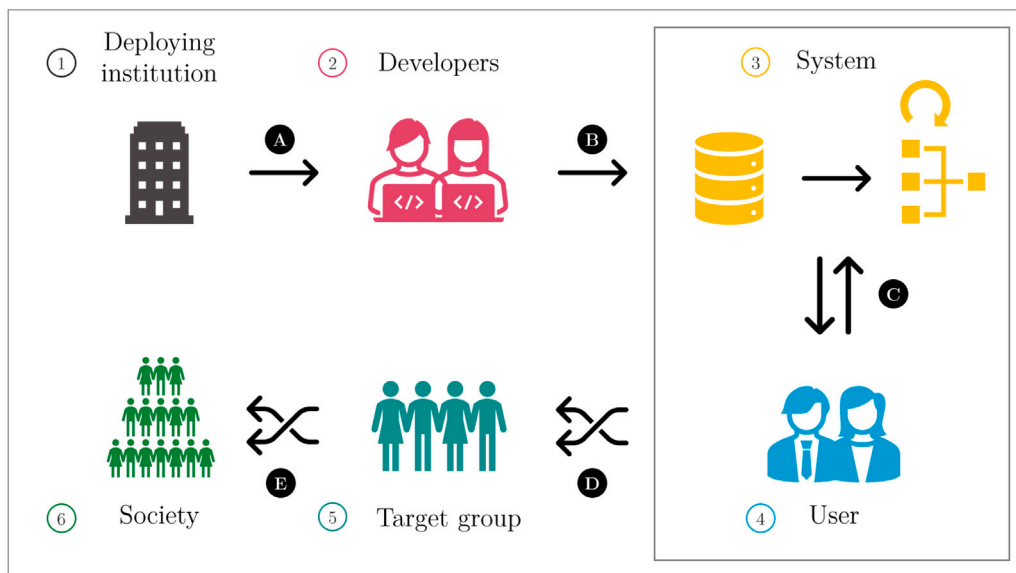


Fig. 3. The XAI Novice Question Bank and system-inquiry diagram. Depicted above are four categories of questions that subsume participants' inquiries about the two ADM use cases (Section 3.1). An asterisk (*) indicates that the question is already present in the XAI Question Bank (Liao et al., 2020). Numbers and letters are added to the side of each question to refer to stakeholders (pictograms) and procedures (arrows) in the system deployment process shown below. System and user are framed together to indicate this part of deployment as the core interaction with the system. In the system section, the arrow indicates the separation of data and model. The two-way arrows in procedure C depict interactions between the user and the system, and the twisted arrows in procedures D and E indicate the complex effects that system deployment and usage have on both the target group and society.

4.1.2. System usage : As a tool, as a sociotechnical component, as part of an organization, misuse, examples, costs

This category captures participants' information needs about how the system would be used in practice and as a concrete application. Questions on usage were asked frequently in both use cases (89 questions in total by over 20 participants).

Usage as a tool: How is the system operated? Questions on practical usage details were almost exclusively gathered in the wristband use case (30 of 31 total). These typically concern how the wristband would be used, including waterproofing, charging, operation, etc. Questions like these would likely be covered by a manual accompanying the wristband.

Usage as a sociotechnical component: How will the system impact interpersonal relations? Questions in this category asked who would use the system and how the system would influence relations between people. These were gathered mostly in the employment prediction use case (13 of 14 total). Participants' inquiries focused on how the system would affect interaction between counselor and job-seeker and covered both the counselor's and job-seeker's personal experience as well as their relationship:

Some people open up only after the second, third, or fourth conversation. [...] Here, it's really about subjective perception. What do I learn from the interview? What do I learn from the person? And what do I make of it? With the algorithm, I would need to know that this dialogue is still possible. (A2)

Usage as part of an organization: How is the system integrated into existing structures? Questions that focused on the system's interaction with existing organizational structures were present in both use cases (28 total) and covered a wide range of topics, including data processing, data access, information flow, and contestability: "What are the options for correction? Is there a system loop that handles this?" (A7).

Misuse: How can the system be misused? A noticeably low number of questions (2 total) addressed the potential of misusing the system; both emerged in the employment prediction use case. While one participant asked about using the system to commit welfare fraud, the other was concerned about the system being used by the police.

Examples: How would the system handle [this case]? Questions in this subcategory aimed to examine how the system would behave given a specific example case and emerged in low count in both use

Table 2
XAI Novice Question Bank: Tabular view. Collection of participants' questions as depicted in Fig. 3 and the type of information provided in response. An asterisk (*) indicates that the question is already present in the XAI Question Bank (Liao et al., 2020).

Category	Subcategory	Question	Information provided
System context	Intention	What is the intention of deploying it?	Design idea and underlying motivation
	Deployment process	How does the deployment process work?	Process of preparing the system for operation
	Consequences	What are the consequences after deployment?	Possible effects of deployment
	Development	How is the system developed?	Development process and institution
	Target group	Who is the system's intended target group?	Scope and features of the target group
	Responsibility	Who is responsible for the system's deployment?	Responsible actors
	Ethical considerations	Is the system's deployment right?	None, answer relies on personal assessment
System usage	Tool	How is the system operated?	Usage in practice, everyday application
	Sociotechnical component	How will the system impact interpersonal relations?	Interpersonal work, substitution of people
	Part of an organization	How is the system integrated into existing structures?	Interfaces, processes, organization's influence
	Misuse	How can the system be misused?	Scenarios of misuse
	Example	How would the system handle [this case]?	Specific example walkthrough
	Costs	What are the costs?	Financing, personal costs
Data	Type	What kind of data was the system trained on?*	Content and structure of training data
	Source	What is the source of the training data?*	Collection and storage of training data
	Attributes	Are the data correct/representative/safe?	Quality, biases, safety of training data
System specs.	Features	What features does the system consider and why?*	Information fed into system
	Functionality	What is the scope of the system's capability?*	Supported range of functions
	Output	What kind of output does the system give?*	Information received from system
	Learning	How does the system learn?*	Automated improvement of system
	Logic	What is the system's overall logic?*	General process overview
	Algorithm	What kind of algorithm was used?*	Specific model used
	Errors	What kind of mistakes is the system likely to make?*	Expected errors
	Limitations	What are the limitations of the system?*	Known boundaries
	Reliability	How reliable are the predictions?*	Dependability of output
	Reflexive	What does [a machine learning concept] mean?	ML concept information

cases (6 total). Specifically, several participants were interested in how the algorithm would handle their case.

Costs: What are the costs? Questions about costs pertained to both personal costs (“*What would it cost me?*” (B5)) as well as organizational expenses (“*If the state pays, what is the cost? Is the money well invested?*” (A4)). This subcategory appeared more often in the wristband use case (6 of 8 total).

4.1.3. Data : Type, source, attributes

This category captures questions about the data used by the system. These questions emerged in both use cases (43 questions total) and mostly occurred only after participants received the XAI Question Bank (Liao et al., 2020) as reference. The question bank already contains questions about **data type** (*What kind of data was the system trained on?*) and **data source** (*What is the source of the training data?*). In our study, multiple participants also asked about data privacy and misuse (11 questions by 8 participants). Therefore, we extend the data category with a subcategory on **data safety** (*Are the data correct/representative / safe?*).

4.1.4. System specifications : Features, functionality, performance, output, predictions, logic, reflexive inquiry

This category covers all questions focusing on specific aspects of the ADM system and has the highest question count of all categories (187 total). In contrast to the previous categories, many questions on system information emerged only after participants were provided with the XAI Question Bank (Liao et al., 2020) in the second inquiry

phase. We focus on describing questions about the system's features, functionality, performance, and reflexive inquiry, as the remaining subcategories mostly contain questions adopted verbatim from the XAI Question Bank.

Features: What features does the system consider and why? Questions on the system's input features emerged almost exclusively in the employment prediction use case (28 of 31 total). Most questions on features emerged before participants were given the XAI Question Bank (Liao et al., 2020) but mirrored questions covered in the bank, e.g., *Which features are considered?*, *How is a certain feature weighed?*, and *Why is this feature included?* This points to shared information needs between AI novices and experts, as questions on features in the question bank were sourced from stakeholders *with* technical knowledge. However, the use case appears to have a strong impact on whether features are inquired at all.

Functionality: What is the scope of the system? Can it do...? Questions on the system's functionality emerged almost exclusively in the wristband use case (30 of 34 total). The 'functionality' subcategory is close to the 'usage as a tool' subcategory but differs in its focus: Questions in this subcategory asked about functional scope in general (“*How does the wristband work?*” (B5)), whereas questions in the “usage as a tool” subcategory instead asked how the system would be used specifically and practically (“*Does that mean you can also shower with it?*” (B4)).

Errors, limitations, reliability Questions about errors, limitations, and reliability (21, 23, and 7; 51 total) emerged in both use cases equally and covered topics such as error frequency and susceptibility,

Table 3

Change in self-reported understanding and decision confidence. Participants indicated their perceived understanding of the system and confidence in their decision on a 5-point Likert scale before and after receiving question-driven explanations in the inquiry phases (Section 3.2). This table summarizes the change between their two self-reports. Notably, changes in perceived understanding did not necessarily lead to increases and decreases in confidence. Changes between reports are encoded with symbols and colors, + (increase), - (decrease), and = (no change). Columns indicate participant ID and use case.

Employment prediction use case	A6	A11	A4	A1	A8	A3	A10	A5	A2	A12	A9	A7
Δ Self-reported understanding	++	+	=	=	=	=	=	=	-	-	- - -	n/a
Δ Confidence	+	=	+	+	=	=	=	- -	=	=	=	n/a
Health wristband use case	B10	B8	B7	B2	B4	B5	B6	B9	B12	B1	B3	B11
Δ Self-reported understanding	+	+	=	=	=	=	=	=	=	-	-	-
Δ Confidence	+	-	++	+	=	=	=	=	-	++	+	+

handling incomplete data, constraints of usage scenarios, and precision. While questions on reliability only emerged after reference to the bank, questions on limitations and errors emerged organically even in the first inquiry phase: “I mean, the question is, if the system really does make mistakes, how quickly can you give the feedback: ‘Oops, that wasn’t a fall, everything’s okay.’ Is that possible?” (B4).

Reflexive inquiry: What does [a machine learning concept] mean? This subcategory captures questions about the meaning of certain notions in machine learning, which emerged in both use cases (14 total) after participants read through the XAI Question Bank (Liao et al., 2020). Questions were often prompted by terms such as *algorithm*, *biases*, and *predictions*. However, a lack of understanding in vocabulary often correlated with a lack of understanding of the corresponding concept: “What kind of data was the system trained on? I don’t understand the ‘trained’ part. That’s just an input, isn’t it?” (A3). We thus adapt the original bank’s question slightly such that it captures questions about machine learning *concepts* instead of *terminologies*.

4.2. RQ2-understanding: How do question-driven explanations interact with participants’ understanding?

This section describes how the explanations given in response to participants’ questions interacted with their understanding and decision confidence. First, we focus on the effects of question-driven explanations on self-reported understanding and decision confidence (Section 4.2.1). Second, we describe how participants’ questions can be analyzed using the framework of “explanatory stances” and which challenges in understanding arose (Section 4.2.2).

4.2.1. How question-driven explanations impact understanding and confidence

Participants reported their perceived understanding and decision confidence on 5-point Likert scales before and after the inquiry phases, i.e., before and after receiving explanations (overview depicted in Table 3).⁹ We describe three main aspects of these self-reports: participants’ overall perceived understanding, how their understanding changed throughout the study, and how their understanding interacted with their decision confidence.

Self-reported understanding was high overall, but participants define ‘high understanding’ differently. In the first self-reported understanding, before receiving explanations, only 4 of 24 participants indicated that they did not understand the ADM system well. In the second self-report, after receiving explanations, there even was only 1 participant who indicated that they did not understand the system and who perceived a substantial decrease in their understanding (“Because we talked, I realized that I don’t understand it.” (A9)). Multiple participants experienced this decrease (A2, A12, B1, B3, B11).

While the overall self-reported understanding was, therefore, high, the level of detail in their explanations for these reports differed noticeably. For example, domain expert A2 explained: “The system should aid

in deciding which measure to offer to a job-seeker. [...] The question is how the algorithm is created, but the basic classification, I think, makes sense. (A2). In contrast, some participants explained their perceived high understanding with more general observations, such as their reception of AI in the media: “There is a lot of talk about artificial intelligence in the media. Mostly at night. I watch it when I can’t sleep.” (B5). Other participants spoke about the general employment market (A5), the consequences of digitization (A9), and comparable systems (e.g., B8, A12). This illustrates a discrepancy in what people define as “high” understanding, possibly resulting from the difficulty of assessing one’s understanding (Keil, 2006) (Section 5).

Question-driven explanations mainly did not influence self-reported understanding. After proceeding through the inquiry phases, most participants reported their understanding to have remained the same; 4 participants reported increased understanding, and 6 reported declines (as seen in Table 3). As expected, increased self-reported understanding was explained through information gained in the inquiry phases. In contrast, decreases in understanding seemed to have various reasons, e.g., for A2, the decrease in understanding was accompanied by heightened scrutiny:

Interesting, I was so convinced about it at first, but now there’s something missing from my point of view. Basically, I’m still convinced that it’s a good thing, but it’s not ready yet. (A2)

Previous studies have observed that information gain can paradoxically result in decreased perceived understanding, even though objective understanding increased, likely resulting from a new awareness for unknown information (Cheng et al., 2019).

Confidence increased for 9 participants but seems not to depend on understanding. Compared to understanding, question-driven explanations increased participants’ decision confidence more often (as seen in Table 3). Several participants noted that it would increase their confidence further if they saw the system in practice: “By and large, I know how the system works. Otherwise, to be sure, I would have to see how it actually works in practice.” (A4). Participants, however, offered little explanation as to why understanding did not increase together with their confidence.

4.2.2. How question-driven explanations interacted with explanatory stances and which challenges arose

The cognitive sciences outline several processes governing how people seek and process information (described in Section 2). In the following, we analyze two main aspects of participant understanding: the role of explanatory stances and the challenges that prevented participants from acquiring a “working understanding” (Keil, 2006).

Explanatory stances offer a way of matching explanations to participants’ information needs. Explanatory stances are a way of describing how people acquire information to develop a predictive strategy of a system’s (or person’s) behavior (Dennett, 1998). To this end, people can take different strategies, which might be focused on a system’s physical attributes, on its design and functions, or, if those are inaccessible, on the “beliefs [an] agent ought to have given its place in the world and its context” (Dennett, 1998), which define its intentions and actions.

⁹ One participant decided not to vote on the use case and did not answer the second query out of concern about how the study results would be used.

To analyze if participants preferred a specific strategy of seeking information, we apply the concept of the three explanatory stances (mechanical, design, and intentional) to their questions during the inquiry phases.¹⁰ We conduct this analysis by assigning participants' questions the likeliest of the three stances. This assignment is an interpretation of their chosen stance derived from the question's phrasing and the conversation's context. Consequently, several questions might be assigned to more than one stance, depending on the interpretation of the information needs (discussed in Section 5.2.2). In the following, we describe our analysis of which stances participants applied for their questions about the ADM systems:

- **Questions about data and system specifications often applied the mechanical stance.** The mechanical stance focuses on parts, processes, and mechanisms and does not consider design or intentions (Lombrozo and Wilkenfeld, 2019). Participants tended to apply this approach to learn about the technical details of the system, such as the data basis, features, predictions, and output. This direction of inquiry often resulted from reference to the XAI Question Bank (Liao et al., 2020), which contains many questions about technical details, including "What is the sample size of the training data?" (A4), "What are the features with the evaluation criteria?" (A7), and "How sensitive is it?" (B8). However, small differences in question phrasing can lead to different explanations. Compare: "What is the sample size of the training data?" (A4) and "The sample size of the data, I assume, is large enough?" (A5). While the first one is taken verbatim from the XAI Question Bank (Liao et al., 2020) and can be said to focus on a "part" of the system (the training data), the second raises the question of what the sample size should be *large enough for*, which rather applies the design stance by asking about functional information.
- **Questions about the system's usage and context applied the design stance.** The design stance focuses on how functions and goals are realized in terms of design (Keil, 2006). Therefore, we assign participants' questions about the system's design's practical and conceptual aspects to this stance. Examples of practical design aspects include questions about usage details and operation instructions ("Do you wear it day and night?" (B1)). In contrast, examples of conceptual design aspects include questions on the motivation behind certain functionalities ("What's the benefit of knowing whether the patient was at the refrigerator at 3:00?" (B4)). Further, we assign questions about context information to this stance when they acquire information about, e.g., the characteristics of the target group ("This thing is only supposed to be distributed to care home residents, right?" (B12)) or control mechanisms ("What options for correction are there?" (A7)).
- **Questions about context and values applied the intentional stance.** The intentional stance describes a system's behavior based on beliefs and motives embedded in it, which are assumed (in the absence of mechanical and design information) to drive its actions (Dennett, 1998). We assigned participants' questions to this stance when they addressed topics such as the intentions of the deploying institutions ("Would that mean that the job counselors would need to perform better?" (A4)) and the rights of affected persons ("If I'm analyzed, I would like to have the results. Is there such a thing as a right of appeal?" (A8)). Further, the intentional stance also tended to be used when participants suspected that their values would not be considered in the system ("When this comes, who guarantees that it is not forced on you?" (B6)). Notably, many questions in the intentional stance focused on the beliefs and values of the deploying institution, which connects to literature emphasizing the importance of people's perceptions towards the institution (Cavaliere and Romeo, 2022). This showed that

participants regarded the job agency critically and is reflected in participants' perceptions of the systems' risks and benefits (Section 4.3).

Challenges in acquiring information to develop understanding. Explanations bring the incompleteness of one's understanding into focus and prompt specific reactions: filling the gap with further details, "outsourcing" it to another mind, or mentally rendering it irrelevant (Keil, 2006). The explanation setting used in this study explicitly supported filling gaps with information by letting participants flexibly inquire about any topic. However, participants also met challenges in acquiring information and developing understanding. In the following, we describe three of these challenges:

- **Blind spots inhibit further inquiry.** In the closing interviews, multiple participants expressed that they did not understand all relevant aspects of the system but could not articulate which information was missing. Reasons given for this effect included a lack of mathematical knowledge (A6), too little preparation (A8), pressure of the study situation (B5, A12), and uncertainty about the system's practical behavior (A1, A4, B1, B9): "Experience has taught me that it's always the practical application that raises issues." (B1). We call understanding gaps that could not be closed "blind spots".
- **Outsourcing understanding divides cognitive labor but risks over-dependence.** Multiple participants asked about other people's judgment of the system, e.g., the job center's employees (A8), digital natives (A9), their families (B11), and the study examiner (A11). This "outsourcing" (Keil, 2003) is a common process to divide the cognitive labor of understanding between trusted people. Critically, some participants stated intention to outsource their full understanding: "The first question is: Would you have said yes to this system?" (A9). Trust is thus a crucial part of stakeholders' understanding when outsourcing is used, especially as it risks over-dependence on the entrusted person.
- **Abandoning of understanding undermines the explanation effort.** Some participants expressed that they would likely not care about the system even if it was used on them, resulting from a lack of interest or the belief that their voice would be irrelevant: "I still think that it's not my interest that is being represented here, but actually the interest of the job agency." (A8). A pronounced disinterest in topics despite their relevance for oneself is known in political science as "civic ignorance" (Lupia, 2016).¹¹ Research has been conducted on ways to remedy this disinterest and will be briefly discussed in Section 5.

4.3. RQ3-perception: How do question-driven explanations impact perceptions of the systems' risks and benefits?

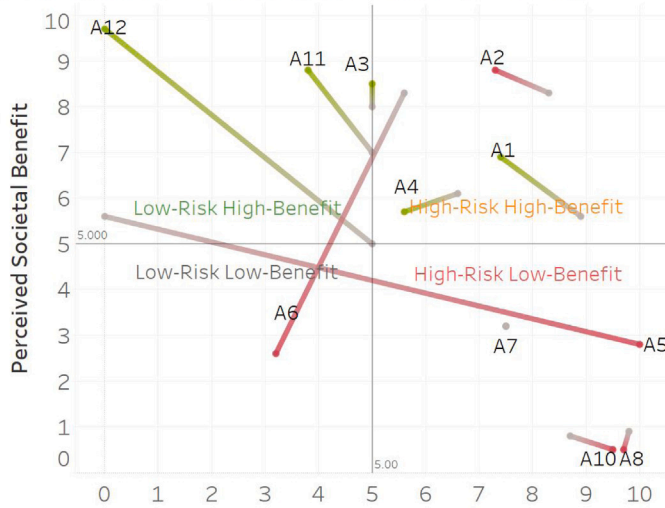
Participants were asked how they perceived the risks and benefits of the ADM system before receiving explanations and after making the deployment decision. The two-dimensional grids used to elicit these perceptions divide the space into four quadrants, corresponding to four combinations of risk and benefits (low-risk low-benefits, low-risk high-benefits, high-risk high-benefits, and high-risk low-benefits; Section 3.2.4). Changes in their perceptions between the first and second report are visualized in Fig. 4, showing solid changes in the employment prediction use case and lesser changes in the health wristband use case. In the following, we describe participants' perceptions in both use cases and outline potential reasons for the changes.

Perceptions changed enormously in employment prediction use case. Participants in the employment prediction use case showed

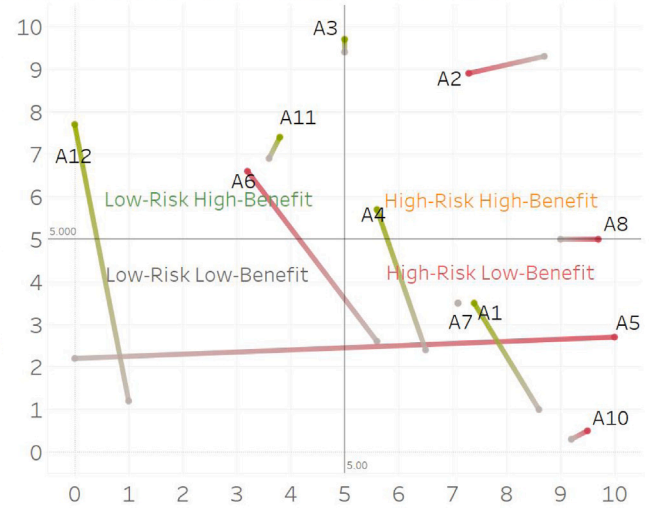
¹⁰ An introduction is provided in Section 2.

¹¹ The term describes that ignorance on public matters applies to almost everyone (Lupia, 2016).

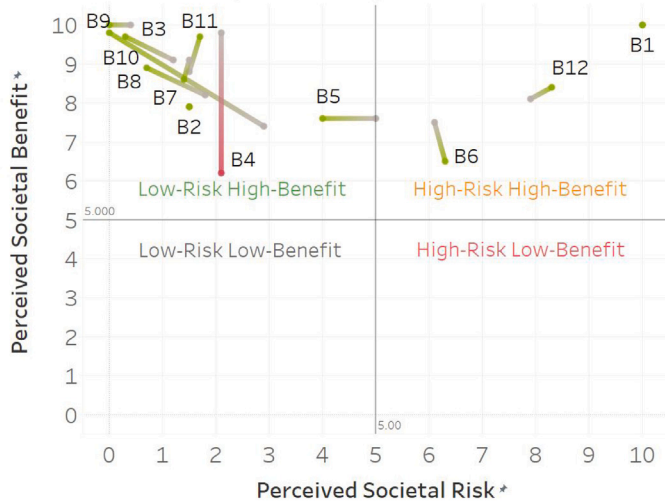
Employability Prediction (Use Case A) - Society




Employability Prediction (Use Case A) - Personal



Health Wristband (Use Case B) - Society



Health Wristband (Use Case B) - Personal



Decision

■ Yes ■ No ■ Origin

Dimensions

0 = Low risk / benefit, 10 = High risk / benefit

Fig. 4. Changes in perceived risks and benefits. Participants were asked two times for their perceptions of the given ADM system’s risks and benefits for both society and them personally (as described in Section 3) on scales from 0 (low) to 10 (high). Lines depict perception changes from before (gray) to after (colored) the inquiry phases, split between the employment prediction use case (A) above and the health wristband use case (B) below. While perceptions changed drastically in A and are distributed throughout all quadrants, perceptions in B changed comparatively little and remained mostly in the upper left quadrant.

very different initial perceptions of the system, ranging over all combinations of risks and benefits and showing little focus on specific attitudes (as seen in Fig. 4). Further, participants enormously changed their perceptions of the ADM systems after receiving explanations, often transitioning into other perception quadrants due to newfound information.

A6: What does “regional labor market” mean? That a large number of jobs are registered as vacant?

Examiner: Yes, the city was divided into regional labor market types. It was intended to show how many jobs are available, but it is also a placeholder variable for the place of residence.

A6: I find that quite discriminatory, I have to say.

Frequently mentioned benefits included a potentially reduced workload, reduced administrative overhead, and guidance for personnel

counselors, while frequently mentioned risks included a lack of care for the individual, fear of discrimination, and overreliance on algorithmic decisions. A11, for example, reported an increase in perceived societal benefit, explaining:

I believe that the state could save money in the broadest sense. And this money could be used to plug other holes. [...] It would only be a risk because soft skills and the emotional part are not included. That's the only risk I can see. (A11)

In contrast, A2 reported a slight decrease in perceived risk and varying changes in benefit while also voting “no”, stating:

What I'm missing here is that no consideration is given to whether the person looking for work can continue in their core business or needs a new career path. [...] I miss the jobseeker's personal perspective. The hard facts are the smaller issue for me. (A2)

We observed that question-driven explanations had mixed effects on participants' perceptions in the employment prediction use case. While the provided information indeed enabled a shift in perspective, the trajectory of that shift seems to differ heavily, even though the information provided did not change between participants.

Perceptions changed less in the health wristband use case. Participants in the health wristband use case saw the system clearly as having low risk and high benefit. Changes in perceptions were minor and never resulted in transitions between perception quadrants, showing that participants perceived the system positively before and after receiving explanations.

Through the question and answer sessions, my idea of the wristband has become more concrete. I realized that the thing can do other things than I had imagined. But with risks and benefits, I'm sticking with what I had before, that hasn't changed. (B7)

For some participants, the wristband presented a trade-off between the values of safety and privacy, which they weighed against each other. Despite little overall change in perceptions, this highlights that participants considered the risk of surveillance and the benefit of the wearer's safety and came to different conclusions. B11, for example, personally did not see any issues, while B4 specifically mentioned the feeling of being surveyed:

For me, I don't see any disadvantages. But I could imagine that there are people who say: "I don't want to be monitored. Even if I fall down and nobody comes and I bleed to death, then I'll just bleed to death." (B11)

I would rate the risk as high because I believe it's going in the direction of surveillance. (B4)

In summary, participants thus weighed the given information and decided based on it, which suggests that the explanations fulfilled their purpose in revealing potential value conflicts in the system's deployment. Making this conflict explicit, B12 stated:

The risk of data collection is a risk inherent to the system. Either you accept this risk and have the benefit of these systems, or you do without it. You probably can't have one without the other. (B12)

5. Discussion

This section outlines the theoretical and practical implications of our study. It addresses AI novices' information needs, the interaction between question-driven explanations and understanding, and the inclusion of domain experts and decision subjects in explanation design. These implications are summarized at the end (Section 5.3).

5.1. Theoretical implications

5.1.1. Designing explanations for AI novices

One of XAI's enduring key questions is: What makes a "good" explanation? Recent work outlines that, to produce understanding, explanations should be contrastive, selected, and social while being first and foremost *contextual* (Miller, 2019), that they should interface directly with people's causal reasoning (Byrne, 2023) and consider the learner's aim and prior knowledge (Mueller et al., 2019). Fulfilling these requirements puts a high demand on explanation design. Further, existing explanation approaches do not typically address the needs of AI novices (Shin, 2023), even though they compose multiple stakeholder groups, including people affected by algorithmic decisions and end-users. We thus argue that one step in designing explanations for AI novices is to identify their information needs and examine how they depend on participants' perceptions, stakeholder roles, and domain expertise. In the following, we present how the findings of our study support this design approach.

Our findings show that **AI novices' information needs cover all parts of an ADM system's "lifecycle"** (Dhanorkar et al., 2021), from technical details over deployment consequences to different dimensions of usage. Our proposed "XAI Novice Question Bank" (Fig. 3) provides a summary of these topics, which can inform future explanation design. However, we also observe differences in information needs between use cases and differences to Liao et al. (2020)'s study with AI practitioners. This is in line with previous work, which shows that various contextual factors mediate which information needs participants articulate, including familiarity with the presented system (Kramer et al., 2018), tangibility of the system (Long et al., 2021), perceived risks and benefits (Araujo et al., 2020) (e.g., trade-off between safety and privacy in the wristband use case), domain expertise (Wang and Yin, 2021), and technical expertise (Cheng et al., 2019).

5.1.2. Anticipating and handling explanations of intention

Explanations have the purpose of enabling stakeholders to attain their aims (Langer et al., 2021), such as facilitating the evaluation of an ADM system's values (Shin, 2021, 2023), or performing an action, like contesting a decision or requesting human intervention (Alfrink et al., 2023). In our study, participants were asked to vote over adopting the ADM systems to incentivize information acquisition through questions. I.e., the "purpose" (Freiesleben and König, 2023) of the explanations was to enable a well-informed decision. For several participants, this succeeded, as they used the explanations to underpin their decision and reported a slight increase in decision confidence (Table 3).

At the same time, as participants' information needs differed between use cases, their perception of what a well-informed decision would need to consider differed. Participants tended to inquire more about societal and structural factors in the employment prediction use case and more about practical and operational aspects in the wristband use case. Framed in the context of the explanatory stances, many participants used the "intentional stance" (Dennett, 1998) to ask questions about the *AMS algorithm* and the beliefs and values it represents (Section 4.2). This demand for explanations of intention was likely due to the system's close association with the job agency, as previous work has shown that the reputation of the deploying institution can affect perceptions of ADM systems (Brown et al., 2019; Woodruff et al., 2018; Lee et al., 2019).

However, delivering explanations of the intention behind a system's deployment is difficult, as it means interpreting the values involved and choosing how to present them — a process that rather should be performed by each participant individually due to its inherent subjectivity. Further, information about the true purpose of an ADM system's deployment can be hard to attain, as it is often part of a deploying institution's non-public decision-making process and can be kept intentionally opaque (Rudin, 2019).

Consequently, **whether explanations of intention should be included needs to be matched to the explanation's purpose.** For example, if the explanation's purpose is solely a mechanical or functional description of the system to enable operation, intentional information is likely negligible. In contrast, when stakeholders should form value assessments or investigate possibilities for contestation (Alfrink et al., 2023), intentional explanations can contribute meaningful information, e.g., by using "justifications" (Biran and Cotton, 2017; Biran and McKown, 2017; de Fine Licht and de Fine Licht, 2020) to explain how decisions are made. As stakeholders' interest in intentional information was central to our study, future research should examine how this information can be conveyed meaningfully.

5.1.3. Taking steps to connect explainability to civic education

Our studies show that our participants were not always interested in understanding ADM systems, even though they might affect them. Assuming that democratic civil societies have the right to understand public AI systems (Züger and Asghari, 2023) and that explanations are a means to realize this right (de Fine Licht and de Fine Licht, 2020), the

question of how to redeem the interest of these stakeholders becomes essential. Research on democratic participation uses the term “civic ignorance”, which is the common phenomenon that citizens are disinterested in laws and regulations shaping their lives (Lau and Redlawsk, 2001; Lupia, 2016). Civic ignorance is not a stigma but an effect that applies to most people, including political experts (Lupia, 2016). Lupia (2016) describes that interest can be remedied by developing educational material that covers the “information that matters”, which is the intersection of (i) information that advances the core concerns of the public and (ii) information that enables the performance of a task, such as judging a political topic.

This paper aims to accomplish this step of identifying information relevant to people affected by algorithmic decisions and sufficient to form a judgment. We argue that future work should combine these empirical insights with conceptual work on AI literacy (Long and Magerko, 2020; Ng et al., 2021) to design explanations that effectively convey information to these stakeholder groups. Our study outlines the first steps towards this approach: Orientating explanations along AI novices’ information needs, allowing for flexible switching between different topics and explanatory stances, using guidance material such as the XAI Question Bank (Liao et al., 2020) and the XAI Novice Question Bank (Fig. 3), and considering to include contextual information such as intention.

5.2. Practical implications

5.2.1. Resolving understanding challenges in question-driven explanations

The direct aim of explanations is to increase a person’s understanding of an ADM system (Langer et al., 2021), but both designing explanations that reliably increase understanding and evaluating if they succeeded is surprisingly difficult (Schmude et al., 2023). In the following, we analyze how our question-driven explanations interacted with participant understanding and derive two concrete implications for future explanation design.

First, **question-driven explanations succeed in letting participants independently select relevant information and guiding them to non-explored aspects of the ADM system**, for example, by providing the XAI Question Bank (Liao et al., 2020) as guidance. This supports the notion of personalized explanations (Conati et al., 2021; Shulner-Tal et al., 2022; Martijn et al., 2022) and could be further elaborated in a design study of a digital question-driven explanation tool based on the XAI Novice Question Bank’s categories. This would allow for a better overview and more straightforward navigation of information while potentially being suitable and relevant for non-expert users. This implementation could further explore the possibility of including a generative model using an information base to generate explanations in response to users’ input queries, as has already been proposed for other domains (Zhao et al., 2024).

While this generative explainer would allow for flexible and comprehensive explanations that could vary in detail, it also poses several challenges: First, the information used for our explanations was “scavenged” (Wieringa, 2023) from various sources, meaning that it needs to be decided which information should be prioritized if sources are ambiguous or contradict each other. Second, non-factual information, such as explanations of intention, would require the model to reflect multiple points of view clearly and neutrally represent value conflicts likely to arise in high-risk ADM systems (Raji et al., 2022). Third, non-deterministic generative models are known to produce content that can vary in quality and might even contain wrong information (Karinhshak et al., 2023), which would be very difficult to spot for AI novices. Lastly, multiple participants in our study reported that they appreciated the social nature of the explanation format and the exchange with another person, which is in line with previous research (Miller, 2019) and which would be mostly lost when transferred to a generative explainer. Providing a viable implementation of generative question-driven explanations would thus require thoroughly considering these issues.

Second, **several processes can interfere with increasing understanding through question-driven explanation, including blind spots, outsourcing, and understanding self-assessment**. We use the term *blind spots* for information that participants felt was missing but that they could not acquire. This might mean a lack of “intelligibility types” (Lim and Dey, 2009), i.e., a lack of guidance on which information was available and how to ask about it. Documents such as our proposed XAI Novice Question Bank (Fig. 3) and the XAI Question Bank (Liao et al., 2020) should help in addressing this issue by showing people the relevant categories of information.

Outsourcing, on the other hand, refers to sharing the process of understanding with other people to make sense of it collaboratively, a topic that seems underexplored in current XAI research. Future work should thus consider if the effectiveness of explanations can be increased by providing them to groups of people who can rely on the “expertise in other minds” (Keil, 2006), i.e., discuss and collaborate. An essential aspect of outsourcing understanding is the reliance on other people to aid in the understanding process (Keil, 2006), especially when the topic seems complex and overloaded. In our study, this prompted participants to skip parts of the understanding process entirely and instead ask the study examiner for their assessment. Considering that civil society organizations are already entrusted with providing support for people regarding finances, education, legal aid, and more, they are an intuitive point of orientation when it comes to providing explanations about high-risk ADM systems to those affected by their decisions (Scott et al., 2022). We argue that future work in XAI should thus cooperate with these organizations to deploy explanations that enable AI novices to understand, negotiate, and contest decisions of ADM systems (Kaun and Taranu, 2020).

Lastly, question-driven explanations also require *understanding self-assessment*, meaning people decide when their understanding is sufficiently developed. However, this can be difficult to assess (Keil, 2006), as perceived understanding can paradoxically decrease despite a gain in information (Cheng et al., 2019) and can also fall prone to the “illusion of explanatory depth” (Rozenblit and Keil, 2002; Chromik et al., 2021), an overestimation of one’s understanding. Previous work has outlined various strategies to offset these issues, including deliberate self-explanation (Duckworth, 2001), better initial introductions to the system (Chromik et al., 2021), and checking for people’s factual understanding by asking examination questions (Buçinca et al., 2020). To offset the decrease in self-reported understanding stemming from the gain in new information, we propose to include an elicitation of participants’ information gain in addition to their perceived understanding. A survey question could be formulated as “How much knowledge did you gain about the system?”.

5.2.2. Matching explanations to stakeholders’ needs

For the reasons outlined, people can take very different approaches to information acquisition and understanding, depending partly on their stakeholder role and aims. As explanatory stances (mechanical, design, and intentional (Dennett, 1998)) describe how people make sense of information and build predictive strategies from it, they can inform explanation design by considering which stance people apply when acquiring information. While previous XAI research explored explanatory stances conceptually (Páez, 2019; Miller, 2019; Zerilli, 2022), we argue that future work should investigate explanatory stances empirically for two main reasons:

First, our analysis shows that participants tended to apply different stances in their questions depending on the topic of inquiry: questions about data and system details tended to aim for technical information (mechanical), while questions about usage and context instead aimed for information on the system’s practical and conceptual design aspects (design) or the values that it might enforce due to how and why it was developed (intentional). Importantly, participants can assume any of the three stances in their questions, and similar-sounding questions might aim for different information, such as “How often does the

system make mistakes?”. **Identifying a question’s explanatory stance can thus help match the provided information to the person’s information need** by considering the question’s context or disambiguating which category of information it aims for. This identification can be facilitated by context information, such as the participants’ previous questions, or by explicitly inquiring how the information is relevant to the participant.

Second, **encouraging the use of different explanatory stances can support a change in perspective**. Work in the cognitive sciences, for example, describes the possibility that the design stance might be preferred over the mechanical stance by default (Lombrozo and Wilkenfeld, 2019).¹² In our study, we observed that participants used the design and intentional stance instead of the mechanical stance, especially before receiving the XAI Question Bank (Liao et al., 2020). However, the boundaries might not be as clear-cut due to the interpretation required when assigning explanatory stances to participants’ questions. Still, considering which information might prompt participants to take a specific stance could be useful for the design of explanations, for example, by providing material that prompts a variety of stances and steers attention to information that is not explored autonomously. For example, introducing the XAI Question Bank (Liao et al., 2020) in the second inquiry phase of our study encouraged participants to ask questions aiming for technical details (mechanical), which they reported to be very useful as they had not considered these perspectives before.

5.3. Summary of implications

In this section, we briefly summarize the implications that our findings and discussion have for the design of explanations.

Theoretical

1. AI novices’ information needs vary by use case. The proposed XAI Novice Question Bank (Fig. 3) expands the XAI Question Bank (Liao et al., 2020) with questions on system context and usage, highlighting relevant information for novices. Future work should explore integrating this information into formats that can be flexibly adapted to different systems.
2. Information on *why* an ADM system is deployed can be relevant for various stakeholders. Typical explanations focus on the system’s workings but not on this contextual information. XAI research should consider how to best convey information on the intention behind a system’s development and deployment.
3. Researchers developing explanations for those affected by algorithmic decisions should note that people might be disinterested in explanations due to a perceived inability to understand ADM or feelings of powerlessness. Future research should further examine how these factors impact explanations for groups who are affected by ADM systems.

Practical

1. Question-driven explanations let stakeholders acquire information until their understanding meets their aims. A digital implementation with a generative explainer model could apply these explanations practically. Challenges to understanding can be addressed by guiding available information, considering if understanding can be outsourced to others, and eliciting information gain in addition to perceived understanding.
2. Information can be understood from different angles, and some may be more intuitive for specific stakeholders: domain experts might ask about a system’s design aspects. In contrast, job-seekers might ask about the beliefs governing its deployment.

Explanatory stances can help to identify which angle someone will take, allowing for better matching of explanations to their needs. Encouraging stance-switching, such as providing technical information to AI novices, can lead to valuable perspective changes.

6. Limitations

Our participants were recruited from a local job agency and an apartment complex from the same geographical region, perhaps resulting in regional or cultural biases. Due to the limited sample size, we did not analyze the impact of sex and/or gender on our results. While this limits the generalizability of our results regarding sex and/or gender aspects, it does not limit the overall validity of our results. Further, the articulated information needs and observed processes of understanding are tied to the presented use cases, as demonstrated in Section 4.2, and to the verbal exchange with the study examiner, meaning that participants would likely ask different questions if the information was presented in another fashion or by another person. Concerning the size of the participant sample, we are guided by research on qualitative methods, which suggests that 16 to 24 interviews are needed to reach adequate code and meaning saturation (Hennink et al., 2017).

7. Conclusion

The use of public AI systems in society can have a considerable impact on people affected by algorithmic decisions. These stakeholders should be able to understand AI systems in order to take action, such as contesting decisions or asking for human intervention. Explanations can support this empowerment, but the information needs of this stakeholder group – who tend to be AI novices – have rarely been covered in explanation approaches. To address this gap, we provide a collection of information needs titled the “XAI Novice Question Bank” (Fig. 3), compiled from interviews with 24 participants and covering two ADM use cases (employment prediction and a health wristband). Our approach aims to extend the concept of question-driven explanations, inspired by the XAI Question Bank (Liao et al., 2020), to lay audiences and highlights the relevance of information needs about an ADM system’s context and usage. Further, we analyze participant understanding by applying cognitive theories to examine from which perspective participants approach information acquisition. We provide suggestions on matching explanations to how stakeholders make sense of information and strategies to overcome understanding challenges. Lastly, we examine participants’ perceptions of the systems’ risk and benefits, finding that higher perceived risk likely prompts participants to inquire about intentional information, such as the motivation and consequences of deploying a system. We close with five critical implications that our findings have for the design of explanations for AI novices affected by algorithmic decisions. More research is needed to create suitable explanations for this crucial stakeholder group.

CRedit authorship contribution statement

Timothée Schmude: Writing – original draft, Visualization, Methodology, Investigation, Conceptualization. **Laura Koesten:** Writing – review & editing, Supervision, Methodology. **Torsten Möller:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition. **Sebastian Tschitschek:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

¹² Preference in explanatory stances is a controversial discussion (Keil, 2021; Lombrozo and Wilkenfeld, 2019; Kelemen et al., 2013).

Data availability

The authors do not have permission to share data.

Acknowledgments

This work has been funded by the Vienna Science and Technology Fund (WWTF) [10.47379/ICT20058] as well as [10.47379/ICT20065]. We further thank the Vienna Job-TransFair team for their valuable contributions.

Appendix A. Supplementary data


Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.ijhcs.2024.103380>.

References

- Adadi, A., Berrada, M., 2018. Peeking inside the black-box: A survey on explainable artificial intelligence (XAI). *IEEE Access* 6, 52138–52160. <http://dx.doi.org/10.1109/ACCESS.2018.2870052>.
- Alfrink, K., Keller, I., Kortuem, G., Doorn, N., 2023. Contestable AI by design: Towards a framework. *Minds Mach.* 33 (4), 613–639. <http://dx.doi.org/10.1007/s11023-022-09611-z>.
- Allhutter, D., Cech, F., Fischer, F., Grill, G., Mager, A., 2020a. Algorithmic profiling of job seekers in Austria: How austerity politics are made effective. *Front. Big Data* 3, <http://dx.doi.org/10.3389/fdata.2020.00005>.
- Allhutter, D., Mager, A., Cech, F., Fischer, F., Grill, G., 2020b. Der AMS-Algorithmus: Eine Soziotechnische Analyse des Arbeitsmarktchancen-Assistenz-Systems (AMAS). Technical Report, [epub.oew.ac.at](http://pub.oew.ac.at).
- Ananny, M., Crawford, K., 2018. Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability. *New Media Soc.* 20 (3), 973–989. <http://dx.doi.org/10.1177/146144816676645>.
- Araujo, T., Helberger, N., Kruikeimer, S., de Vreese, C.H., 2020. In AI we trust? Perceptions about automated decision-making by artificial intelligence. *AI Soc.* 35 (3), 611–623. <http://dx.doi.org/10.1007/s00146-019-00931-w>.
- Barredo Arrieta, A., Díaz-Rodríguez, N., Del Ser, J., Bannetot, A., Tabik, S., Barbado, A., Garcia, S., Gil-Lopez, S., Molina, D., Benjamins, R., Chatila, R., Herrera, F., 2020. Explainable artificial intelligence (XAI): Concepts, taxonomies, opportunities and challenges toward responsible AI. *Inf. Fusion* 58, 82–115. <http://dx.doi.org/10.1016/j.inffus.2019.12.012>.
- Baumberger, C., Beisbart, C., Brun, G., 2017. What is understanding? An overview of recent debates in epistemology and philosophy of science. In: Grimm, S., Baumberger, C., Ammon, S. (Eds.), *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*. Routledge, pp. 1–34.
- Bertrand, A., Eagan, J.R., Maxwell, W., 2023. Questioning the ability of feature-based explanations to empower non-experts in robo-advised financial decision-making. In: *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency. FAccT '23*, Association for Computing Machinery, New York, NY, USA, pp. 943–958. <http://dx.doi.org/10.1145/3593013.3594053>.
- Biran, O., Cotton, C.V., 2017. Explanation and justification in machine learning : A survey. In: *Proceedings of the IJCAI 2017 Workshop on Explainable Artificial Intelligence. XAI*, cs.columbia.edu/~orb/papers/xai_survey_paper_2017.pdf.
- Biran, O., McKeown, K., 2017. Human-centric justification of machine learning predictions. In: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence. IJCAI-17*, pp. 1461–1467. <http://dx.doi.org/10.24963/ijcai.2017.202>.
- Bove, C., Aigrain, J., Lesot, M.-J., Tijus, C., Detyniecki, M., 2022. Contextualization and exploration of local feature importance explanations to improve understanding and satisfaction of non-expert users. In: *Proceedings of the 27th International Conference on Intelligent User Interfaces. IUI '22*, Association for Computing Machinery, New York, NY, USA, pp. 807–819. <http://dx.doi.org/10.1145/3490099.3511139>.
- Braun, V., Clarke, V., 2006. Using thematic analysis in psychology. *Qualitat. Res. Psychol.* 3 (2), 77–101. <http://dx.doi.org/10.1191/1478088706qp063oa>.
- Brown, A., Chouldechova, A., Putnam-Hornstein, E., Tobin, A., Vaithianathan, R., 2019. Toward algorithmic accountability in public services: A qualitative study of affected community perspectives on algorithmic decision-making in child welfare services. In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. ACM, Glasgow Scotland UK*, pp. 1–12. <http://dx.doi.org/10.1145/3290605.3300271>.
- Buçinca, Z., Lin, P., Gajos, K.Z., Glassman, E.L., 2020. Proxy tasks and subjective measures can be misleading in evaluating explainable AI systems. In: *Proceedings of the 25th International Conference on Intelligent User Interfaces. ACM, Cagliari Italy*, pp. 454–464. <http://dx.doi.org/10.1145/3377325.3377498>.
- Bundesagentur für Arbeit, 2021. Bearbeiten von bewerberdaten durch träger. arbeitsagentur.de/datei/dok_ba013193.pdf.
- Burns, A., Lee, C., Chawla, R., Peck, E., Mahyar, N., 2023. Who do we mean when we talk about visualization novices? In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. CHI '23*, Association for Computing Machinery, New York, NY, USA, <http://dx.doi.org/10.1145/3544548.3581524>.
- Byrne, R.M., 2023. Good explanations in explainable artificial intelligence (XAI): Evidence from human explanatory reasoning. In: *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence. Macau, SAR China*, pp. 6536–6544. <http://dx.doi.org/10.24963/ijcai.2023/733>.
- Capel, T., Brereton, M., 2023. What is human-centered about human-centered AI? A map of the research landscape. In: *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. ACM, Hamburg Germany*, pp. 1–23. <http://dx.doi.org/10.1145/3544548.3580959>.
- Cavaliere, P., Romeo, G., 2022. From poisons to antidotes: Algorithms as democracy boosters. *Eur. J. Risk Regulat.* 13 (3), 421–442. <http://dx.doi.org/10.1017/err.2021.57>, URL https://www.cambridge.org/core/product/identifier/S1867299X2100057X/type/journal_article.
- Chang, F., Östlund, B., Kuoppamäki, S., 2023. Domesticating social alarm systems in nursing homes: Qualitative study of differences in the perspectives of assistant nurses. *J. Med. Internet Res.* 25, e44692. <http://dx.doi.org/10.2196/44692>.
- Cheng, H.-F., Wang, R., Zhang, Z., O'Connell, F., Gray, T., Harper, F.M., Zhu, H., 2019. Explaining decision-making algorithms through UI: Strategies to help non-expert stakeholders. In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. ACM, Glasgow Scotland UK*, pp. 1–12. <http://dx.doi.org/10.1145/3290605.3300789>.
- Chromik, M., Eiband, M., Buchner, F., Krüger, A., Butz, A., 2021. I think I get your point, AI! the illusion of explanatory depth in explainable AI. In: *26th International Conference on Intelligent User Interfaces. ACM, College Station TX USA*, pp. 307–317. <http://dx.doi.org/10.1145/3397481.3450644>.
- Conati, C., Barral, O., Putnam, V., Rieger, L., 2021. Toward personalized XAI: A case study in intelligent tutoring systems. *Artificial Intelligence* 298, 103503. <http://dx.doi.org/10.1016/j.artint.2021.103503>.
- Corvite, S., Roemmich, K., Rosenberg, T.I., Andalibi, N., 2023. Data subjects' perspectives on emotion artificial intelligence use in the workplace: A relational ethics lens. *Proc. ACM Human-Computer Interact.* 7 (CSCW1), 1–38. <http://dx.doi.org/10.1145/3579600>.
- de Fine Licht, K., de Fine Licht, J., 2020. Artificial intelligence, transparency, and public decision-making: Why explanations are key when trying to produce perceived legitimacy. *AI Soc.* 35 (4), 917–926. <http://dx.doi.org/10.1007/s00146-020-00960-w>.
- Dennett, D.C., 1998. *The intentional stance*, 7. printing In: *A Bradford book*, MIT Press, Cambridge, Mass..
- Dennett, D., 2006. *Intentional systems theory*. In: McLaughlin, B., Beckermann, A., Walter, S. (Eds.), *The Oxford Handbook of Philosophy of Mind*. Oxford University Press.
- Desiere, S., Struyven, L., 2021. Using artificial intelligence to classify jobseekers: The accuracy-equity trade-off. *J. Soc. Policy* 50 (2), 367–385. <http://dx.doi.org/10.1017/S0047279420000203>.
- Dhanorkar, S., Wolf, C.T., Qian, K., Xu, A., Popa, L., Li, Y., 2021. Who needs to know what, when?: Broadening the explainable AI (XAI) design space by looking at explanations across the AI lifecycle. In: *Designing Interactive Systems Conference 2021. ACM, Virtual Event USA*, pp. 1591–1602. <http://dx.doi.org/10.1145/3461778.3462131>.
- Duckworth, E. (Ed.), 2001. *"Tell me more": Listening to learners explain*. Teachers College Press, New York.
- Ehsan, U., Liao, Q.V., Muller, M., Riedl, M.O., Weisz, J.D., 2021. Expanding explainability: Towards social transparency in AI systems. In: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems. ACM, Yokohama Japan*, pp. 1–19. <http://dx.doi.org/10.1145/3411764.3445188>.
- Ehsan, U., Saha, K., De Choudhury, M., Riedl, M.O., 2023. Charting the sociotechnical gap in explainable AI: A framework to address the gap in XAI. *Proc. ACM Hum.-Comput. Interact.* 7 (CSCW1), <http://dx.doi.org/10.1145/3579467>.
- Ehsan, U., Wintersberger, P., Liao, Q.V., Watkins, E.A., Manger, C., Daumé Iii, H., Riene, A., Riedl, M.O., 2022. Human-centered explainable AI (HCXAI): Beyond opening the black-box of AI. In: *CHI Conference on Human Factors in Computing Systems Extended Abstracts. ACM, New Orleans LA USA*, pp. 1–7. <http://dx.doi.org/10.1145/3491101.3503727>.
- European Commission, 2021. Laying down harmonised rules on artificial intelligence. digital-strategy.ec.europa.eu/en/library/proposal-regulation-laying-down-harmonised-rules-artificial-intelligence.
- European Parliament. Directorate General for Parliamentary Research Services, 2019. Understanding algorithmic decision-making: opportunities and challenges. Publications Office, LU, <http://dx.doi.org/10.2861/536131>.
- Freiesleben, T., König, G., 2023. Dear XAI community, we need to talk!. In: Longo, L. (Ed.), *Explainable Artificial Intelligence*. Springer Nature Switzerland, Cham, pp. 48–65. http://dx.doi.org/10.1007/978-3-031-44064-9_3.
- Gerke, S., Minssen, T., Cohen, G., 2020. Chapter 12 - ethical and legal challenges of artificial intelligence-driven healthcare. In: Bohr, A., Memarzadeh, K. (Eds.), *Artificial Intelligence in Healthcare*. Academic Press, pp. 295–336. <http://dx.doi.org/10.1016/B978-0-12-818438-7.00012-5>.

- Golpayegani, D., Pandit, H.J., Lewis, D., 2023. To be high-risk, or not to be—Semantic specifications and implications of the AI act's high-risk AI applications and harmonised standards. In: 2023 ACM Conference on Fairness, Accountability, and Transparency. ACM, Chicago IL USA, pp. 905–915. <http://dx.doi.org/10.1145/3593013.3594050>.
- Grimm, S.R., 2019. Varieties of understanding. Oxford University Press, pp. 1–14. <http://dx.doi.org/10.1093/oso/9780190860974.003.0001>.
- Henman, P., 2020. Improving public services using artificial intelligence: possibilities, pitfalls, governance. Asia Pacific J. Public Administration 42 (4), 209–221. <http://dx.doi.org/10.1080/23276665.2020.1816188>.
- Hennink, M.M., Kaiser, B.N., Marconi, V.C., 2017. Code saturation versus meaning saturation: How many interviews are enough? Qualit. Health Res. 27 (4), 591–608. <http://dx.doi.org/10.1177/1049732316665344>.
- Holl, J., Kernbeiß, G., Wagner-Pinter, M., 2018. Das AMS-arbeitsmarktchancen-modell. Jakesch, M., Bućinca, Z., Amershi, S., Olteanu, A., 2022. How different groups prioritize ethical values for responsible AI. In: 2022 ACM Conference on Fairness, Accountability, and Transparency. ACM, Seoul Republic of Korea, pp. 310–323. <http://dx.doi.org/10.1145/3531146.3533097>.
- Karinshak, E., Liu, S.X., Park, J.S., Hancock, J.T., 2023. Working with AI to persuade: Examining a large language model's ability to generate pro-vaccination messages. Proc. ACM Human-Comput. Interact. 7 (CSCW1), 1–29. <http://dx.doi.org/10.1145/3579592>.
- Kaun, A., Taranu, G., 2020. Automating society report 2020 / Sweden. automatingsociety.algorithmwatch.org/report2020/sweden/.
- Keil, F.C., 2003. Folkscience: Coarse interpretations of a complex reality. Trends in Cognitive Sciences 7 (8), 368–373. [http://dx.doi.org/10.1016/S1364-6613\(03\)00158-X](http://dx.doi.org/10.1016/S1364-6613(03)00158-X).
- Keil, F.C., 2006. Explanation and understanding. Annu. Rev. Psychol. 57 (1), 227–254. <http://dx.doi.org/10.1146/annurev.psych.57.102904.190100>.
- Keil, F.C., 2021. The challenges and benefits of mechanistic explanation in folk scientific understanding. In: Wilkenfeld, D.A., Samuels, R. (Eds.), Advances in Experimental Philosophy of Science, Paperback In: Advances in experimental philosophy, Bloomsbury Academic, London New York Oxford New Delhi Sydney, pp. 41–56.
- Kelemen, D., Rottman, J., Seston, R., 2013. Professional physical scientists display tenacious teleological tendencies: Purpose-based reasoning as a cognitive default. J. Exper. Psychol. General 142 (4), 1074–1083. <http://dx.doi.org/10.1037/a0030399>.
- Kim, M., Kim, S., Kim, J., Song, T.-J., Kim, Y., 2024. Do stakeholder needs differ? Designing stakeholder-tailored explainable artificial intelligence (XAI) interfaces. Int. J. Hum.-Comput. Stud. 181, 103160. <http://dx.doi.org/10.1016/j.ijhcs.2023.103160>.
- Kim, S.S.Y., Watkins, E.A., Russakovsky, O., Fong, R., Monroy-Hernández, A., 2023. “Help me help the AI”: Understanding how explainability can support human-AI interaction. In: Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems. CHI '23, Association for Computing Machinery, New York, NY, USA, <http://dx.doi.org/10.1145/3544548.3581001>.
- Kramer, M.F., Schaich Borg, J., Conitzer, V., Sinnott-Armstrong, W., 2018. When do people want AI to make decisions? In: Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society. AIES '18, Association for Computing Machinery, New York, NY, USA, pp. 204–209. <http://dx.doi.org/10.1145/3278721.3278752>.
- Kutsarova, V., 2020. Managing Alarming Situations with Mobile Crowdsensing Systems and Wearable Devices. KTH, School of Electrical Engineering and Computer Science (EECS), p. 65.
- Kuziemski, M., Misuraca, G., 2020. AI governance in the public sector: Three tales from the frontiers of automated decision-making in democratic settings. Telecommun. Policy 44 (6), 101976. <http://dx.doi.org/10.1016/j.telpol.2020.101976>.
- Langer, M., Oster, D., Speith, T., Hermanns, H., Kästner, L., Schmidt, E., Sesing, A., Baum, K., 2021. What do we want from explainable artificial intelligence (XAI)? – a stakeholder perspective on XAI and a conceptual model guiding interdisciplinary XAI research. Artificial Intelligence 296, 103473. <http://dx.doi.org/10.1016/j.artint.2021.103473>.
- Lau, R.R., Redlawsk, D.P., 2001. Advantages and disadvantages of cognitive heuristics in political decision making. Am. J. Political Sci. 45 (4), 951. <http://dx.doi.org/10.2307/2669334>.
- Lee, M.K., Kusbit, D., Kahng, A., Kim, J.T., Yuan, X., Chan, A., See, D., Noothigattu, R., Lee, S., Psomas, A., Procaccia, A.D., 2019. Webuildai: Participatory framework for algorithmic governance. Proc. ACM Human-Comput. Interact. 3 (CSCW), 1–35. <http://dx.doi.org/10.1145/3359283>.
- Liao, Q.V., Gruen, D., Miller, S., 2020. Questioning the AI: Informing design practices for explainable AI user experiences. In: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. CHI '20, Association for Computing Machinery, New York, NY, USA, pp. 1–15. <http://dx.doi.org/10.1145/3313831.3376590>.
- Liao, Q.V., Pribić, M., Han, J., Miller, S., Sow, D., 2021. Question-driven design process for explainable AI user experiences. [arXiv:https://arxiv.org/abs/2104.03483](https://arxiv.org/abs/2104.03483).
- Lim, B.Y., Dey, A.K., 2009. Assessing demand for intelligibility in context-aware applications. In: Proceedings of the 11th International Conference on Ubiquitous Computing. UbiComp '09, Association for Computing Machinery, New York, NY, USA, pp. 195–204. <http://dx.doi.org/10.1145/1620545.1620576>.
- Lima, G., Grgic-Hlaca, N., Jeong, J.K., Cha, M., 2023. Who should pay when machines cause harm? Laypeople's expectations of legal damages for machine-caused harm. In: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency. FAccT '23, Association for Computing Machinery, New York, NY, USA, pp. 236–246. <http://dx.doi.org/10.1145/3593013.3593992>.
- Lombrozo, T., Gwynne, N.Z., 2014. Explanation and inference: Mechanistic and functional explanations guide property generalization. Front. Human Neurosci. 8, <http://dx.doi.org/10.3389/fnhum.2014.00700>.
- Lombrozo, T., Wilkenfeld, D., 2019. Mechanistic versus functional understanding. In: Varieties of Understanding. Oxford University Press, pp. 209–230. <http://dx.doi.org/10.1093/oso/9780190860974.003.0011>.
- Long, D., Magerko, B., 2020. What is AI literacy? Competencies and design considerations. In: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems. ACM, Honolulu HI USA, pp. 1–16. <http://dx.doi.org/10.1145/3313831.3376727>.
- Long, D., Padiyath, A., Teachey, A., Magerko, B., 2021. The role of collaboration, creativity, and embodiment in AI learning experiences. In: Creativity and Cognition. ACM, Virtual Event Italy, pp. 1–10. <http://dx.doi.org/10.1145/3450741.3465264>.
- Lopez, P., 2019. Reinforcing intersectional inequality via the AMS algorithm in Austria. In: Proceedings of the 18th Annual STS Conference. Graz, pp. 289–309. <http://dx.doi.org/10.3217/978-3-85125-668-0-16>.
- Lucaj, L., van der Smagt, P., Benbouzid, D., 2023. AI regulation is (not) all you need. In: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency. FAccT '23, Association for Computing Machinery, New York, NY, USA, pp. 1267–1279. <http://dx.doi.org/10.1145/3593013.3594079>.
- Lupia, A., 2016. Uninformed: Why people know so little about politics and what we can do about it. Oxford University Press, New York, <http://dx.doi.org/10.1093/oso/9780190263720.001.0001>.
- Martijn, M., Conati, C., Verbert, K., 2022. “Knowing me, knowing you”: personalized explanations for a music recommender system. User Model. User-Adapt. Interact. 32 (1–2), 215–252. <http://dx.doi.org/10.1007/s11257-021-09304-9>.
- Miller, T., 2019. Explanation in artificial intelligence: Insights from the social sciences. Artificial Intelligence 267, 1–38. <http://dx.doi.org/10.1016/j.artint.2018.07.007>.
- Miller, T., 2023. Explainable AI is dead, long live explainable AI! hypothesis-driven decision support using evaluative AI. In: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency. FAccT '23, Association for Computing Machinery, New York, NY, USA, pp. 333–342. <http://dx.doi.org/10.1145/3593013.3594001>.
- Mueller, S.T., Hoffman, R.R., Clancey, W., Emrey, A., Klein, G., 2019. Explanation in human-AI systems: A literature meta-review, synopsis of key ideas and publications, and bibliography for explainable AI. [arXiv:https://arxiv.org/abs/1902.01876](https://arxiv.org/abs/1902.01876).
- Nahavandi, D., Alizadehsani, R., Khosravi, A., Acharya, U.R., 2022. Application of artificial intelligence in wearable devices: Opportunities and challenges. Comput. Methods Programs Biomed. 213, 106541. <http://dx.doi.org/10.1016/j.cmpb.2021.106541>.
- Ng, D.T.K., Leung, J.K.L., Chu, K.W.S., Qiao, M.S., 2021. AI literacy: Definition, teaching, evaluation and ethical issues. Proc. Assoc. Inf. Sci. Technol. 58 (1), 504–509. <http://dx.doi.org/10.1002/pra2.487>.
- Niklas, J., Sztandar-Sztanderska, K., Szymielewicz, K., 2015. Profiling the unemployed in Poland: Social and political implications of algorithmic decision making. panoptykon.org/sites/default/files/leadimage-biblioteka/panoptykon_profiling_report_final.pdf.
- Páez, A., 2019. The pragmatic turn in explainable artificial intelligence (XAI). Minds Mach. 29 (3), 441–459. <http://dx.doi.org/10.1007/s11023-019-09502-w>.
- Phillips, P.J., Hahn, C.A., Fontana, P.C., Yates, A.N., Greene, K., Broniatowski, D.A., Przybocci, M.A., 2021. Four principles of explainable artificial intelligence. Technical Report NIST IR 8312, National Institute of Standards and Technology (U.S.), Gaithersburg, MD, <http://dx.doi.org/10.6028/NIST.IR.8312>.
- Raji, I.D., Kumar, I.E., Horowitz, A., Selbst, A., 2022. The fallacy of AI functionality. In: 2022 ACM Conference on Fairness, Accountability, and Transparency. ACM, Seoul Republic of Korea, pp. 959–972. <http://dx.doi.org/10.1145/3531146.3533158>.
- Ram, A., 1989. Question-driven understanding: An integrated theory of story understanding, memory and learning (Ph.D. thesis). Yale University, USA, <http://dx.doi.org/10.5555/916448>.
- Rozenblit, L., Keil, F., 2002. The misunderstood limits of folk science: An illusion of explanatory depth. Cogn. Sci. 26 (5), 521–562. <http://dx.doi.org/10.1207/s15516709cog2605.1>.
- Rudin, C., 2019. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. Nat. Mach. Intell. 1 (5), 206–215. <http://dx.doi.org/10.1038/s42256-019-0048-x>.
- Rugg, G., Petre, M., 2007. A gentle guide to research methods. Open University Press.
- Sahlen, M., Alam, E., Silva, P., 2018. Ny teknik i välfärden - vår tids välfärdskliv (New Technologies in Welfare – The Welfare Leap of Our Times). Technical Report, Svenskt Näringsliv, svensktnaringsliv.se/sakomraden.
- Sato, B.K., Hill, C.F.C., Lo, S.M., 2019. Testing the test: Are exams measuring understanding? Biochem. Mol. Biol. Educ. 47 (3), 296–302. <http://dx.doi.org/10.1002/bmb.21231>.

- Schmude, T., Koesten, L., Möller, T., Tschischek, S., 2023. On the impact of explanations on understanding of algorithmic decision-making. In: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency. FAccT '23, Association for Computing Machinery, New York, NY, USA, pp. 959–970. <http://dx.doi.org/10.1145/3593013.3594054>.
- Schoeffer, J., Kuehl, N., Machowski, Y., 2022. “There is not enough information”: On the effects of explanations on perceptions of informational fairness and trustworthiness in automated decision-making. In: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency. ACM, Seoul Republic of Korea, pp. 1616–1628. <http://dx.doi.org/10.1145/3531146.3533218>.
- Scott, K.M., Wang, S.M., Miceli, M., Delobelle, P., Sztandar-Sztanderska, K., Berendt, B., 2022. Algorithmic tools in public employment services: Towards a jobseeker-centric perspective. In: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency. FAccT '22, ACM, Seoul Republic of Korea, pp. 2138–2148. <http://dx.doi.org/10.1145/3531146.3534631>.
- Seneviratne, S., Hu, Y., Nguyen, T., Lan, G., Khalifa, S., Thilakarathna, K., Hassan, M., Seneviratne, A., 2017. A survey of wearable devices and challenges. IEEE Commun. Surv. Tutor. PP, 1. <http://dx.doi.org/10.1109/COMST.2017.2731979>.
- Shen, H., Jin, H., Cabrera, A.A., Perer, A., Zhu, H., Hong, J.I., 2020. Designing alternative representations of confusion matrices to support non-expert public understanding of algorithm performance. Proc. ACM Human-Comput. Interact. 4 (CSCW2), 1–22. <http://dx.doi.org/10.1145/3415224>.
- Shin, D., 2021. The effects of explainability and causability on perception, trust, and acceptance: Implications for explainable AI. Int. J. Hum.-Comput. Stud. 146, 102551. <http://dx.doi.org/10.1016/j.ijhcs.2020.102551>.
- Shin, D., 2022. How do people judge the credibility of algorithmic sources? AI Soc. 37 (1), 81–96. <http://dx.doi.org/10.1007/s00146-021-01158-4>.
- Shin, D.D., 2023. Algorithms, humans, and interactions: How do algorithms interact with people? designing meaningful AI experiences, first ed. Routledge, Boca Raton, FL. <http://dx.doi.org/10.1201/b23083>.
- Shin, D., Kee, K.F., Shin, E.Y., 2022a. Algorithm awareness: Why user awareness is critical for personal privacy in the adoption of algorithmic platforms? Int. J. Inf. Manage. 65, 102494. <http://dx.doi.org/10.1016/j.ijinfomgt.2022.102494>.
- Shin, D., Lim, J.S., Ahmad, N., Ibahrine, M., 2022b. Understanding user sensemaking in fairness and transparency in algorithms: Algorithmic sensemaking in over-the-top platform. AI & Soc. 39 (2), 477–490. <http://dx.doi.org/10.1007/s00146-022-01525-9>.
- Shin, D., Shin, E.Y., 2023. Human-centered AI: A framework for green and sustainable AI. Computer 56 (6), 16–25. <http://dx.doi.org/10.1109/MC.2023.3241071>.
- Shneiderman, B., 2022. Human-centered AI. Oxford University Press, Oxford.
- Shulner-Tal, A., Kuflik, T., Klinger, D., 2022. Enhancing fairness perception – towards human-centred AI and personalized explanations understanding the factors influencing laypeople’s fairness perceptions of algorithmic decisions. Int. J. Hum.–Comput. Int. 1–28. <http://dx.doi.org/10.1080/10447318.2022.2095705>.
- Swiss Federal Council, 2023. Popular vote. admin.ch/gov/en/start/documentation/votes.html.
- Szymanski, M., Millicamp, M., Verbert, K., 2021. Visual, textual or hybrid: the effect of user expertise on different explanations. In: 26th International Conference on Intelligent User Interfaces. ACM, College Station TX USA, pp. 109–119. <http://dx.doi.org/10.1145/3397481.3450662>.
- Urqhart, C., Lehmann, H., Myers, M.D., 2009. Putting the ‘theory’ back into grounded theory: guidelines for grounded theory studies in information systems: Guidelines for grounded theory studies in information systems. Inf. Syst. J. 20 (4), 357–381. <http://dx.doi.org/10.1111/j.1365-2575.2009.00328.x>.
- Wang, X., Yin, M., 2021. Are explanations helpful? A comparative study of the effects of explanations in AI-assisted decision-making. In: 26th International Conference on Intelligent User Interfaces. IUI '21, Association for Computing Machinery, New York, NY, USA, pp. 318–328. <http://dx.doi.org/10.1145/3397481.3450650>.
- Weiss, R.S., 1995. Learning from strangers: The art and method of qualitative interview studies, 1. Free Press paperback Free Press, New York, NY.
- Wieringa, M., 2023. “Hey syri, tell me about algorithmic accountability”: Lessons from a landmark case. Data & Policy 5, e2. <http://dx.doi.org/10.1017/dap.2022.39>.
- Wiggins, G.P., McTighe, J., 2005. Understanding by design, Expanded 2nd Association for Supervision and Curriculum Development, Alexandria, VA.
- Woodruff, A., Fox, S.E., Rouso-Schindler, S., Warshaw, J., 2018. A qualitative exploration of perceptions of algorithmic fairness. In: Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems. ACM, Montreal QC Canada, pp. 1–14. <http://dx.doi.org/10.1145/3173574.3174230>.
- Xu, W., 2019. Toward human-centered AI: A perspective from human-computer interaction. Interactions 26 (4), 42–46. <http://dx.doi.org/10.1145/3328485>.
- Yurrita, M., Murray-Rust, D., Balayn, A., Bozzon, A., 2022. Towards a multi-stakeholder value-based assessment framework for algorithmic systems. In: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency. FAccT '22, Association for Computing Machinery, New York, NY, USA, pp. 535–563. <http://dx.doi.org/10.1145/3531146.3533118>.
- Zagzebski, L., 2019. Toward a theory of understanding. In: Varieties of Understanding. Oxford University Press, pp. 123–136. <http://dx.doi.org/10.1093/oso/9780190860974.003.0007>.
- Zerilli, J., 2022. Explaining machine learning decisions. Philos. Sci. 89 (1), 1–19. <http://dx.doi.org/10.1017/psa.2021.13>.
- Zhao, H., Chen, H., Yang, F., Liu, N., Deng, H., Cai, H., Wang, S., Yin, D., Du, M., 2024. Explainability for large language models: A survey. ACM Trans. Intell. Syst. Technol. 3639372. <http://dx.doi.org/10.1145/3639372>.
- Züger, T., Asghari, H., 2023. AI for the public. how public interest theory shifts the discourse on AI. AI Soc. 38 (2), 815–828. <http://dx.doi.org/10.1007/s00146-022-01480-5>.



Timothée Schmude is a Ph.D. student at the University of Vienna, Austria. He holds a master's degrees in Information Processing and Professional Writing from University of Cologne and a bachelor's degree in Media Studies from University of Bayreuth, Germany. His research is situated in the interface between Machine Learning and Social Sciences and employs empirical methods to examine how explanations of AI systems can be made suitable for a larger population. He further collaborates with legal scholars to investigate the tensions between regulatory and practical requirements of explainable AI.



Laura Koesten is a postdoctoral researcher at the Faculty of Computer Science at the University of Vienna. She researches various facets of human data interaction, including sensemaking with data and visualizations, data discovery and reuse, along with ethical and collaborative dimensions of data-centric work. She received her Ph.D. degree in computer science from the University of Southampton and the Open Data Institute, UK; and her MSc. in Human Factors from Loughborough University, UK.



Torsten Möller is a professor of computer science at the University of Vienna, Austria, since 2013. Between 1999 and 2012 he served as a Computing Science faculty member at Simon Fraser University, Canada. He received his Ph.D. in Computer and Information Science from Ohio State University in 1999 and a Vordiplom (B.Sc.) in mathematical computer science from Humboldt University of Berlin, Germany. He is a senior member of IEEE and ACM, and a member of Eurographics. His research interests include algorithms and tools for analyzing and displaying data with principles rooted in computer graphics, human-computer interaction, data science, and visualization.



Sebastian Tschischek is assistant professor for Machine Learning at the Faculty of Computer Science at the University of Vienna. He received his Ph.D. from Graz University of Technology in 2014. After his Ph.D., he was a postdoctoral fellow at ETH Zurich and a senior researcher at Microsoft Research Cambridge. He develops machine learning algorithms for structured objects and sequential decision-making and aims to understand the impact of the usage of AI on society. His research focuses on expressive models for heterogeneous data and on the impact of uncertainty and its quantification, for instance on sequential decision making and human-machine interaction.