

A Taxonomy of Questions for Critical Reflection in Machine-Assisted Decision-Making

Simon W.S. Fischer¹, Hanna Schraffenberger², Serge Thill¹, Pim Haselager¹

¹Donders Institute for Brain, Cognition, and Behaviour, Dpt. Human-Centred Intelligent Systems, Nijmegen, The Netherlands

²iHub, Radboud University, Nijmegen, The Netherlands
simon.fischer@donders.ru.nl

Abstract

Decision-makers run the risk of relying too much on machine recommendations, which is associated with lower cognitive engagement. Reflection has been shown to increase cognitive engagement and improve critical thinking and therefore decision-making. Questions are a means to stimulate reflection, but there is a research gap regarding the systematic creation and use of relevant questions for machine-assisted decision-making. We therefore present a taxonomy of questions aimed at promoting reflection and cognitive engagement in order to stimulate a deliberate decision-making process. Our taxonomy builds on the Socratic questioning method and a question bank for explainable AI. As a starting point, we focus on clinical decision-making. Brief discussions with two medical and three educational researchers provide feedback on the relevance and expected benefits of our taxonomy. Our work contributes to research on mitigating overreliance in human-AI interactions and aims to support effective human oversight as required by the European AI Act.

1 Introduction

Reflection, also referred to as reflective practice or reflective thought, is an important skill in professional practices (Schön 1983). It is beneficial in, for example, management (Cunliffe 2002), law (Casey 2014), and healthcare (Mann, Gordon, and MacLeod 2009), and is part of many educational curricula (Boud and Walker 1998). As the American philosopher John Dewey (1933, p.9) states, reflection is an

“Active, persistent, and careful consideration of any belief or supposed form of knowledge in the light of the grounds that support it and further conclusion to which it tends”.

As reflection helps to scrutinise the validity of assumptions and information (Mezirow 1990), it has the potential to improve reasoning, judgement, and problem-solving (Ghanizadeh 2017; Khoshgoftar and Barkhordari-Sharifabad 2023; Mamede and Schmidt 2023). Consequently, reflection has been shown to improve strategic decisions (Walger, Roglio, and Abib 2016) and diagnostic accuracy (Mamede, Schmidt, and Penaforte 2008; Hess et al. 2015; Prakash,

Sladek, and Schuwirth 2019). In addition, reflection enables decision-makers to recognise and articulate the reasons for their decisions, allowing them to direct their actions accordingly and take responsibility (forward-looking) and accountability (backward-looking). Responsible decision-making ideally leads to more desirable outcomes for all parties involved. In view of this, it seems sensible to encourage and support reflection during decision-making, especially as many high-impact decisions nowadays are informed by machine recommendations or predictions.

Decision-support systems (DSS) provide recommendations that assist decision-makers in domains like law, finance and healthcare. Physicians, for example, can use clinical DSSs to make a diagnosis or find a treatment option. Studies show, however, that operators tend to rely too much on these systems and accept incorrect recommendations (Dratsch et al. 2023), which is known as overreliance (Passi and Vorvoreanu 2022). In order to reduce harmful consequences of overreliance, the European AI Act (European Parliament 2024, §14) requires:

[...] that natural persons to whom human oversight is assigned are enabled, as appropriate and proportionate: [...] (b) to remain aware of the possible tendency of automatically relying or over-relying on the output produced by a high-risk AI system (automation bias), in particular for high-risk AI systems used to provide information or recommendations for decisions to be taken by natural persons;

Various approaches therefore aim to mitigate overreliance by increasing the decision-maker’s cognitive engagement through different interactions, such as presenting evidence for and against a decision (Miller 2023) or presenting explanations in the form of questions (Danry et al. 2023). Others argue that DSS should support deliberation and reflection rather than make recommendations (Schmidt and Mamede 2023; Ma et al. 2024). Similarly, in previous work we have introduced the idea of a *reflection machine*, which asks questions to the decision-maker (Cornelissen et al. 2022; Haselager et al. 2023; Fischer 2025). However, it remains unclear what kind of questions should be asked in order to effectively stimulate cognitive engagement and reflection.

To address this gap, we provide a taxonomy of questions that helps to systematically formulate relevant questions that stimulate critical reflection on the decision at hand,

including the machine recommendation and the data that led to it. To identify relevant questions, the Socratic questioning method, named after the Greek philosopher Socrates, proves to be a useful starting point as it involves systematically asking questions to explore complex ideas, uncover assumptions, and gain a better understanding of the topic under investigation (Paul 2019). As such, Socratic questioning has been shown to increase students' critical thinking (Dalim 2022; Ho, Chen, and Li 2023). Although the Socratic method is slowly gaining traction in the field of human-machine interaction (Lara 2021; Ang, Gollapalli, and Ng 2023; Liu et al. 2024), there is as yet no systematic approach for applying this questioning method to the context of machine-assisted decision-making.

We therefore adapt a taxonomy of Socratic questions based on concepts of critical thinking (Paul 2019) to identify aspects to which questions should relate, and transfer it to the context of machine-assisted decision-making by synthesising it with a question bank for the design of human-centred explainable AI (Liao, Gruen, and Miller 2020; Liao et al. 2021; Liao and Varshney 2022). We also draw inspiration, albeit only marginally, from a revised version of Bloom's taxonomy (Krathwohl 2002; Forehand 2010) to identify cognitive processes to which questions can refer.

While we aim for a domain-independent taxonomy by building on general concepts of critical thinking, we take as a starting point clinical decision-making, where a physician uses a DSS to find a treatment option among several possibilities with uncertain outcomes for patients. To assess the transferability of our taxonomy, we apply it to the educational domain, where teachers use learning analytics dashboards to make decisions about students. We have discussed our taxonomy with medical and educational researchers, who have emphasised the relevance and expected benefits of our taxonomy. Nevertheless, an empirical evaluation of the effectiveness of certain questions in a specific decision-making context remains for future work. The contributions of this paper are as follows:

- With regard to mitigating overreliance (section 2.1) and human-centred XAI (section 2.2), we add to the growing literature on approaches to increase cognitive engagement, and, in particular, expand on existing proposals to promote reflection (section 2.3).
- We discuss literature that shows that reflection increases cognitive engagement and reasoning, and thus decision-making (section 2.4), and that questions can stimulate reflection (section 2.5).
- We present a new taxonomy to help formulate relevant questions that stimulate reflection in machine-assisted decision-making (section 4). This taxonomy supports the design and development of interventions that go beyond machine recommendations and explanations. In addition, our taxonomy can benefit stakeholders such as developers, auditors, or consumers who need to remain critical in the development, evaluation, and use of DSS (section 6).

2 Background

In the following, we will discuss background work that led us to the creation of our question taxonomy. In order to sketch the problem of overreliance on machine recommendations and explanations (section 2.1), we begin with a short review of the relevant literature on human-machine interaction, particularly in relation to the human factor in explainable AI (section 2.2). From alternative interaction approaches, we derive that promoting *cognitive engagement* is a promising method to reduce overreliance (section 2.3). We identify the benefits of *reflective thought* for cognitive engagement as it promotes critical thinking, reasoning, and thus responsible decision-making (section 2.4). Finally, we suggest that *questions* can be useful to stimulate reflection and thus cognitive engagement, and introduce the Socratic questioning method, which is widely used in education (section 2.5). In short, to arrive at our taxonomy, we combine relevant insights from empirical findings and literature on human-machine interaction, education, and philosophy.

2.1 Overreliance on DSS and Explainable AI

Decision-support systems (DSS) provide recommendations based on information by using (simple) mathematical models, or more advanced machine learning models, including deep learning approaches. A meta-study of 106 experimental studies on machine-assisted decision-making found that decision-making accuracy generally decreases when decision-makers use DSS (Vaccaro, Almaatouq, and Malone 2024). Decision-makers may tend to prefer machine recommendations over their own judgement (Tschandl et al. 2020), even when the recommendations are wrong (Jacobs et al. 2021). The acceptance of incorrect recommendations, referred to as overreliance, can be because of low vigilance or lack of cognitive engagement (Passi and Vorvoreanu 2022).

A common approach to mitigating overreliance is explainable AI (XAI) (Vasconcelos et al. 2023; Longo et al. 2024), which can, for example, provide information about which input was decisive for a particular output. It is assumed that these explanations will help the operator to assess the DSS recommendation and thus make an informed decision as to whether or not to follow the recommendation.

The aforementioned meta-study, however, found that explanations do not lead to significant synergy effects in human-machine decision-making (Vaccaro, Almaatouq, and Malone 2024). On the contrary, explanations may increase the likelihood of operators accepting a machine recommendation regardless of its correctness (Bansal et al. 2021; Bussonne, Stumpf, and O'Sullivan 2015; Ehsan et al. 2021), thereby increasing overreliance (Vered et al. 2023). One reason for this is that the focus on numerical explanations can give a false sense of objectivity, leading operators, regardless of their expertise, to overestimate the capabilities of a DSS (Ehsan et al. 2024). Consequently, providing additional information in the form of explanations does not automatically lead to more cognitive engagement of the decision-maker and to more accurate decisions (Gajos and Mamykina 2022).

2.2 A Question Bank for XAI

To provide more relevant and actionable explanations, a question bank for explainable AI helps designers and developers select the appropriate XAI technique that supports the operator's information needs (Liao, Gruen, and Miller 2020; Liao et al. 2021; Liao and Varshney 2022). To this end, the question bank clusters and maps prototypical user questions to XAI techniques. Feature importance, for example, can provide insights into why-questions, such as "Why was this prediction given?"

Answering the (technical) question of how a DSS arrived at a particular output is however not always desired or helpful. As one participant working on clinical decision-support systems mentioned in the study on the XAI question bank (Liao, Gruen, and Miller 2020, p.6):

[explanations by system rationale] are essentially 'this is how I do it, take it or leave it'. But doctors don't like this approach...Thinking that [AI is] giving treatment recommendations is the wrong place to start, because doctors know how to do it. It's everything that happens around that decision they need help with... more discussions about the output, rather than how you get there.

Although the XAI question bank allows for the design of more human-centred XAI, it is still very algorithmic- and data-centred, as the authors themselves state and as we will discuss below. Even if it is known how a DSS arrived at an output or recommendation, operators might be more interested in factors regarding the broader socio-technical context outside of the model (Zednik 2019; Ehsan et al. 2021; Ghassemi, Oakden-Rayner, and Beam 2021; Ehsan and Riedl 2024), which could relate to questions like, "What is missing from the dataset?", "Are there any other overlooked alternative solutions?", or "How to integrate this recommendation into the decision-process?"

With our proposed taxonomy, we want to stimulate "more discussions about the output" (study participant in Liao, Gruen, and Miller 2020) by identifying and categorising relevant questions. The XAI question bank will prove useful as it contains questions of operators and maps them to XAI techniques. As we will show, we utilise the mapping from questions to XAI techniques and reverse the process to re-purpose explanations to formulate possible questions, e.g., "Is it relevant to focus on factor x ?". As such, in line with actionability, we consider "explanations [informing questions] as a means to help a data subject [and decision-maker] *act* rather than merely understand" (Wachter, Mittelstadt, and Russell 2017, p.843), where 'act' in this case is deliberate decision-making through reflection.

2.3 Increasing Cognitive Engagement

To mitigate overreliance, scholars explore different interactions with machine recommendations and explanations in addition to XAI that attempt to cognitively engage the decision-maker in the decision-making process. We will briefly mention four relevant approaches: 1) promoting cognitive engagement through interventions, 2) supporting the decision-maker to make their own decision, 3) formulating

explanations as questions, and 4) proposals to stimulate reflection.

First, cognitive interventions are based on the dual-process model of reasoning, and aim to interrupt the decision-maker's habitual thinking (system 1) in order to encourage analytical thinking (system 2) (Lambe et al. 2016). Cognitive interventions can take the form of checklists, instructions for analytical thinking, or reflection, for example (Lambe et al. 2016; Prakash, Sladek, and Schuwirth 2019). In the context of human-machine interaction, one study found that such interventions, i.e., presenting explanations delayed, on-demand, or only after the operator made a decision, can reduce, yet not eliminate, overreliance on DSS (Bućinca, Malaya, and Gajos 2021).

Second, instead of presenting a recommendation and justifying it with an explanation that the decision-maker must accept or reject, a so-called hypothesis-driven recommender shows evidence for and against a decision (Miller 2023). In this way, the system supports the decision-maker's reasoning process and allows them to make an informed decision themselves. A study found that while presenting evidence for and against did little to improve diagnostic accuracy, physicians valued the reflective aspect and had more confidence in their final decisions (Cabitza et al. 2023).

Third, one study found that explanations formulated as questions improve human judgement of the logical validity of the information provided (Danry et al. 2023). Causal explanations were presented both as statements and questions. For example, the statement: "The patient has the flu because of a headache", is formulated as a question: "If the patient has a headache, does it follow that they have the flu?". Similar to the hypothesis-driven recommender, questions help the decision-maker to think for themselves (Ryan and Deci 2000), which leads to more cognitive engagement (Gajos and Mamykina 2022).

Fourth, it has been suggested that DSS should take on roles other than that of a recommender, such as a devil's advocate raising objections and challenging the decision-maker (Chiang et al. 2024; Ma et al. 2024; Sarkar 2024). Insights from decision-making in human groups show that a devil's advocate has positive effects on decision performance (Schwenk 1984; Schweiger, Sandberg, and Rechner 1989). Similarly, others argue that DSS should support deliberation and facilitate reflection (Cornelissen et al. 2022; Haselager et al. 2023; Schmidt and Mamede 2023; Zhang, Yang, and Li 2023; Ma et al. 2025).

In view of the above, we understand questions as cognitive interventions (1), which facilitate own reasoning and decision-making (2&3). With our taxonomy, we therefore want to both build on and expand these four approaches by offering more potential questions compared to (3) and supporting the design of suggested interactions that promote reflection (4).

2.4 Reflection and Responsible Decision-Making

Amongst various forms of increasing cognitive engagement, such as checklists or feedback, stimulating reflection proved to be the most effective (Lambe et al. 2016; Prakash, Sladek, and Schuwirth 2019). While reflection can occur be-

fore, during or after an action (Rogers 2001), we focus on reflection during decision-making, i.e., reflection-in-action (Schön 1983). As mentioned earlier, reflection can promote reasoning, critical thinking and problem-solving, and thus improve decision-making.

In areas where decisions have far-reaching consequences, such as law or healthcare, decision-makers have great professional responsibility. Physicians, for example, are committed to the well-being of patients. This professional responsibility includes an epistemological responsibility that requires decision-makers to gather and evaluate relevant information, and to know the reasons for a particular course of action (van Baalen and Boon 2015). Being able to provide reasons for (past) actions allows for backward-looking responsibility or accountability. More importantly, knowing the reasons enables one to determine and direct future actions to achieve a desired state, and thus to take forward-looking responsibility (Loughran 1996).

Reflection-in-action can therefore increase epistemological and thus professional responsibility, both in a backward-looking and, more importantly, in a forward-looking sense. Ideally, this forward-looking responsibility leads to fairer and more desirable decision outcomes for all parties involved (Ueda et al. 2024). Reflection-in-action could, for example, prevent a physician from jumping to conclusions after seeing seemingly similar patient cases throughout the day, which could lead to tunnel vision (Mamede and Schmidt 2023). Reflection makes it possible to distinguish between relevant and irrelevant information, and, to quote Dewey (1933, p.17) again, “converts action that is merely appetitive, blind and impulsive into intelligent action”.

2.5 Types of Questions

A common technique to facilitate reflection is to ask questions (Osmond and Darlington 2005). Questions encourage independent thinking, which can increase confidence in one’s own decision-making, i.e., decision self-efficacy. This in turn can affect the motivation and well-being of the decision-maker (Ryan and Deci 2000). Findings from educational research suggest that motivation functions as an antecedent to cognitive engagement and also increases it (Blumenfeld, Kempner, and Krajcik 2006; Rotgans and Schmidt 2011; Singh et al. 2022). Questions thus seem to promote reflection, which can increase decision self-efficacy and motivation, which in turn can increase cognitive engagement.

Given the breadth of possible questions, we want to point out two works from the field of education that categorise questions in different ways and which we also use for our purposes, namely a taxonomy of Socratic questions (Paul 2019), and Bloom’s taxonomy for educational objectives (Anderson and Krathwohl 2001; Krathwohl 2002).

First, Socratic questioning is widely used in education to develop and promote critical thinking skills (Torabizadeh, Homayuni, and Moattari 2018; Makhene 2019; Dalim 2022; Ho, Chen, and Li 2023). It consists of systematically asking questions that help to clarify concepts, improve understanding, and uncover gaps in knowledge (Paul 2019). Accordingly, Socratic questioning helps to discover one’s thoughts, to analyse and evaluate assumptions, information and infer-

ences, and to arrive at judgements through own reasoning. To support teachers in the classroom, a taxonomy of Socratic questions has been developed, defining eight types of questions that probe different parts of thinking, such as assumptions, reasons, or implications (Paul 2019).

Second, Bloom’s taxonomy categorises educational learning objectives hierarchically into six levels of cognitive processes, which are: *remembering, understanding, applying, analysing, evaluating, and creating*. The last three levels, analysing, evaluating, and creating, are also components of critical thinking (Paul 2019). This categorisation will be helpful when formulating questions that address different cognitive processes.

In the context of human-machine interaction, the Socratic questioning method is slowly gaining traction. Based on the taxonomy of Socratic questions, a dataset of 110k questions that could be used in counselling and coaching is proposed, including a generation method of prompt-tuning a language model (Ang, Gollapalli, and Ng 2023). The resulting questions fall into one of the six categories defined by Paul (2019): 1) clarification (“What do you mean by...?”), 2) probing assumptions (“Why do you assume...?”), 3) probing reasons (“How do you know that...?”), 4) probing implications (“If..., what is likely to happen?”), 5) probing alternatives (“What else should we consider?”), and 6) others that do not conform to the Socratic categories (Ang, Gollapalli, and Ng 2023). In addition, a virtual assistant that fosters moral enhancement through Socratic questions is proposed (Lara and Deckers 2020), as well as chatbots based on language models that utilise the Socratic method for thought-provoking tutoring (Favero et al. 2024; Liu et al. 2024).

However, there is currently no systematic approach to the use of (Socratic) questions to promote critical reflection in decision-making, particularly in conjunction with decision-support systems. Given the link between critical reflection and reasoning and the associated positive effect on decision-making, we see potential for the use of questions to promote decision self-efficacy and cognitive engagement. We therefore aim to address this gap by adapting the taxonomy of Socratic questions (Paul 2010, 2019) to the context of machine-assisted decision-making. In doing so, we also aim for a higher level of granularity as compared to Ang, Gollapalli, and Ng (2023), who utilise parts of the Socratic questions taxonomy in unchanged form to generate counselling and coaching questions.

3 Methodology

Our taxonomy is based on the taxonomy of Socratic questions (Paul 2019), the XAI question bank (Liao and Varshney 2022), and a revision of Bloom’s taxonomy (Anderson and Krathwohl 2001) discussed in the previous section. To construct our taxonomy we proceeded in two steps: 1) identifying elements that questions can relate to, and 2) identifying information for formulating and enriching questions. In order to identify elements (1), we synthesised the taxonomy of Socratic questions with the XAI question bank through a procedure similar to affinity diagramming. In order to help formulate and enrich questions (2), we reversed the mapping of the XAI question bank to repurpose information de-

rived from explanations. In addition, we used the cognitive processes identified in Bloom’s taxonomy to determine the scope of questions.

3.1 Procedure

Phase 1: Identifying Elements for Reflection First, we identified relevant aspects of machine-assisted decision-making to which potential questions should relate, i.e., *what* to reflect on. In a deductive (top-down) approach, the categories from the taxonomy of Socratic questions served as a starting point. As Elder and Paul (2010) note:

“Whenever we think we think for a *purpose* within a *point of view* based on *assumptions* leading to *implications and consequences*. We use *data, facts and experiences* to make *inferences* and judgements based on *concepts* and theories to answer a *question* or solve a problem.”

Based on this structure of thinking, the authors categorise questions into elements of thought or reasoning as follows:

- **Purpose:** Questions about the agenda, goal or objective.
- **Question at Issue:** Questions about the problem or issue that gave rise to the question.
- **Information:** Questions about background information, such as data, facts, observations or experiences.
- **Interpretation and Inferences:** Questions about how meaning was derived and conclusions drawn.
- **Concepts:** Questions about underlying theories, definitions, or models that define a thought.
- **Assumptions:** Questions about taken for granted presuppositions of a thought.
- **Implications and Consequences:** Questions about implications.
- **Point of View:** Questions about the frame of reference or perspective of a thought.

While these elements of thought, including the provided example questions, are helpful in general, they cannot be directly applied to the domain of decision-making (with DSSs). This is because the taxonomy of Socratic Questions is primarily intended for teachers and students in the classroom.

We thus took the elements of thought as a starting point for our taxonomy and mapped them to the context of machine-assisted decision-making. For this mapping and to identify relevant aspects of machine-assisted decision-making, we used the XAI question bank. The question bank contains more than 50 common questions that operators have towards a DSS, categorised as follows (Liao and Varshney 2022):

- **How** (global model-wide): Questions about the general logic of the model.
- **Why** (a given prediction): Questions about the reasons for a prediction.
- **Why Not** (a different prediction): Questions about the difference to an expected outcome.

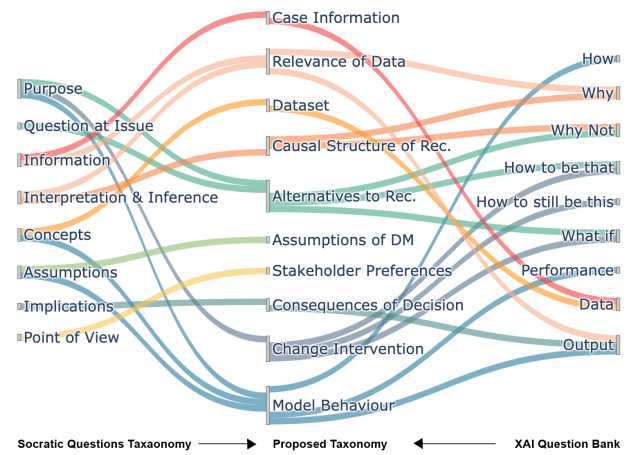


Figure 1: Our elements for reflection in machine-assisted decision-making in the middle column (see section 4) are based on the taxonomy for Socratic questions (Paul 2019), in the left column, and the XAI question bank (Liao and Varshney 2022), in the right column.

- **How to be That** (a different prediction): Questions about how to change an instance to get a different outcome.
- **How to still be This** (the current prediction): Questions about possible changes to still get the same outcome.
- **What if:** Questions about changed outcome based on changed input.
- **Performance:** Questions about the performance of the DSS.
- **Data:** Questions about training data.
- **Output:** Questions about how to apply or use the DSS output.

Overall, questions from the XAI question bank relate to the logic of the model, its performance, the reasons for a prediction, possible changes, the data, and the output.

This leaves us with the categorisation of questions based on thought structure on the one hand, and the components of machine-assisted decision-making represented by clusters of common questions in the context of explainable AI on the other (Figure 1). The element *Information*, for example, can be mapped, as is also mentioned in the definition, to *Data* the DSS is trained on, as well as (input) data of a specific case. Through affinity diagramming we arrived at the following 10 elements, which are shown in the middle column of Figure 1: *case information, relevance of data, dataset, causal structure of recommendation, alternatives to recommendation, assumptions of decision-maker, stakeholder preferences, consequences of decision, change intervention, and model behaviour*.

Phase 2: Formulating Questions After identifying *what* to question, the next step was to specify *how* to formulate and enrich specific questions. For this, we again used the XAI question bank, in particular the mapping of common questions to XAI techniques. Instead of starting with questions from operators in order to arrive at suitable explana-

tions, however, we reversed the process. We moved from explanations to questions, in order to extract information using XAI techniques, which then flows into possible questions. On the one hand, this means that (some) explanations can be (re-)formulated as questions, e.g., “Does outcome Y follow from information x ?” (Danry et al. 2023). On the other hand, questions can relate to information extracted from explanations, e.g., “Is information x a relevant factor?”.

Finally, to capture the scope of questions, we used, albeit only marginally, a revised version of Bloom’s taxonomy. In Bloom’s taxonomy, questions are categorised into levels of cognitive processes to support the formulation of learning objectives. Three of these categories or cognitive abilities are components of critical thinking (Paul 2019), namely, analysing, evaluating, and creating, where (Anderson and Krathwohl 2001):

- **Analysing** means to break down information into parts and determining how these parts are connected, and identifying which parts are relevant or irrelevant. Verbs that relate to this cognitive process are: *compare, criticize, differentiate, discriminate, deconstruct, inspect*. For example, “What evidence is there for Y ?”.
- **Evaluating** means to judge information and its appropriateness, and to detect inconsistencies. Verbs that relate to this cognitive process are: *assess, support, defend, validate, evaluate*. For example, “How to justify Y ?”
- **Creating** means compiling information in a new way, or considering alternative hypotheses. Verbs that relate to this cognitive process are: *create, develop, formulate*. For example, “What could be done to minimise Y ?”

Based on this cognitive distinction, we suggest that questions should address these higher-level cognitive processes in order to effectively stimulate critical reflection.

3.2 Expert Feedback

While one author created the initial taxonomy, we discussed the categories with their overlaps and boundaries among the authors throughout the process. Besides, we created the taxonomy in an iterative way. Once we had a first version, we sent 10 initial categories including some sample questions to one orthopaedic surgeon and one senior researcher in orthopaedics, who both work at different hospitals and are part of an ongoing collaboration. To get some early feedback, we asked them if they saw any missing question types. They confirmed the completeness of the derived question types and mentioned that it seems we “covered all the aspects” and that they “don’t see any obvious questions missing”. We then refined the categories, without adding or removing any, expanded details, added descriptions, and added sample questions to a separate table. We presented the final taxonomy again to the same surgeon and researcher (section 5).

4 Taxonomy of Questions

Our taxonomy consists of 10 question types or elements for reflection (Table 1). In order to demonstrate how to apply our taxonomy, we provide some sample questions from the medical and educational domain. A description of the individual question type follows, focussing on the medical field,

where a physician uses a DSS to make a decision about a patient’s treatment.

Q1. Questions can relate to the *case information* provided, which represents information in the form of data. Previous operations on the spine, for example, might be indicated by integers. This value, however, does not indicate how long ago the operation was performed and at which location of the spine it was carried out, both of which can influence the effectiveness of further operations. Questions can help the decision-maker to inspect and contextualise the data provided to give it semantic depth. To ensure data quality, questions can also relate to the reliability of the data provided to rule out inconsistencies, such as measurement errors, and to ensure that no relevant data is missing. As such, Q1 helps the decision-maker get a more holistic picture of the case (van Baalen and Boon 2015).

Q2. Questions can relate to the *relevance of data*. In doing so, questions can encourage the decision-maker to contextualise the information, e.g., symptom, in a specific case, which can lead to validating the machine recommendation by other means, e.g., by conducting additional tests. So while XAI techniques, such as feature importance, can show which features have contributed to an outcome, domain knowledge is necessary to evaluate whether these data points are important in a particular case. It is possible to use the information extracted with XAI to generate an evaluative question, like “How relevant is data x ?”. The question of the relevance of the most important data point(s) also raises the question of whether there are other factors that could be overlooked in the decision. As different XAI techniques can result in different rankings of the contributing effects of the individual features (Saarela and Jauhiainen 2021), it is important for decision-makers not to take the information obtained through XAI at face value. The discrepancies between two opposing feature ranks from two different XAI methods could also lead to a question like “Is x or z more relevant?”.

Q3. Questions can relate to the *dataset* on which the model was built. Cardiovascular disease, for example, has different symptoms in men and women, yet most of the data for training DSSs stem from male patients, which is why DSSs often misdiagnose women (Li et al. 2023). Certain populations or phenomena are often overrepresented in data, which leads to a skewed data distribution (Ueda et al. 2024). Accordingly, a reflective question could be raised if there are large differences between the case information and the dataset. Besides, data is often only a partial representation of the phenomena under investigation (Broussard 2019; Binns et al. 2018), as it is impossible to quantify every aspect of the complexity of social life (Birhane 2021). Questions can help to consider limitations of a dataset and take into account information that is not yet included and is (now) known to contribute to the outcome. So compared to Q2, which is more concerned with the question of what information is relevant for a single recommendation (local), Q3 asks whether the dataset contains relevant information (global).

Q4. Questions can relate to the (apparent) *causal structure* between provided case information and resulted machine recommendation. Again, based on feature contribution it is possible to ask whether and how the machine recommenda-

ID	Element for Reflection (What?)	Description	Useful Information (How?)
Q1	Case Information (<i>Information</i>)	Questions that help further assess, inspect, and contextualise data points to ensure its quality, reliability, and completeness.	Input Data (e.g., tabular patient data), Datasheets (Gebru et al. 2021)
Q2	Relevance of Data (<i>Information, Interpretation & Inference</i>)	Questions that help evaluate whether the data adequately supports the hypothesis / recommendation.	Feature Contribution (e.g., SHAP (Lundberg and Lee 2017), LIME (Ribeiro, Singh, and Guestrin 2016))
Q3	Dataset (<i>Concepts</i>)	Questions that help analyse and assess the data and to ensure it adequately represents the phenomenon (by highlighting limitations or characteristics of the dataset).	Training Data, Datasheets (Gebru et al. 2021), FactSheets (Arnold et al. 2019)
Q4	Causal Structure of Recommendation (<i>Interpretation & Inference</i>)	Questions that help evaluate whether the outcome follows from data to ensure that the causal structure of the model / recommendation is sound.	Feature Contribution (Lundberg and Lee 2017; Ribeiro, Singh, and Guestrin 2016), Counterfactuals (Karimi, Schölkopf, and Valera 2021; Poyiadzi et al. 2020)
Q5*	Alternatives to Recommendation (<i>Question, Purpose</i>)	Questions that help to consider other possibilities ensuring that the larger solution space is considered.	Contextual information (e.g., differential diagnosis or patient preferences (Q7)), Partial Dependence Plot, Counterfactuals
Q6	Assumptions and Expectations of Decision-Maker (<i>Assumptions</i>)	Questions that elicit taken-for-granted assumptions to ensure that the decision-maker is aware of their reasons and possible cognitive biases.	General questions, Cognitive biases, Confidence of decision-maker
Q7*	Stakeholder Preferences (<i>Point of View</i>)	Questions that help the decision-maker enquire and take into account the preferences or needs of the people concerned (e.g., patient), in order to ensure a better or broader understanding of the problem at hand.	General questions, Preference elicitation, Values
Q8*	Consequences of Decision (<i>Implications</i>)	Questions that help elicit anticipation and forward-looking responsibility to ensure that (unintended) consequences and trade-offs are considered and mitigated.	Limitations of Model, FactSheets (Arnold et al. 2019), Model Cards (Mitchell et al. 2019)
Q9*	Change Intervention (<i>Purpose</i>)	Questions that support the exploration of feasible (smaller) interventions to make a desired outcome more likely.	Feature Perturbation, Counterfactuals
Q10	Model Behaviour (<i>Assumptions, Purpose, Concepts</i>)	Questions that help to assess and evaluate the assumptions, rules and thresholds built into the model. In particular, areas of decision limits, i.e. until when does the result remain the same and when does it change.	Decision Boundaries, Feature Perturbation, Counterfactuals, Model Cards (Mitchell et al. 2019), Global Feature Importance (Lundberg et al. 2020)

Table 1: The taxonomy we propose helps to systematically formulate questions that stimulate critical reflection in machine-assisted decision-making. From a taxonomy of Socratic questions we take the elements of thought (Paul 2019), i.e., the *What*, indicated in *italic*, and adapt them to the context of machine-assisted decision-making. To identify information for enriching questions, i.e., the *How*, we use the XAI question bank (Liao, Gruen, and Miller 2020; Liao et al. 2021; Liao and Varshney 2022) and reverse the mapping between questions and explanations. Identifiers (IDs) marked with an asterisk (*) indicate questions that address the level of creating, according to Bloom’s taxonomy (Krathwohl 2002).

tion follows from data point x , for instance “Does diagnosis Y follow from symptom x ?”. So although this question type merely reformulates a causal explanation (Danry et al. 2023), the aim is to help the decision-maker to scrutinise the information.

Q5. Questions can help identify and consider *alternatives to the recommendation*. This calls into question the appropriateness of a recommendation. It might be possible to treat the same condition in different ways, where the feasibility of each option depends on various constraints like available resources or the patient’s preferences (Q7). For example, one purpose might be full recovery, for which a DSS might recommend surgery that involves a long hospital stay; another purpose might be to promote the patient’s well-being, who might prefer to spend as little time as possible in hospital, so that the more appropriate treatment might be a non-surgical intervention. At the same time, as different disease have common symptoms, the physician wants to ensure that another possible diagnosis is not overlooked, e.g., through a differential diagnosis (Mamede and Schmidt 2023).

Q6. Questions can relate to the *assumptions of the decision-maker*. On the one hand, these questions can prevent the decision-maker from jumping to conclusions by helping to uncover tacit knowledge, taken for granted assumptions, and cognitive biases and re-evaluate them case by case. On the other hand, questions can help to increase the confidence and self-efficacy of decision-makers as they allow them to make their own decision.

Q7. Questions can relate to *stakeholder preferences*. If the decision-maker is not the decision-subject, as is the case with physicians and patients, it is important to consider the different perspectives. Ideally, enquiring about the preferences of stakeholders is part of the decision-maker’s routine practice, yet these questions can function as a form of safeguard. If the decision-maker is the decision-subject, questions can help elicit preferences and values of the decision-maker themselves.

Q8. Questions can help anticipate (unintended) *consequences of a decision*. While the decision-maker has ideally weighed up the implications, these questions can promote forward-looking responsibility of the decision-maker. Furthermore, with insights from the dataset (Q3) attention could be drawn to certain limitations of a DSS. Compared to Q3, Q8 focuses on assessing the impact of a recommendation, weighing trade-offs, and supporting the decision-maker in considering appropriate mitigation strategies.

Q9. Questions can relate to *change interventions* in order to further investigate the feasibility of desired alternatives. For such questions, hypothetical scenarios can be created by feature perturbation and counterfactual explanations. Unlike counterfactual explanations, which provide information about what the outcome would have been if the data had been like this, these questions are about whether it is possible to change certain data points to make a particular outcome more likely. For example, a question might help the physician to consider prior and smaller interventions such as smoking cessation, which could increase the expected effectiveness of a desired treatment such as surgery. Although these questions ideally lead to more actionable insights, they

might imply unfeasible or unattainable changes and actions, e.g., changing the age of a patient. These hypothetical scenarios can nevertheless provide insights into inbuilt associations and decision boundaries of the DSS.

Q10. Questions can relate to the *model behaviour*. These questions help to clarify the assumptions built into the DSS, which may be based on the data set (Q3) or other assumptions made during model development. As such, these questions can also provide the decision-maker with a better understanding of the workings and (apparent) causal structure built into the model (Q4). For example, if a DSS makes a decision based on age with a threshold of 50 years, the decision-maker may want to re-evaluate the decision, if the patient is close to that threshold, such as 48 years old.

5 Evaluation

To formatively evaluate our question taxonomy, we presented a first version with sample questions and the final version in two phases to the same orthopaedic surgeon and orthopaedic researcher. In individual online meetings of 30 minutes each, which were not recorded or transcribed, we asked them whether the taxonomy and the derived questions could help practitioners, whether the considerations associated with these questions already correspond to common practice, whether certain questions might be more relevant than others, and whether they could provide example questions or example cases.

Overall, both recognised the relevance of our taxonomy and the use of questions. They were confident that questions could help to re-evaluate initial assumptions of the physician, and counteract habitual priming, e.g., seeing similar yet different patient cases throughout the day. Importantly, as the surgeon noted, this priming also happens through the information presented by a DSS, which questions could help to mitigate. Moreover, they mentioned that the questions derived from the taxonomy could help to treat different patients more objectively by reducing discrepancies between cases, such as different patients describing their pain differently, i.e., pain catastrophising. As one remarked, given that many question types from the taxonomy are derived from data, it is difficult to ask about data points that are not contained in the dataset (Q3).¹ To derive relevant questions, they suggested to talk to experts and ask for their reasoning process, as the usefulness of specific questions, as they noted, depend on the specific use case. Accordingly, they also pointed out that questions should help and not irritate.

5.1 Transferability of Our Taxonomy

To assess the applicability of our taxonomy in another domain, we discussed the finalised version with two post-doctoral researchers in the field of education and technology, and one doctoral researcher who is also a primary school teacher. These individual and informal conversations, each lasting about 30 minutes, were neither recorded nor transcribed. The aim of these conversations was to clarify

¹In this case, general questions like “Were all relevant information considered?” or “Did you consider factor x which is known to influence the outcome, but is not part of the data set?” could help.

whether the taxonomy is relevant for the educational domain, and to familiarise ourselves with some exemplary use cases and questions.

The three researchers shared two examples of so-called learning analytics dashboards (LAD), which educators can use to plan lessons, define learning objectives, track the learning progress of students, create exercises accordingly, and grade students.² These LADs also recommend exercises and learning objectives based on student performance that teachers can use. In view of this, one researcher commented, “Teachers feel the need to be more critical, but do not know how” as they wonder whether they use the LAD correctly. Accordingly, the researchers noted that questions could contribute to the interpretation of the student’s progress shown in the LAD, and help consider factors that go beyond data points, as these systems “do not know everything”. All three researchers therefore confirmed the expected relevance and applicability of our taxonomy for stimulating teachers’ critical reflection in making decisions about students and their futures using recommendations of LADs.

After the three feedback sessions, we created sample questions for the educational domain. We sent these to the three researchers and two of them added some questions and comments that helped us refine our derived questions. We then presented the sample questions, including our taxonomy, in a lab session attended by 12 people, from doctoral researchers to full professors. The research group, which explores the impacts of technologies on learning and teaching practices, recognised the usefulness of our taxonomy. As one participant commented, the sample questions we created are very close to the questions teachers already ask themselves. Nevertheless, the same researcher continued, our taxonomy could be informative to see which questions are covered and which are left out. In addition, another researcher noted that it is very promising that our taxonomy aligns with the questions that teachers ask in relation to algorithmic learning goal recommendations. Furthermore, one attendee noted that students doing the exercises in a LAD could also benefit from questions for critical reflection, such as considering the benefit of doing an exercise that might go against (short-term) preferences.

In comparison to the medical sample questions, the wording of the educational sample questions differs. The structure and elements of critical reflection however remain the same for both. We therefore maintain that our taxonomy is transferable and can be applied to other decision-making domains without adaptation as it builds on general concepts of critical thinking.

5.2 Effectiveness of Questions

Another type of evaluation, beyond the scope of this paper, is to assess the effectiveness of the questions created based on the taxonomy in decision-making, and the extent to which they promote critical reflection and cognitive engagement of the decision-maker. Existing studies, primarily from the educational domain, point to the positive effects of Socratic questioning strategies on moral reasoning in nursing (Tora-

²<https://snappet.org/> and <https://www.gynzy.com/en>.

bizadeh, Hodayuni, and Moattari 2018) and critical thinking skills during (healthcare) education of students (Ishtiaq et al. 2024; Jameel, Gul, and Ishtiaq 2025; Yang, Newby, and Bill 2005; Makhene 2019; Ho, Chen, and Li 2023). These studies demonstrate the validity and robustness of the Socratic questioning approach on which our proposed taxonomy is based, which makes the prospects for our taxonomy promising.

Nevertheless, assessing the effect of questions on decision-making remains an open challenge. Current evaluation methods for decision-making and overreliance on DSS typically only focus on the outcome of a decision, i.e., decision accuracy (Guo et al. 2024; Klingbeil, Grützner, and Schreck 2024). To measure reflection and cognitive engagement, however, it is important to consider the entire decision-making process, including the reasoning process, whether alternatives were considered, the decision-maker’s confidence in the final decision, and satisfaction with the outcome. Existing methods for measuring cognitive engagement are associated with considerable effort and technical barriers, such as electroencephalography (EEG) (Hassib et al. 2017; Berka et al. 2007), or the measurement of pupil dilation (Van Der Wel and Van Steenbergen 2018). Consequently, an accessible evaluation framework would need to be developed in future work, e.g., a self-report scale that takes the entire decision-making process into account. For this, inspiration can be drawn from existing scales for factors associated with cognitive engagement, such as self-efficacy, motivation, and confidence (Lee et al. 2025; Hart, Spiva, and Mareno 2014; Greene 2015), as well as from methods in the educational domain for assessing student’s cognitive engagement with technology (Vongkulluksn et al. 2022).

6 Discussion

We have presented a taxonomy of question types to support the systematic creation of questions that stimulate critical reflection in machine-assisted decision-making. As Figure 1 shows, the questions identified in the XAI question bank for human-centred XAI are very model- and data-centric. As a result, our taxonomy is more data-oriented. Nevertheless, our taxonomy can help to take into account human factors such as assumptions or cognitive biases of the decision-maker and stakeholder preferences as well as consequences of a decision to overcome the current algorithmic focus of XAI (section 2.2). In doing so, questions can stimulate discussion about the machine recommendation and its integration into the broader decision-making process.

6.1 Different Stakeholders

Our taxonomy is primarily targeted at decision-makers who operate DSS to make informed decisions. Accordingly, the aim of our taxonomy is not to formulate questions for the decision-maker to assess the performance of the DSS as this would unfairly shift the responsibility from developers and regulators to operators. The deployed models must be evaluated and tested beforehand, which requires a shift from individual human oversight to institutional oversight (Green 2022). Considering this, other stakeholders could also benefit from critical reflection and thus our taxonomy. *Developers* constructing a model, for example, might ask themselves

what assumptions they are making and whether the dataset used is suitable for solving a given problem. *Auditors* assessing compliance could benefit from our taxonomy when defining objectives for an audit of the model or dataset (Raji et al. 2020; Birhane et al. 2024). Moreover, the taxonomy could be used as a form of supporting AI literacy to help people remain more critical of machine-generated information, especially in view of misinformation (Zavolokina et al. 2024). *Patients* could be supported in identifying (treatment) preferences through reflection, thereby asking relevant questions to their healthcare providers (operating DSS), or being critical while interacting with a medical chatbot themselves. *Students* could be taught to question the relevance of retrieved information and could be helped to remain cognitively engaged in their technology-aided learning. Our proposed taxonomy could therefore benefit various stakeholders involved in the design, development, testing, and use of DSS or machine learning systems.

6.2 Creation of Questions

To facilitate the formulation of questions, we suggested that questions should address the levels of critical thinking according to Bloom’s taxonomy, i.e., evaluating, analysing, and creating. As such, it might be that questions cannot be answered directly and instead lead to further questions (Paul 2019). To motivate further inquiry and engagement of the decision-maker, a brief explanation of why a question should be entertained could be given, as mentioned by the orthopaedic surgeon we spoke to.

It must be considered that questions can be leading and possibly steer the decision-maker towards an undesirable decision. Hence, we advise against the use of large language models, as the generated questions could be irrelevant and unreliable (Weidinger et al. 2022). Instead, researchers or designers ideally work together with experts to create useful questions. When creating the sample questions for the educational domain we realised that it is difficult to do so without having much knowledge about the respective use case, access to data, or the functionality of the DSS or learning analytics dashboard (LAD). Nevertheless, uncovering these blind spots might enable further inquiry into how the system works, and open up a dialogue between developers or service providers and domain experts or affected organisations, such as hospitals or schools.

Furthermore, certain questions might fall into multiple categories of our taxonomy, as different question types can overlap or build on other categories. The causal structure of a recommendation (Q4), for example, can entail questions about the data (Q1&Q3). We believe, however, that this is not a problem as it shows how different elements are connected to each other and it could rather help to formulate questions more precisely so that they can be delineated.

6.3 Presentation of Questions

To present the questions to decision-makers, one of the educational researchers we spoke to suggested that a poster of our taxonomy with sample questions could be displayed in the staff room to sensitise teachers to the different elements of the machine-assisted decision-making process. Alongside

an analogue form and building on the above-mentioned proposals of DSS supporting deliberation and reflection (section 2.3), our taxonomy could inform the design of a technical implementation. As such, we introduced and envision a *reflection machine* (RM) that functions as an additional system or component of a DSS (Haselager et al. 2023; Cornelissen et al. 2022; Fischer 2025). Based on input data, the recommendations, the functioning of the DSS derived from XAI, and the use of our taxonomy relevant questions can be created and presented to the decision-maker accordingly.

Supporting the decision-making process with questions might be an unfamiliar interaction. The design of the interface and the integration into existing workflows therefore influence the acceptance and perceived usefulness of a specific prototype (Lund 2001), and thus the effectiveness of questions on cognitive engagement.

6.4 Cognitive Load

While future work needs to evaluate the effectiveness of questions, questions could interrupt the workflow and increase the cognitive load of the decision-maker. In the aforementioned study, in which explanations were presented at different times as a form of cognitive intervention, participants reported increased cognitive load, which is why they favoured these interventions less (Buçinca, Malaya, and Gajos 2021). Ultimately, however, the cognitive burden is unavoidable and inherent to critical reflection, as Dewey (1933, p.17) notes:

[...] yet thinking need not be reflective. For the person may not be sufficiently critical about the ideas that occur to him. He may jump at a conclusion without weighing the grounds on which it rests; he may forego or unduly shorten the act of hunting, inquiring; he may take the first ‘answer’ or solution, that comes to him because of mental sloth, torpor, impatience to get something settled. One can think reflectively only when one is willing to endure suspense and to undergo the trouble of searching.

Nevertheless, the right balance between effort and effectiveness must be found. Questions must be well integrated into decision-making processes in order to enable the operator to constructively take the posed question into account. The competences, decision-making style, and need for cognition of the decision-maker play a role here.

7 Conclusion

Our taxonomy of 10 question types shifts the emphasis from machine recommendations and explanations to questions and human reflection. Careful selection and presentation of questions has the potential to increase cognitive engagement, mitigate overreliance on DSS and thus support the effective human oversight required by the European AI Act.

A Appendices

For the tables with the sample questions, we refer to: <https://arxiv.org/abs/2504.12830>.

Acknowledgments

We thank Sebastiaan van de Groes, at Radboud UMC, and Miranda van Hooff, at Sint Maartenskliniek and Radboud UMC, for their input on the construction of the taxonomy. Many thanks to Josh Ring and the Adaptive Learning Lab, in particular to Susan Verbunt-Janssen, Susanne de Mooij, and Zowi Vermeire for their valuable insights, and helping with the sample questions for the educational domain. Thanks to Cătălina Răţală and the Decision Science Group for their questions and feedback on a presentation of our paper. Comments from three ACM FAccT reviewers on a previous version, as well as comments from three AAAI/ACM AIES reviewers helped us to improve the paper. This research is funded by the Donders Centre for Cognition.

References

- Anderson, L.; and Krathwohl, D. 2001. *A Taxonomy for Learning, Teaching, and Assessing: A Revision of Bloom's Taxonomy of Educational Objectives*. New York: Longman. ISBN 978-0-8013-1903-7.
- Ang, B. H.; Gollapalli, S. D.; and Ng, S.-K. 2023. Socratic Question Generation: A Novel Dataset, Models, and Evaluation. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, 147–165. Dubrovnik, Croatia: Association for Computational Linguistics.
- Arnold, M.; Bellamy, R. K. E.; Hind, M.; Houde, S.; Mehta, S.; Mojsilovic, A.; Nair, R.; Ramamurthy, K. N.; Reimer, D.; Olteanu, A.; Piorkowski, D.; Tsay, J.; and Varshney, K. R. 2019. FactSheets: Increasing Trust in AI Services through Supplier's Declarations of Conformity. arXiv:1808.07261.
- Bansal, G.; Wu, T.; Zhou, J.; Fok, R.; Nushi, B.; Kamar, E.; Ribeiro, M. T.; and Weld, D. 2021. Does the Whole Exceed Its Parts? The Effect of AI Explanations on Complementary Team Performance. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–16. Yokohama Japan: ACM. ISBN 978-1-4503-8096-6.
- Berka, C.; Levendowski, D. J.; Lumicao, M. N.; Yau, A.; Davis, G.; Zivkovic, V. T.; Olmstead, R. E.; Tremoulet, P. D.; and Craven, P. L. 2007. EEG Correlates of Task Engagement and Mental Workload in Vigilance, Learning, and Memory Tasks. *Aviation, Space, and Environmental Medicine*, 78(5).
- Binns, R.; Van Kleek, M.; Veale, M.; Lyngs, U.; Zhao, J.; and Shadbolt, N. 2018. 'It's Reducing a Human Being to a Percentage': Perceptions of Justice in Algorithmic Decisions. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems*, 1–14. Montreal QC Canada: ACM. ISBN 978-1-4503-5620-6.
- Birhane, A. 2021. The Impossibility of Automating Ambiguity. *Artificial Life*, 27(1): 44–61.
- Birhane, A.; Steed, R.; Ojewale, V.; Vecchione, B.; and Raji, I. D. 2024. AI Auditing: The Broken Bus on the Road to AI Accountability. In *2024 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, 612–643. Toronto, ON, Canada: IEEE. ISBN 979-8-3503-4950-4.
- Blumenfeld, P. C.; Kempler, T. M.; and Krajcik, J. S. 2006. Motivation and Cognitive Engagement in Learning Environments. In *The Cambridge Handbook of: The Learning Sciences.*, 475–488. New York, NY, US: Cambridge University Press. ISBN 0-521-84554-8 (Hardcover); 0-521-60777-9 (Paperback); 978-0521-84554-0 (Hardcover); 978-0-521-60777-3 (Paperback).
- Boud, D.; and Walker, D. 1998. Promoting Reflection in Professional Courses: The Challenge of Context. *Studies in Higher Education*, 23(2): 191–206.
- Broussard, M. 2019. *Artificial Unintelligence: How Computers Misunderstand the World*. Cambridge, Massachusetts London, England: The MIT Press, first mit press paperback edition edition. ISBN 978-0-262-53701-8 978-0-262-03800-3.
- Buçinca, Z.; Malaya, M. B.; and Gajos, K. Z. 2021. To Trust or to Think: Cognitive Forcing Functions Can Reduce Overreliance on AI in AI-assisted Decision-making. *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW1): 1–21.
- Bussone, A.; Stumpf, S.; and O'Sullivan, D. 2015. The Role of Explanations on Trust and Reliance in Clinical Decision Support Systems. In *2015 International Conference on Healthcare Informatics*, 160–169. Dallas, TX, USA: IEEE. ISBN 978-1-4673-9548-9.
- Cabitzza, F.; Campagner, A.; Famiglini, L.; Natali, C.; Caccavella, V.; and Gallazzi, E. 2023. Let Me Think! Investigating the Effect of Explanations Feeding Doubts About the AI Advice. In Holzinger, A.; Kieseberg, P.; Cabitzza, F.; Campagner, A.; Tjoa, A. M.; and Weippl, E., eds., *Machine Learning and Knowledge Extraction*, volume 14065, 155–169. Cham: Springer Nature Switzerland. ISBN 978-3-031-40836-6 978-3-031-40837-3.
- Casey, T. 2014. Reflective Practice in Legal Education: The Stages of Reflection. *Clinical Law Review*, 20(2): 317–354.
- Chiang, C.-W.; Lu, Z.; Li, Z.; and Yin, M. 2024. Enhancing AI-Assisted Group Decision Making through LLM-Powered Devil's Advocate. In *Proceedings of the 29th International Conference on Intelligent User Interfaces*, 103–119. Greenville SC USA: ACM. ISBN 979-8-4007-0508-3.
- Cornelissen, N. A. J.; van Eerd, R. J. M.; Schraffenberger, H. K.; and Haselager, W. F. G. 2022. Reflection Machines: Increasing Meaningful Human Control over Decision Support Systems. *Ethics and Information Technology*, 24(2): 19.
- Cunliffe, A. L. 2002. Reflexive Dialogical Practice in Management Learning. *Management Learning*, 33(1): 35–61.
- Dalim, S. 2022. Promoting Students' Critical Thinking Through Socratic Method: The Views and Challenges. *Asian Journal of University Education*, 18(4): 1034–1047.
- Danry, V.; Pataranutaporn, P.; Mao, Y.; and Maes, P. 2023. Don't Just Tell Me, Ask Me: AI Systems That Intelligently Frame Explanations as Questions Improve Human Logical Discernment Accuracy over Causal AI Explanations. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, CHI '23, 1–13. Hamburg Germany: ACM. ISBN 978-1-4503-9421-5.

- Dewey, J. 1933. *How We Think: A Restatement Of The Relation Of Reflective Thinking to the Educative Process*. Lexington, Massachusetts: D.C. Heath and Company, 2 edition.
- Dratsch, T.; Chen, X.; Rezazade Mehrizi, M.; Kloeckner, R.; Mähringer-Kunz, A.; Püsken, M.; Baeßler, B.; Sauer, S.; Maintz, D.; and Pinto dos Santos, D. 2023. Automation Bias in Mammography: The Impact of Artificial Intelligence BI-RADS Suggestions on Reader Performance. *Radiology*, 307(4): e222176.
- Ehsan, U.; Liao, Q. V.; Muller, M.; Riedl, M. O.; and 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*, 1–19. Yokohama Japan: ACM. ISBN 978-1-4503-8096-6.
- Ehsan, U.; Passi, S.; Liao, Q. V.; Chan, L.; Lee, I.-H.; Muller, M.; and Riedl, M. O. 2024. The Who in XAI: How AI Background Shapes Perceptions of AI Explanations. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–32. Honolulu HI USA: ACM. ISBN 979-8-4007-0330-0.
- Ehsan, U.; and Riedl, M. O. 2024. Explainable AI Reloaded: Challenging the XAI Status Quo in the Era of Large Language Models. In *Proceedings of the Halfway to the Future Symposium*, HttF '24, 8. Santa Cruz, CA, USA: Association for Computing Machinery.
- European Parliament. 2024. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying down Harmonised Rules on Artificial Intelligence and Amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2020/1828 (Artificial Intelligence Act) (Text with EEA Relevance).
- Favero, L.; Pérez-Ortiz, J. A.; Käser, T.; and Oliver, N. 2024. Enhancing Critical Thinking in Education by Means of a Socratic Chatbot. arXiv:2409.05511.
- Fischer, S. W. 2025. A Reflection Machine to Support Critical Reflection During Decision-Making. In *Companion Proceedings of the 30th International Conference on Intelligent User Interfaces*, 199–201. Cagliari Italy: ACM. ISBN 979-8-4007-1409-2.
- Forehand, M. 2010. Blooms Taxonomy. *Emerging perspectives on learning, teaching, and technology*, 41(4): 47–56.
- Gajos, K. Z.; and Mamykina, L. 2022. Do People Engage Cognitively with AI? Impact of AI Assistance on Incidental Learning. In *27th International Conference on Intelligent User Interfaces*, 794–806. Helsinki Finland: ACM. ISBN 978-1-4503-9144-3.
- Geburu, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for Datasets. *Communications of the ACM*, 64(12): 86–92.
- Ghanizadeh, A. 2017. The Interplay between Reflective Thinking, Critical Thinking, Self-Monitoring, and Academic Achievement in Higher Education. *Higher Education*, 74(1): 101–114.
- Ghassemi, M.; Oakden-Rayner, L.; and Beam, A. L. 2021. The False Hope of Current Approaches to Explainable Artificial Intelligence in Health Care. *The Lancet Digital Health*, 3(11): e745–e750.
- Green, B. 2022. The Flaws of Policies Requiring Human Oversight of Government Algorithms. *Computer Law & Security Review*, 45: 105681.
- Greene, B. A. 2015. Measuring Cognitive Engagement With Self-Report Scales: Reflections From Over 20 Years of Research. *Educational Psychologist*, 50(1): 14–30.
- Guo, Z.; Wu, Y.; Hartline, J. D.; and Hullman, J. 2024. A Decision Theoretic Framework for Measuring AI Reliance. In *The 2024 ACM Conference on Fairness, Accountability, and Transparency*, 221–236. Rio de Janeiro Brazil: ACM. ISBN 979-8-4007-0450-5.
- Hart, P. L.; Spiva, L.; and Mareno, N. 2014. Psychometric Properties of the Clinical Decision-Making Self-Confidence Scale. *Journal of Nursing Measurement*, 22(2): 312–322.
- Haselager, P.; Schraffenberger, H.; Thill, S.; Fischer, S.; Lanillos, P.; van de Groes, S.; and van Hooff, M. 2023. Reflection Machines: Supporting Effective Human Oversight Over Medical Decision Support Systems. *Cambridge Quarterly of Healthcare Ethics*, 33(3): 380–389.
- Hassib, M.; Khamis, M.; Friedl, S.; Schneegass, S.; and Alt, F. 2017. Brainatwork: Logging Cognitive Engagement and Tasks in the Workplace Using Electroencephalography. In *Proceedings of the 16th International Conference on Mobile and Ubiquitous Multimedia*, 305–310. Stuttgart Germany: ACM. ISBN 978-1-4503-5378-6.
- Hess, B. J.; Lipner, R. S.; Thompson, V.; Holmboe, E. S.; and Graber, M. L. 2015. Blink or Think: Can Further Reflection Improve Initial Diagnostic Impressions? *Academic Medicine*, 90(1): 112–118.
- Ho, Y.-R.; Chen, B.-Y.; and Li, C.-M. 2023. Thinking More Wisely: Using the Socratic Method to Develop Critical Thinking Skills amongst Healthcare Students. *BMC Medical Education*, 23(1): 173.
- Ishtiaq, M.; Gul, R.; Holt, J.; and Babar, M. 2024. Questioning Training and Critical Thinking of Undergraduate Students of Health and Social Sciences: A Scoping Review. *Teaching and Learning in Nursing*, 19(4): e602–e609.
- Jacobs, M.; Pradier, M. F.; McCoy, T. H.; Perlis, R. H.; Doshi-Velez, F.; and Gajos, K. Z. 2021. How Machine-Learning Recommendations Influence Clinician Treatment Selections: The Example of Antidepressant Selection. *Translational Psychiatry*, 11(1): 108.
- Jameel, S.; Gul, R.; and Ishtiaq, M. 2025. The Complexity of Questioning Strategy in Clinical Teaching from the Perspective of Nurse Educators. *Teaching and Learning in Nursing*, 20(2): e520–e527.
- Karimi, A.-H.; Schölkopf, B.; and Valera, I. 2021. Algorithmic Recourse: From Counterfactual Explanations to Interventions. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 353–362. Virtual Event Canada: ACM. ISBN 978-1-4503-8309-7.

- Khoshgoftar, Z.; and Barkhordari-Sharifabad, M. 2023. Medical Students' Reflective Capacity and Its Role in Their Critical Thinking Disposition. *BMC Medical Education*, 23(1): 198.
- Klingbeil, A.; Grützner, C.; and Schreck, P. 2024. Trust and Reliance on AI — An Experimental Study on the Extent and Costs of Overreliance on AI. *Computers in Human Behavior*, 160: 108352.
- Krathwohl, D. R. 2002. A Revision of Bloom's Taxonomy: An Overview. *Theory Into Practice*, 41(4): 212–218.
- Lambe, K. A.; O'Reilly, G.; Kelly, B. D.; and Curristan, S. 2016. Dual-Process Cognitive Interventions to Enhance Diagnostic Reasoning: A Systematic Review. *BMJ Quality & Safety*, 25(10): 808–820.
- Lara, F. 2021. Why a Virtual Assistant for Moral Enhancement When We Could Have a Socrates? *Science and Engineering Ethics*, 27(4): 42.
- Lara, F.; and Deckers, J. 2020. Artificial Intelligence as a Socratic Assistant for Moral Enhancement. *Neuroethics*, 13(3): 275–287.
- Lee, H.-P. H.; Sarkar, A.; Tankelevitch, L.; Drosos, I.; Rintel, S.; Banks, R.; and Wilson, N. 2025. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–22. Yokohama Japan: ACM. ISBN 979-8-4007-1394-1.
- Li, F.; Wu, P.; Ong, H. H.; Peterson, J. F.; Wei, W.-Q.; and Zhao, J. 2023. Evaluating and Mitigating Bias in Machine Learning Models for Cardiovascular Disease Prediction. *Journal of Biomedical Informatics*, 138: 104294.
- Liao, Q. V.; Gruen, D.; and 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, 1–15. Honolulu HI USA: ACM. ISBN 978-1-4503-6708-0.
- Liao, Q. V.; Pribić, M.; Han, J.; Miller, S.; and Sow, D. 2021. Question-Driven Design Process for Explainable AI User Experiences. arXiv:2104.03483.
- Liao, Q. V.; and Varshney, K. R. 2022. Human-Centered Explainable AI (XAI): From Algorithms to User Experiences. arXiv:2110.10790.
- Liu, J.; Huang, Z.; Xiao, T.; Sha, J.; Wu, J.; Liu, Q.; Wang, S.; and Chen, E. 2024. SocraticLM: Exploring Socratic Personalized Teaching with Large Language Models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, volume 37, 85693–85721. Curran Associates, Inc.
- Longo, L.; Brcic, M.; Cabitza, F.; Choi, J.; Confalonieri, R.; Ser, J. D.; Guidotti, R.; Hayashi, Y.; Herrera, F.; Holzinger, A.; Jiang, R.; Khosravi, H.; Lecue, F.; Malgieri, G.; Páez, A.; Samek, W.; Schneider, J.; Speith, T.; and Stumpf, S. 2024. Explainable Artificial Intelligence (XAI) 2.0: A Manifesto of Open Challenges and Interdisciplinary Research Directions. *Information Fusion*, 106: 102301.
- Loughran, J. J. 1996. *Developing Reflective Practice: Learning about Teaching and Learning Through Modelling*. London: Falmer Press.
- Lund, A. M. 2001. Measuring Usability with the USE Questionnaire. *Usability Interface*, 8(2): 3–6.
- Lundberg, S. M.; Erion, G.; Chen, H.; DeGrave, A.; Prutkin, J. M.; Nair, B.; Katz, R.; Himmelfarb, J.; Bansal, N.; and Lee, S.-I. 2020. From Local Explanations to Global Understanding with Explainable AI for Trees. *Nature Machine Intelligence*, 2(1): 56–67.
- Lundberg, S. M.; and Lee, S.-I. 2017. A Unified Approach to Interpreting Model Predictions. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Ma, S.; Chen, Q.; Wang, X.; Zheng, C.; Peng, Z.; Yin, M.; and Ma, X. 2025. Towards Human-AI Deliberation: Design and Evaluation of LLM-Empowered Deliberative AI for AI-Assisted Decision-Making. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, 1–23. Yokohama Japan: ACM.
- Ma, S.; Zhang, C.; Wang, X.; Ma, X.; and Yin, M. 2024. Beyond Recommender: An Exploratory Study of the Effects of Different AI Roles in AI-Assisted Decision Making. arXiv:2403.01791.
- Makhene, A. 2019. The Use of the Socratic Inquiry to Facilitate Critical Thinking in Nursing Education. *Health SA Gesondheid*, 24.
- Mamede, S.; and Schmidt, H. G. 2023. Deliberate Reflection and Clinical Reasoning: Founding Ideas and Empirical Findings. *Medical Education*, 57(1): 76–85.
- Mamede, S.; Schmidt, H. G.; and Penaforte, J. C. 2008. Effects of Reflective Practice on the Accuracy of Medical Diagnoses. *Medical Education*, 42(5): 468–475.
- Mann, K.; Gordon, J.; and MacLeod, A. 2009. Reflection and Reflective Practice in Health Professions Education: A Systematic Review. *Advances in Health Sciences Education*, 14(4): 595–621.
- Mezirow, J. 1990. *Fostering Critical Reflection in Adulthood: A Guide to Transformative and Emancipatory Learning*. The Jossey-Bass Higher Education Series. San Francisco: Jossey-Bass Publishers. ISBN 978-1-55542-207-3.
- Miller, T. 2023. Explainable AI Is Dead, Long Live Explainable AI!: Hypothesis-driven Decision Support Using Evaluative AI. In *2023 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '23, 333–342. Chicago IL USA: ACM. ISBN 979-8-4007-0192-4.
- Mitchell, M.; Wu, S.; Zaldivar, A.; Barnes, P.; Vasserman, L.; Hutchinson, B.; Spitzer, E.; Raji, I. D.; and Gebru, T. 2019. Model Cards for Model Reporting. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, 220–229. Atlanta GA USA: ACM. ISBN 978-1-4503-6125-5.
- Osmond, J.; and Darlington, Y. 2005. Reflective Analysis: Techniques for Facilitating Reflection. *Australian Social Work*, 58(1): 3–14.

- Passi, S.; and Vorvoreanu, M. 2022. Overreliance on AI: Literature Review. Technical Report MSR-TR-2022-12, Microsoft.
- Paul, R. 2010. *The Art of Asking Essential Questions: Based on Critical Thinking Concepts and Socratic Principles*. Thinker's Guide Library. Blue Ridge Summit: Foundation for Critical Thinking, 5th ed edition. ISBN 978-0-944583-16-6 978-1-5381-3380-4.
- Paul, R. 2019. *The Thinker's Guide to Socratic Questioning*. Thinker's Guide Library. Blue Ridge Summit: Rowman & Littlefield Publishers. ISBN 978-0-944583-31-9 978-1-5381-3381-1.
- Poyiadzi, R.; Sokol, K.; Santos-Rodriguez, R.; De Bie, T.; and Flach, P. 2020. FACE: Feasible and Actionable Counterfactual Explanations. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, 344–350. New York NY USA: ACM. ISBN 978-1-4503-7110-0.
- Prakash, S.; Sladek, R. M.; and Schuwirth, L. 2019. Interventions to Improve Diagnostic Decision Making: A Systematic Review and Meta-Analysis on Reflective Strategies. *Medical Teacher*, 41(5): 517–524.
- Raji, I. D.; Smart, A.; White, R. N.; Mitchell, M.; Gebru, T.; Hutchinson, B.; Smith-Loud, J.; Theron, D.; and Barnes, P. 2020. Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*, 33–44. Barcelona Spain: ACM. ISBN 978-1-4503-6936-7.
- Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. "Why Should I Trust You?": Explaining the Predictions of Any Classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 1135–1144. San Francisco California USA: ACM. ISBN 978-1-4503-4232-2.
- Rogers, R. R. 2001. Reflection in Higher Education: A Concept Analysis. *Innovative Higher Education*, 26(1): 37–57.
- Rotgans, J. I.; and Schmidt, H. G. 2011. Cognitive Engagement in the Problem-Based Learning Classroom. *Advances in Health Sciences Education*, 16(4): 465–479.
- Ryan, R. M.; and Deci, E. L. 2000. Self-Determination Theory and the Facilitation of Intrinsic Motivation, Social Development, and Well-Being. *American Psychologist*, 55(1): 68–78.
- Saarela, M.; and Jauhiainen, S. 2021. Comparison of Feature Importance Measures as Explanations for Classification Models. *SN Applied Sciences*, 3(2): 272.
- Sarkar, A. 2024. AI Should Challenge, Not Obey. *Communications of the ACM*, 67(10): 18–21.
- Schmidt, H. G.; and Mamede, S. 2023. Improving Diagnostic Decision Support through Deliberate Reflection: A Proposal. *Diagnosis*, 10(1): 38–42.
- Schön, D. A. 1983. *The Reflective Practitioner: How Professionals Think in Action*. New York: Basic Books. ISBN 978-0-465-06874-6 978-0-465-06878-4.
- Schweiger, D. M.; Sandberg, W. R.; and Rechner, P. L. 1989. Experiential Effects of Dialectical Inquiry, Devil's Advocacy, and Consensus Approaches to Strategic Decision Making. *Academy of Management Journal*, 32(4): 745–772.
- Schwenk, C. R. 1984. DEVIL'S ADVOCACY IN MANAGERIAL DECISION-MAKING. *Journal of Management Studies*, 21(2): 153–168.
- Singh, M.; James, P.; Paul, H.; and Bolar, K. 2022. Impact of Cognitive-Behavioral Motivation on Student Engagement. *Heliyon*, 8(7): e09843.
- Torabizadeh, C.; Homayuni, L.; and Moattari, M. 2018. Impacts of Socratic Questioning on Moral Reasoning of Nursing Students. *Nursing Ethics*, 25(2): 174–185.
- Tschandl, P.; Rinner, C.; Apalla, Z.; Argenziano, G.; Codella, N.; Halpern, A.; Janda, M.; Lallas, A.; Longo, C.; Malvehy, J.; Paoli, J.; Puig, S.; Rosendahl, C.; Soyer, H. P.; Zalaudek, I.; and Kittler, H. 2020. Human-Computer Collaboration for Skin Cancer Recognition. *Nature Medicine*, 26(8): 1229–1234.
- Ueda, D.; Kakinuma, T.; Fujita, S.; Kamagata, K.; Fushimi, Y.; Ito, R.; Matsui, Y.; Nozaki, T.; Nakaura, T.; Fujima, N.; Tatsugami, F.; Yanagawa, M.; Hirata, K.; Yamada, A.; Tsuboyama, T.; Kawamura, M.; Fujioka, T.; and Naganawa, S. 2024. Fairness of Artificial Intelligence in Healthcare: Review and Recommendations. *Japanese Journal of Radiology*, 42(1): 3–15.
- Vaccaro, M.; Almaatouq, A.; and Malone, T. 2024. When Combinations of Humans and AI Are Useful: A Systematic Review and Meta-Analysis. *Nature Human Behaviour*, 8(12): 2293–2303.
- van Baalen, S.; and Boon, M. 2015. An Epistemological Shift: From Evidence-Based Medicine to Epistemological Responsibility: From EBM to Epistemological Responsibility. *Journal of Evaluation in Clinical Practice*, 21(3): 433–439.
- Van Der Wel, P.; and Van Steenbergen, H. 2018. Pupil Dilation as an Index of Effort in Cognitive Control Tasks: A Review. *Psychonomic Bulletin & Review*, 25(6): 2005–2015.
- Vasconcelos, H.; Jörke, M.; Grunde-McLaughlin, M.; Gerstenberg, T.; Bernstein, M. S.; and Krishna, R. 2023. Explanations Can Reduce Overreliance on AI Systems During Decision-Making. *Proceedings of the ACM on Human-Computer Interaction*, 7(CSCW1): 1–38.
- Vered, M.; Livni, T.; Howe, P. D. L.; Miller, T.; and Sonenberg, L. 2023. The Effects of Explanations on Automation Bias. *Artificial Intelligence*, 322: 103952.
- Vongkulluksn, V. W.; Lu, L.; Nelson, M. J.; and Xie, K. 2022. Cognitive Engagement with Technology Scale: A Validation Study. *Educational technology research and development*, 70(2): 419–445.
- Wachter, S.; Mittelstadt, B.; and Russell, C. 2017. Counterfactual Explanations Without Opening the Black Box: Automated Decisions and the GDPR. *Harvard Journal of Law & Technology*, 31(2).
- Walger, C.; Roglio, K. D. D.; and Abib, G. 2016. HR Managers' Decision-Making Processes: A "Reflective Practice" Analysis. *Management Research Review*, 39(6): 655–671.

Weidinger, L.; Uesato, J.; Rauh, M.; Griffin, C.; Huang, P.-S.; Mellor, J.; Glaese, A.; Cheng, M.; Balle, B.; Kasirzadeh, A.; Biles, C.; Brown, S.; Kenton, Z.; Hawkins, W.; Stepleton, T.; Birhane, A.; Hendricks, L. A.; Rimell, L.; Isaac, W.; Haas, J.; Legassick, S.; Irving, G.; and Gabriel, I. 2022. Taxonomy of Risks Posed by Language Models. In *2022 ACM Conference on Fairness, Accountability, and Transparency*, 214–229. Seoul Republic of Korea: ACM. ISBN 978-1-4503-9352-2.

Yang, Y.-T. C.; Newby, T. J.; and Bill, R. L. 2005. Using Socratic Questioning to Promote Critical Thinking Skills Through Asynchronous Discussion Forums in Distance Learning Environments. *American Journal of Distance Education*, 19(3): 163–181.

Zavolokina, L.; Sprenkamp, K.; Katashinskaya, Z.; Jones, D. G.; and Schwabe, G. 2024. Think Fast, Think Slow, Think Critical: Designing an Automated Propaganda Detection Tool. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*, 1–24. Honolulu HI USA: ACM. ISBN 979-8-4007-0330-0.

Zednik, C. 2019. Solving the Black Box Problem: A Normative Framework for Explainable Artificial Intelligence. *Philosophy & Technology*, 34: 265–288.

Zhang, T.; Yang, X. J.; and Li, B. 2023. May I Ask a Follow-up Question? Understanding the Benefits of Conversations in Neural Network Explainability. arXiv:2309.13965.